Introduction to Modern Economic Growth: Parts 1-4

Daron Acemoglu

Department of Economics,

Massachusetts Institute of Technology

Contents

Contents

Preface xi

Part 1. Introduction 1

Chapter 1. Economic Growth and Economic Development: The Questions 3
1.1. Cross-Country Income Differences 3
1.2. Income and Welfare 6
1.3. Economic Growth and Income Differences 9
1.4. Origins of Today's Income Differences and World Economic Growth 12
1.5. Conditional Convergence 16
1.6. Correlates of Economic Growth 18
1.7. From Correlates to Fundamental Causes 21
1.8. The Agenda 24
1.9. References and Literature 26

Chapter 2. The Solow Growth Model 31
2.1. The Economic Environment of the Basic Solow Model 32
2.2. The Solow Model in Discrete Time 40
2.3. Transitional Dynamics in the Discrete Time Solow Model 50
2.4. The Solow Model in Continuous Time 54
2.5. Transitional Dynamics in the Continuous Time Solow Model 58
2.6. A First Look at Sustained Growth 62
2.7. Solow Model with Technological Progress 63
2.8. Comparative Dynamics 74
2.9. Taking Stock 75
2.10. References and Literature 76
2.11. Exercises 77

Chapter 3. The Solow Model and the Data 83
3.1. Growth Accounting 83
3.2. Solow Model and Regression Analyses 86
3.3. The Solow Model with Human Capital 93
3.4. Solow Model and Cross-Country Income Differences: Regression Analyses 98
3.5. Calibrating Productivity Differences 106
3.6. Estimating Productivity Differences 111
3.7. Taking Stock 116
3.8. References and Literature 118
3.9. Exercises 119

Chapter 4. Fundamental Determinants of Differences in Economic Performance 123

4.1. Proximate Versus Fundamental Causes 123
4.2. Economies of Scale, Population, Technology and World Growth 127
4.3. The Four Fundamental Causes 129
4.4. The Effect of Institutions on Economic Growth 139
4.5. What Types of Institutions? 153
4.6. Disease and Development 155
4.7. Political Economy of Institutions: First Thoughts 158
4.8. Taking Stock 159
4.9. References and Literature 159
4.10. Exercises 162

Part 2. Towards Neoclassical Growth 165

Chapter 5. Foundations of Neoclassical Growth 167
5.1. Preliminaries 167
5.2. The Representative Household 169
5.3. Infinite Planning Horizon 175
5.4. The Representative Firm 178
5.5. Problem Formulation 180
5.6. Welfare Theorems 181
5.7. Proof of the Second Welfare Theorem, Theorem 5.7* 188
5.8. Sequential Trading 190
5.9. Optimal Growth 194
5.10. Taking Stock 195
5.11. References and Literature 196
5.12. Exercises 197

Chapter 6. Infinite-Horizon Optimization and Dynamic Programming 203
6.1. Discrete-Time Infinite-Horizon Optimization 203
6.2. Introduction to Stationary Dynamic Programming 206
6.3. Stationary Dynamic Programming Theorems 208
6.4. The Contraction Mapping Theorem and Applications* 212
6.5. Proofs of the Main Dynamic Programming Theorems* 217
6.6. Fundamentals of Stationary Dynamic Programming 224
6.7. Nonstationary Infinite-Horizon Optimization 235
6.8. Optimal Growth in Discrete Time 239
6.9. Competitive Equilibrium Growth 244
6.10. Computation 245
6.11. Taking Stock 246
6.12. References and Literature 246
6.13. Exercises 248

Chapter 7. Review of the Theory of Optimal Control 253
7.1. Variational Arguments 254
7.2. The Maximum Principle: A First Look 262
7.3. Infinite-Horizon Optimal Control 267
7.4. More on Transversality Conditions 278
7.5. Discounted Infinite-Horizon Optimal Control 281
7.6. Existence of Solutions, Concavity and Differentiability* 288

7.7. A First Look at Optimal Growth in Continuous Time 296
7.8. The q-Theory of Investment and Saddle-Path Stability 298
7.9. Taking Stock 304
7.10. References and Literature 305
7.11. Exercises 308

Part 3. Neoclassical Growth 315

Chapter 8. The Neoclassical Growth Model 317
8.1. Preferences, Technology and Demographics 317
8.2. Characterization of Equilibrium 322
8.3. Optimal Growth 327
8.4. Steady-State Equilibrium 328
8.5. Transitional Dynamics 330
8.6. Neoclassical Growth in Discrete Time 333
8.7. Technological Change and the Canonical Neoclassical Model 335
8.8. The Role of Policy 341
8.9. Comparative Dynamics 342
8.10. A Quantitative Evaluation 344
8.11. Extensions 346
8.12. Taking Stock 347
8.13. References and Literature 348
8.14. Exercises 349

Chapter 9. Growth with Overlapping Generations 359
9.1. Problems of Infinity 359
9.2. The Baseline Overlapping Generations Model 361
9.3. The Canonical Overlapping Generations Model 366
9.4. Overaccumulation and Pareto Optimality of Competitive Equilibrium in the Overlapping Generations Model 368
9.5. Role of Social Security in Capital Accumulation 371
9.6. Overlapping Generations with Impure Altruism 373
9.7. Overlapping Generations with Perpetual Youth 377
9.8. Overlapping Generations in Continuous Time 380
9.9. Taking Stock 386
9.10. References and Literature 387
9.11. Exercises 388

Chapter 10. Human Capital and Economic Growth 393
10.1. A Simple Separation Theorem 393
10.2. Schooling Investments and Returns to Education 395
10.3. The Ben-Porath Model 397
10.4. Neoclassical Growth with Physical and Human Capital 401
10.5. Capital-Skill Complementarity in an Overlapping Generations Model 406
10.6. Physical and Human Capital with Imperfect Labor Markets 409
10.7. Human Capital Externalities 415
10.8. The Nelson-Phelps Model of Human Capital 417
10.9. Taking Stock 419
10.10. References and Literature 420

10.11. Exercises 422

Chapter 11. First-Generation Models of Endogenous Growth 425
11.1. The AK Model Revisited 426
11.2. The AK Model with Physical and Human Capital 431
11.3. The Two-Sector AK Model 433
11.4. Growth with Externalities 437
11.5. Taking Stock 441
11.6. References and Literature 443
11.7. Exercises 443

Part 4. Endogenous Technological Change 449

Chapter 12. Modeling Technological Change 451
12.1. Different Conceptions of Technology 451
12.2. Science and Profits 455
12.3. The Value of Innovation in Partial Equilibrium 457
12.4. The Dixit-Stiglitz Model and "Aggregate Demand Externalities" 464
12.5. Individual R&D Uncertainty and the Stock Market 471
12.6. Taking Stock 472
12.7. References and Literature 473
12.8. Exercises 474

Chapter 13. Expanding Variety Models 479
13.1. The Lab-Equipment Model of Growth with Input Varieties 479
13.2. Growth with Knowledge Spillovers 491
13.3. Growth without Scale Effects 493
13.4. Growth with Expanding Product Varieties 496
13.5. Taking Stock 500
13.6. References and Literature 501
13.7. Exercises 502

Chapter 14. Models of Schumpeterian Growth 509
14.1. A Baseline Model of Schumpeterian Growth 510
14.2. A One-Sector Schumpeterian Growth Model 519
14.3. Innovation by Incumbents and Entrants and Sources of Productivity Growth 524
14.4. Step-by-Step Innovations* 536
14.5. Taking Stock 548
14.6. References and Literature 549
14.7. Exercises 551

Chapter 15. Directed Technological Change 559
15.1. Importance of Biased Technological Change 559
15.2. Basics and Definitions 563
15.3. Baseline Model of Directed Technological Change 566
15.4. Directed Technological Change with Knowledge Spillovers 579
15.5. Directed Technological Change without Scale Effects 583
15.6. Endogenous Labor-Augmenting Technological Change 585
15.7. Generalizations and Other Applications 588
15.8. An Alternative Approach to Labor-Augmenting Technological Change * 589

15.9. Taking Stock 594
15.10. References and Literature 595
15.11. Exercises 598

Part 5. Stochastic Growth 605

Chapter 16. Stochastic Dynamic Programming 607
16.1. Dynamic Programming with Expectations 607
16.2. Proofs of the Stochastic Dynamic Programming Theorems* 614
16.3. Stochastic Euler Equations 620
16.4. Generalization to Markov Processes* 622
16.5. Applications of Stochastic Dynamic Programming 624
16.6. Taking Stock 632
16.7. References and Literature 633
16.8. Exercises 634

Chapter 17. Stochastic Growth Models 639
17.1. The Brock-Mirman Model 640
17.2. Equilibrium Growth under Uncertainty 645
17.3. Application: Real Business Cycle Models 654
17.4. Growth with Incomplete Markets: The Bewley Model 657
17.5. The Overlapping Generations Model with Uncertainty 661
17.6. Risk, Diversification and Growth 663
17.7. Taking Stock 681
17.8. References and Literature 682
17.9. Exercises 683

Part 6. Technology Diffusion, Trade and Interdependences 689

Chapter 18. Diffusion of Technology 693
18.1. Productivity Differences and Technology 693
18.2. A Benchmark Model of Technology Diffusion 696
18.3. Technology Diffusion and Endogenous Growth 703
18.4. Appropriate and Inappropriate Technologies and Productivity Differences 708
18.5. Contracting Institutions and Technology Adoption 716
18.6. Taking Stock 729
18.7. References and Literature 731
18.8. Exercises 732

Chapter 19. Trade and Growth 739
19.1. Growth and Financial Capital Flows 739
19.2. Why Doesn't Capital Flow from Rich to Poor Countries? 745
19.3. Economic Growth in a Heckscher-Ohlin World 747
19.4. Trade, Specialization and the World Income Distribution 757
19.5. Trade, Technology Diffusion and the Product Cycle 769
19.6. Trade and Endogenous Technological Change 774
19.7. Learning-by-Doing, Trade and Growth 777
19.8. Taking Stock 781
19.9. References and Literature 783
19.10. Exercises 785

Part 7. Economic Development and Economic Growth 791

Chapter 20. Structural Change and Economic Growth 797
20.1. Non-Balanced Growth: The Demand Side 797
20.2. Non-Balanced Growth: The Supply Side 805
20.3. Agricultural Productivity and Industrialization 818
20.4. Taking Stock 824
20.5. References and Literature 825
20.6. Exercises 826

Chapter 21. Structural Transformations and Market Failures in Development 831
21.1. Financial Development 833
21.2. Fertility, Mortality and the Demographic Transition 838
21.3. Migration, Urbanization and The Dual Economy 846
21.4. Distance to the Frontier and Changes in the Organization of Production 856
21.5. Multiple Equilibria From Aggregate Demand Externalities and the Big Push 865
21.6. Inequality, Credit Market Imperfections and Human Capital 872
21.7. Towards a Unified Theory of Development and Growth? 885
21.8. Taking Stock 890
21.9. References and Literature 891
21.10. Exercises 894

Part 8. Political Economy of Growth 901

Chapter 22. Institutions, Political Economy and Growth 907
22.1. The Impact of Institutions on Long-Run Development 908
22.2. Distributional Conflict and Economic Growth in a Simple Society 913
22.3. Distributional Conflict and Competition 925
22.4. Inefficient Economic Institutions: A First Pass 938
22.5. Distributional Conflict and Economic Growth: Concave Preferences* 942
22.6. Heterogeneous Preferences, Social Choice and the Median Voter* 949
22.7. Distributional Conflict and Economic Growth: Heterogeneity and the Median Voter 968
22.8. The Provision of Public Goods: Weak Versus Strong States 973
22.9. Taking Stock 979
22.10. References and Literature 982
22.11. Exercises 985

Chapter 23. Political Institutions and Economic Growth 993
23.1. Political Regimes and Economic Growth 994
23.2. Political Institutions and Growth-Enhancing Policies 999
23.3. Dynamic Tradeoffs 1003
23.4. Understanding Endogenous Political Change 1021
23.5. Dynamics of Political and Economic Institutions: A First Look 1032
23.6. Taking Stock 1044
23.7. References and Literature 1046
23.8. Exercises 1047

Chapter 24. Epilogue: Mechanics and Causes of Economic Growth 1053
24.1. What Have We Learned? 1053

24.2. A Possible Perspective on Growth and Stagnation over the Past 200 Years 1057
24.3. Many Remaining Questions 1067

Part 9. Mathematical Appendices 1071

Chapter A. Odds and Ends in Real Analysis and Applications to Optimization 1073
A.1. Distances and Metric Spaces 1073
A.2. Mappings, Functions, Sequences, Nets and Continuity 1077
A.3. A Minimal Amount of Topology: Continuity and Compactness* 1082
A.4. The Product Topology* 1088
A.5. Absolute Continuity and Equicontinuity* 1091
A.6. Correspondences and Berge's Maximum Theorem 1094
A.7. Convexity, Concavity, Quasi-Concavity and Fixed Points 1098
A.8. Differentiation, Taylor Series and the Mean Value Theorem 1101
A.9. Functions of Several Variables and the Inverse and Implicit Function Theorems 1105
A.10. Separation Theorems* 1109
A.11. Constrained Optimization 1113
A.12. Exercises 1118

Chapter B. Review of Ordinary Differential Equations 1121
B.1. Review of Eigenvalues and Eigenvectors 1121
B.2. Some Basic Results on Integrals 1122
B.3. Linear Differential Equations 1124
B.4. Solutions to Linear First-Order Differential Equations 1125
B.5. Systems of Linear Differential Equations 1128
B.6. Stability for Nonlinear Differential Equations 1130
B.7. Separable and Exact Differential Equations 1131
B.8. Existence and Uniqueness of Solutions 1133
B.9. Continuity and Differentiability of Solutions 1135
B.10. Difference Equations 1135
B.11. Exercises 1138

Chapter C. Brief Review of Dynamic Games 1139
C.1. Basic Definitions 1139
C.2. Some Basic Results 1143
C.3. Application: Repeated Games With Perfect Observability 1147
C.4. Exercises 1148

Chapter D. List of Theorems 1151
Chapter 2 1151
Chapter 5 1151
Chapter 6 1151
Chapter 7 1152
Chapter 10 1152
Chapter 16 1152
Chapter 22 1153
Appendix Chapter A 1153
Appendix Chapter B 1154
Appendix Chapter C 1154

References (incomplete) 1157

Preface

This book is intended to serve two purposes:

(1) First and foremost, this is a book about economic growth and long-run economic development. The process of economic growth and the sources of differences in economic performance across nations are some of the most interesting, important and challenging areas in modern social science. The primary purpose of this book is to introduce graduate students to these major questions and to the theoretical tools necessary for studying them. The book therefore strives to provide students with a strong background in dynamic economic analysis, since only such a background will enable a serious study of economic growth and economic development. It also tries to provide a clear discussion of the broad empirical patterns and historical processes underlying the current state of the world economy. This is motivated by my belief that to understand why some countries grow and some fail to do so, economists have to move beyond the mechanics of models and pose questions about the fundamental causes of economic growth.

(2) In a somewhat different capacity, this book is also a graduate-level introduction to modern macroeconomics and dynamic economic analysis. It is sometimes commented that, unlike basic microeconomic theory, there is no core of current macroeconomic theory that is shared by all economists. This is not entirely true. While there is disagreement among macroeconomists about how to approach short-run macroeconomic phenomena and what the boundaries of macroeconomics should be, there is broad agreement about the workhorse models of dynamic macroeconomic analysis. These include the Solow growth model, the neoclassical growth model, the overlapping-generations model and models of technological change and technology adoption. Since these are all models of economic growth, a thorough treatment of modern economic growth can also provide (and perhaps should provide) an introduction to this core material of modern macroeconomics. Although there are several good graduate-level macroeconomic textbooks, they typically spend relatively little time on the basic core material and do not develop the links between modern macroeconomic analysis and economic dynamics on the one hand and general equilibrium theory on the other. In contrast, the current book does not cover any of the short-run topics in macroeconomics, but provides a thorough and rigorous introduction to what I view to be the core of macroeconomics. Therefore, the second purpose of the book is to provide a graduate-level introduction to modern macroeconomics.

The selection of topics is designed to strike a balance between the two purposes of the book. Chapters 1, 3 and 4 introduce many of the salient features of the process of economic growth and the sources of cross-country differences in economic performance. Even though these chapters cannot do justice to the large literature on economic growth empirics, they provide a sufficient background for students to appreciate the set of issues that are central to the study of economic growth and also a platform for further study of this large literature.

Chapters 5-7 provide the conceptual and the mathematical foundations of modern macroeconomic analysis. Chapter 5 provides the microfoundations for much of the rest of the book (and for much of modern macroeconomics), while Chapters 6 and 7 provide a quick but relatively rigorous introduction to dynamic optimization. Most books on macroeconomics or economic growth use either continuous time or discrete time exclusively. I believe that a serious study of both economic growth and modern macroeconomics requires the student (and the researcher) to be able to go between discrete and continuous time, and choose whichever one is more convenient or appropriate for the set of questions at hand. Therefore, I have deviated from this standard practice and included both continuous-time and discrete-time material throughout the book.

Chapters 2, 8, 9 and 10 introduce the basic workhorse models of modern macroeconomics and traditional economic growth, while Chapter 11 presents the first generation models of sustained (endogenous) economic growth. Chapters 12-15 cover models of technological progress, which are an essential part of any modern economic growth course.

Chapter 16 generalizes the tools introduced in Chapter 6 to stochastic environments. Using these tools, Chapter 17 presents a number of models of stochastic growth, most notably, the neoclassical growth model under uncertainty, which is the foundation of much of modern macroeconomics (though it is often left out of economic growth courses). The canonical Real Business Cycle model is presented as an application. This chapter also covers another major workhorse model of modern macroeconomics, the incomplete markets model of Bewley. Finally, this chapter also presents a number of other approaches to modeling the interaction between incomplete markets and economic growth and shows how models of stochastic growth can be useful in understanding how economies transition from stagnation or slow growth to an equilibrium with sustained growth.

Chapters 18-21 cover a range of topics that are sometimes left out of economic growth textbooks. These include models of technology adoption, technology diffusion, the interaction between international trade and technology, the process of structural change, the demographic transition, the possibility of poverty traps, the effects of inequality on economic growth and the interaction between financial and economic development. These topics are important for creating a bridge between the empirical patterns we observe in practice and the theory. Most traditional growth models consider a single economy in isolation and often after it has already embarked upon a process of steady economic growth. A study of models that incorporate cross-country interdependences, structural change and the possibility of takeoffs will enable us to link core topics of development economics, such as structural change, poverty traps or the demographic transition, to the theory of economic growth.

Finally, Chapters 22 and 23 consider another topic often omitted from macroeconomics and economic growth textbooks; political economy. This is motivated by my belief that the study of economic growth would be seriously hampered if we failed to ask questions about the fundamental causes of why countries differ in their economic performances. These questions inexorably bring us to differences in economic policies and institutions across nations. Political economy enables us to develop models to understand why economic policies and institutions differ across countries and must therefore be an integral part of the study of economic growth.

A few words on the philosophy and organization of the book might also be useful for students and teachers. The underlying philosophy of the book is that all the results that are stated should be proved or at least explained in detail. This implies a somewhat different organization than existing books. Most textbooks in economics do not provide proofs for

many of the results that are stated or invoked, and mathematical tools that are essential for the analysis are often taken for granted or developed in appendices. In contrast, I have strived to provide simple proofs of almost all results stated in this book. It turns out that once unnecessary generality is removed, most results can be stated and proved in a way that is easily accessible to graduate students. In fact, I believe that even somewhat long proofs are much easier to understand than general statements made without proof, which leave the reader wondering about why these statements are true.

I hope that the style I have chosen not only makes the book self-contained, but also gives the students an opportunity to develop a thorough understanding of the material. In line with this philosophy, I present the basic mathematical tools necessary for the development of the main material within the body of the text. My own experience suggests that a "linear" progression, where the necessary mathematical tools are introduced when needed, makes it easier for the students to follow and appreciate the material. Consequently, analysis of stability of dynamical systems, dynamic programming in discrete time and optimal control in continuous time are all introduced within the main body of the text. This should both help the students appreciate the foundations of the theory of economic growth and also provide them with an introduction to the main tools of dynamic economic analysis, which are increasingly used in every subdiscipline of economics. Throughout, when some material is technically more difficult and can be skipped without loss of continuity, it is marked with a "*". Only material that is tangentially related to the main results in the text or those that should be familiar to most graduate students are left for the Mathematical Appendices.

I have also included a large number of exercises. Students can only gain a thorough understanding of the material by working through the exercises. The exercises that are somewhat more difficult are also marked with a "*".

This book can be used in a number of different ways. First, it can be used in a one-quarter or one-semester course on economic growth. Such a course might start with Chapters 1-4, then depending on the nature of the course, use Chapters 5-7 either for a thorough study of the general equilibrium and dynamic optimization foundations of growth theory or only for reference. Chapters 8-11 cover the traditional growth theory and Chapters 12-15 provide the basics of endogenous growth theory. Depending on time and interest, any selection of Chapters 16-23 can be used for the last part of such a course.

Second, the book can be used for a one-quarter first-year graduate-level course in macroeconomics. In this case, Chapter 1 is optional. Chapters 3, 5-7, 8-11 and 16 and 17 would be the core of such a course. The same material could also be covered in a one-semester course, but in this case, it could be supplemented either with some of the later chapters or with material from one of the leading graduate-level macroeconomic textbooks on short-run macroeconomics, fiscal policy, asset pricing, or other topics in dynamic macroeconomics.

Third, the book can be used for an advanced (second-year) course in economic growth or economic development. An advanced course on growth or development could use Chapters 1-11 as background and then focus on selected chapters from Chapters 12-23.

Finally, since the book is self-contained, I also hope that it can be used for self-study.

Acknowledgments. This book grew out of the first graduate-level introduction to macroeconomics course I have taught at MIT. Parts of the book have also been taught as part of a second-year graduate macroeconomics course. I would like to thank the students who attended these lectures and made comments that have improved the manuscript. I owe a special thanks to Monica Martinez-Bravo, Samuel Pienknagura, Lucia Tian Tian and

especially to Georgy Egorov, Michael Peters and Alp Simsek for outstanding research assistance. In fact, without Georgy, Michael and Alp's help, this book would have taken me much longer and would have contained many more errors. I also thank Lauren Fahey for editorial suggestions and help with the references. I would also like to thank Pol Antras, Kiminori Matsuyama, James Robinson, Jesus Fernandez-Villaverde and Pierre Yared for very valuable suggestions on multiple chapters, and George-Marios Angeletos, Binyamin Berdugo, Olivier Blanchard, Francesco Caselli, Melissa Dell, Leopoldo Fergusson, Peter Funk, Oded Galor, Hugo Hopenhayn, Simon Johnson, Chad Jones, Christos Koulovatianos, Omer Moav, Eduardo Morales, Ismail Saglam, Ekkehart Schlicht, Patricia Waeger and Jesse Zinn for useful suggestions and corrections on individual chapters.

Please note that this is a preliminary draft of the book manuscript. The draft certainly contains mistakes. Comments and suggestions for corrections are welcome.

Version 3: February, 2008.

Part 1

Introduction

CHAPTER 1

Economic Growth and Economic Development: The Questions

1.1. Cross-Country Income Differences

There are very large differences in income per capita and output per worker across countries today. Countries at the top of the world income distribution are more than thirty times as rich as those at the bottom. For example, in 2000, GDP (or income) per capita in the United States was over $34000. In contrast, income per capita is much lower in many other countries: about $8000 in Mexico, about $4000 in China, just over $2500 in India, only about $1000 in Nigeria, and much much lower in some other sub-Saharan African countries such as Chad, Ethiopia and Mali. These numbers are all in 2000 US dollars and are adjusted for purchasing power party (PPP) to allow for differences in relative prices of different goods across countries.[1] The cross-country income gap is considerably larger when there is no PPP-adjustment. For example, without the PPP adjustment, GDP per capita in India and China relative to the United States in 2000 would be lower by a factor of four or so.

Figure 1.1 provides a first look at these differences. It plots estimates of the distribution of PPP-adjusted GDP per capita across the available set of countries in 1960, 1980 and 2000. A number of features are worth noting. First, the 1960 density shows that 15 years after the end of World War II, most countries had income per capita less than $1500 (in 2000 US dollars); the mode of the distribution is around $1250. The rightwards shift of the distributions for 1980 and for 2000 shows the growth of average income per capita for the next 40 years. In 2000, the mode is slightly above $3000, but now there is another concentration of countries between $20,000 and $30,000. The density estimate for the year 2000 shows the considerable inequality in income per capita today.

Part of the spreading out of the distribution in Figure 1.1 is because of the increase in average incomes. It may therefore be more informative to look at the logarithm (log) of income per capita. It is more natural to look at the log of variables, such as income per capita, that grow over time, especially when growth is approximately proportional as suggested by Figure 1.8. This is for the simple reason that when $x(t)$ grows at a proportional rate, $\log x(t)$ grows linearly, and if $x_1(t)$ and $x_2(t)$ both grow by the same proportional amount, $\log x_1(t) - \log x_2(t)$ remains constant, while $x_1(t) - x_2(t)$ increases.

[1] All data are from the Penn World tables compiled by Summers and Heston. Details of data sources and more on PPP adjustment can be found in the References and Literature Section at the end of this chapter.

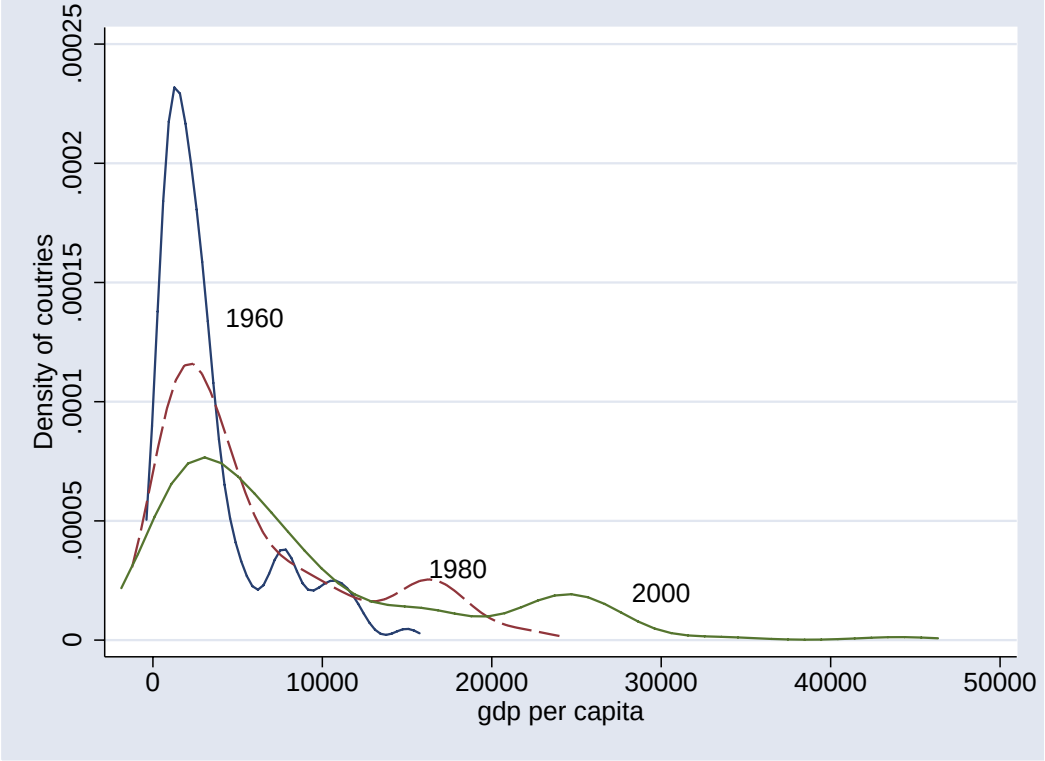

Figure 1.1. Estimates of the distribution of countries according to PPP-adjusted GDP per capita in 1960, 1980 and 2000.

Figure 1.2 shows a similar pattern, but now the spreading-out is more limited. This reflects the fact that while the absolute gap between rich and poor countries has increased considerably between 1960 and 2000, the proportional gap has increased much less. Nevertheless, it can be seen that the 2000 density for log GDP per capita is still more spread out than the 1960 density. In particular, both figures show that there has been a considerable increase in the density of relatively rich countries, while many countries still remain quite poor. This last pattern is sometimes referred to as the "stratification phenomenon", corresponding to the fact that some of the middle-income countries of the 1960s have joined the ranks of relatively high-income countries, while others have maintained their middle-income status or even experienced relative impoverishment.

Figures 1.1 and 1.2 demonstrate that there is somewhat greater inequality among nations. An equally relevant concept might be inequality among individuals in the world economy. Figures 1.1 and 1.2 are not directly informative on this, since they treat each country identically regardless of the size of its population. An alternative is presented in Figure 1.3, which shows the population-weighted distribution. In this case, countries such as China, India, the United States, and Russia receive greater weight because they have larger populations. The picture that emerges in this case is quite different. In fact, the 2000 distribution looks less spread out, with thinner left tail than the 1960 distribution. This reflects the fact that in

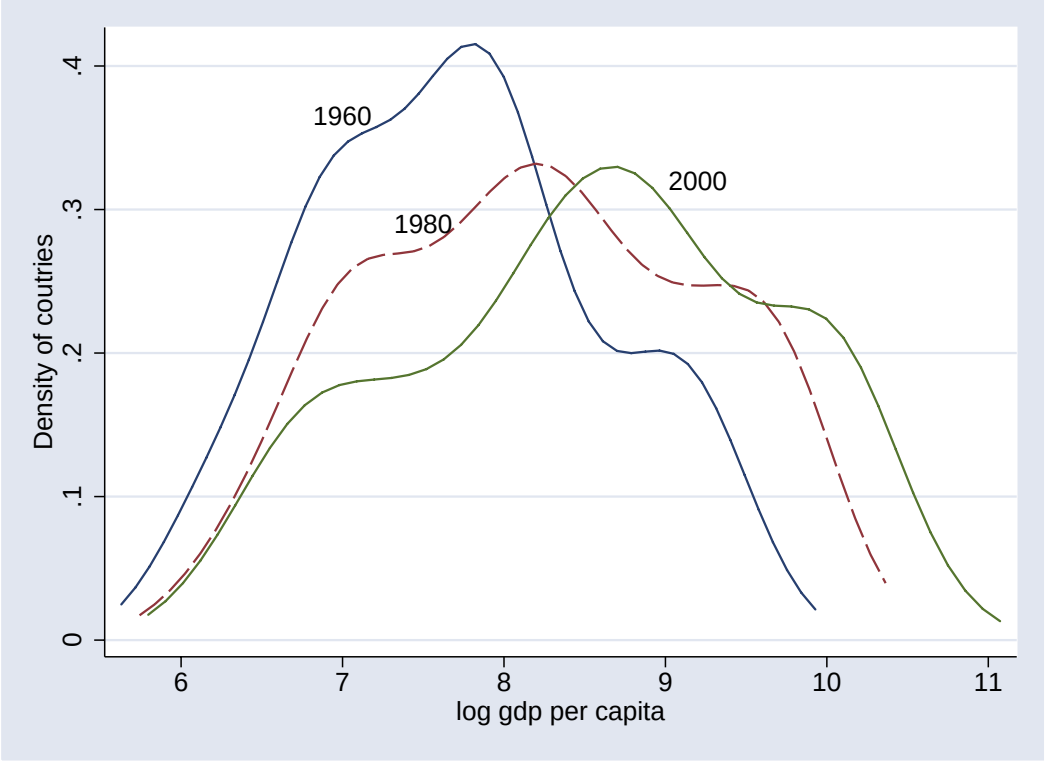

Figure 1.2. Estimates of the distribution of countries according to log GDP per capita (PPP-adjusted) in 1960, 1980 and 2000.

1960 China and India were among the poorest nations in the world, whereas their relatively rapid growth in the 1990s puts them into the middle-poor category by 2000. Chinese and Indian growth has therefore created a powerful force towards relative equalization of income per capita among the inhabitants of the globe.

Figures 1.1, 1.2 and 1.3 look at the distribution of GDP per capita. While this measure is relevant for the welfare of the population, much of growth theory focuses on the productive capacity of countries. Theory is therefore easier to map to data when we look at output (GDP) per worker. Moreover, key sources of difference in economic performance across countries are national policies and institutions. So for the purpose of understanding the sources of differences in income and growth across countries (as opposed to assessing welfare questions), the unweighted distribution is more relevant than the population-weighted distribution. Consequently, Figure 1.4 looks at the unweighted distribution of countries according to (PPP-adjusted) GDP per worker. Since internationally comparable data on employment are not available for a large number of countries, "workers" here refer to the total economically active population (according to the definition of the International Labour Organization). Figure 1.4 is very similar to Figure 1.2, and if anything, shows a greater concentration of countries in the relatively rich tail by 2000, with the poor tail remaining more or less the same as in Figure 1.2.

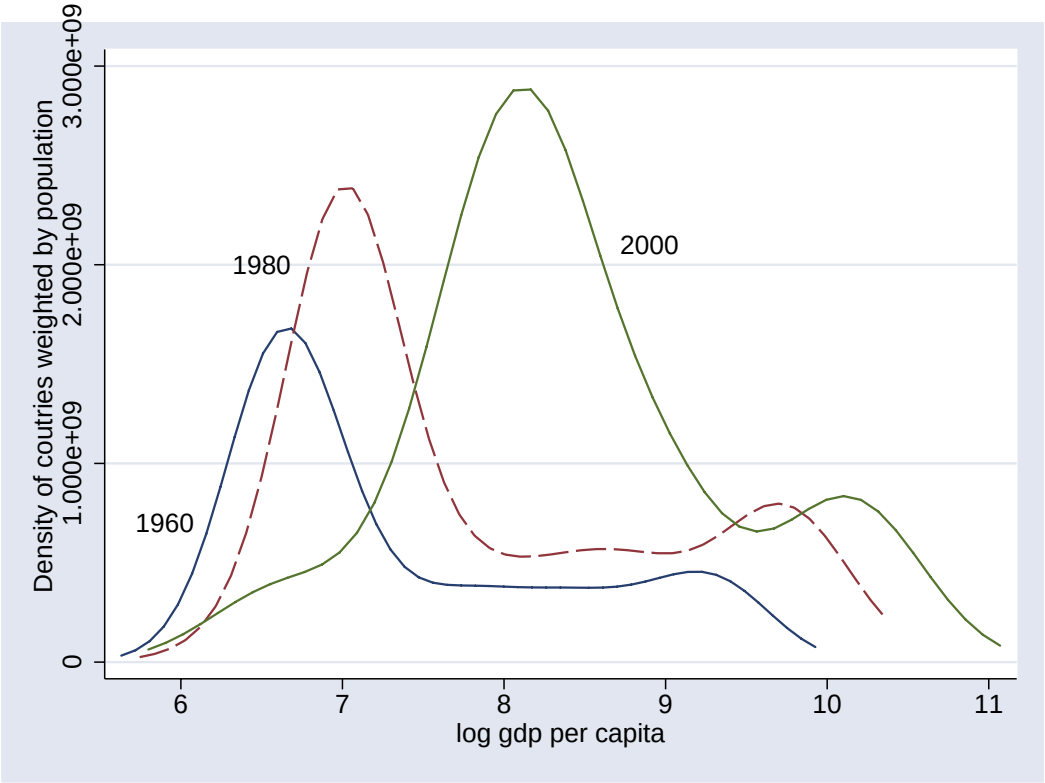

Figure 1.3. Estimates of the population-weighted distribution of countries according to log GDP per capita (PPP-adjusted) in 1960, 1980 and 2000.

Overall, Figures 1.1-1.4 document two important facts: first, there is a large amount of inequality in income per capita and income per worker across countries as shown by the highly dispersed distributions. Second, there is a slight but noticeable increase in inequality across nations (though not necessarily across individuals in the world economy).

1.2. Income and Welfare

Should we care about cross-country income differences? The answer is *definitely yes*. High income levels reflect high standards of living. Economic growth sometimes increases pollution or it may raise individual aspirations, so that the same bundle of consumption may no longer make an individual as happy. But at the end of the day, when one compares an advanced, rich country with a less-developed one, there are striking differences in the quality of life, standards of living and health.

Figures 1.5 and 1.6 give a glimpse of these differences and depict the relationship between income per capita in 2000 and consumption per capita and life expectancy at birth in the same year. Consumption data also come from the Penn World tables, while data on life expectancy at birth are available from the World Bank Development Indicators.

These figures document that income per capita differences are strongly associated with differences in consumption and differences in health as measured by life expectancy. Recall

Figure 1.4. Estimates of the distribution of countries according to log GDP per worker (PPP-adjusted) in 1960, 1980 and 2000.

also that these numbers refer to PPP-adjusted quantities, thus differences in consumption do not (at least in principle) reflect the fact that the same bundle of consumption goods costs different amounts in different countries. The PPP adjustment corrects for these differences and attempts to measure the variation in real consumption. Therefore, the richest countries are not only producing more than thirty times as much as the poorest countries, but are also consuming thirty times as much. Similarly, cross-country differences in health are quite remarkable; while life expectancy at birth is as high as 80 in the richest countries, it is only between 40 and 50 in many sub-Saharan African nations. These gaps represent huge welfare differences.

Understanding how some countries can be so rich while some others are so poor is one of the most important, perhaps *the* most important, challenges facing social science. It is important both because these income differences have major welfare consequences and because a study of these striking differences will shed light on how the economies of different nations function and sometimes how they fail to function.

The emphasis on income differences across countries implies neither that income per capita can be used as a "sufficient statistic" for the welfare of the average citizen nor that it is the only feature that we should care about. As we will discuss in detail later, the efficiency properties of the market economy (such as the celebrated *First Welfare Theorem* or Adam

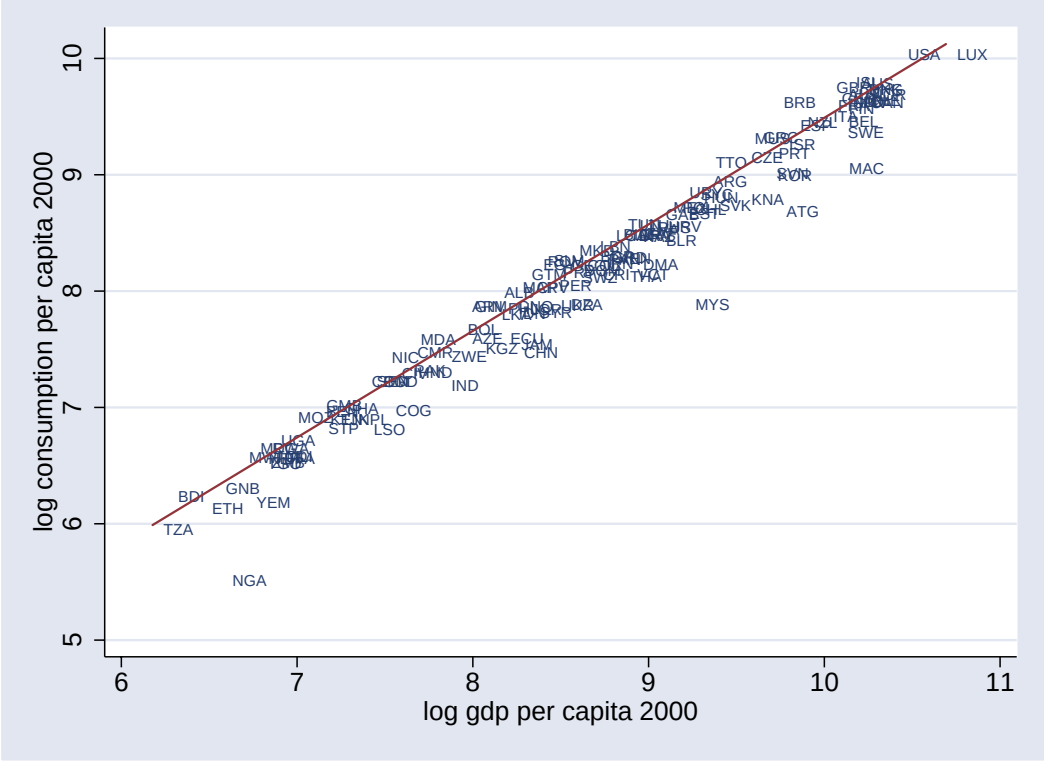

Figure 1.5. The association between income per capita and consumption per capita in 2000.

Smith's *invisible hand*) do not imply that there is no conflict among individuals or groups in society. Economic growth is generally good for welfare but it often creates "winners" and "losers." Joseph Schumpeter's famous notion of *creative destruction* emphasizes precisely this aspect of economic growth; productive relationships, firms and sometimes individual livelihoods will often be destroyed by the process of economic growth because growth is brought about by the introduction of new technologies and creation of new firms, replacing existing firms and technologies. This process creates a natural social tension, even in a growing society. Another source of social tension related to growth (and development) is that, as emphasized by Simon Kuznets and discussed in detail in Part 7 below, growth and development are often accompanied by sweeping structural transformations, which can also destroy certain established relationships and create yet other winners and losers in the process. One of the important questions of political economy, which will be discussed in the last part of the book, concerns how institutions and policies can be arranged so that those who lose out from the process of economic growth can be compensated or prevented from blocking economic progress via other means.

A stark illustration of the fact that growth does not always mean an improvement in the living standards of all or even most citizens in a society comes from South Africa under Apartheid. Available data (from gold mining wages) illustrate that from the beginning of

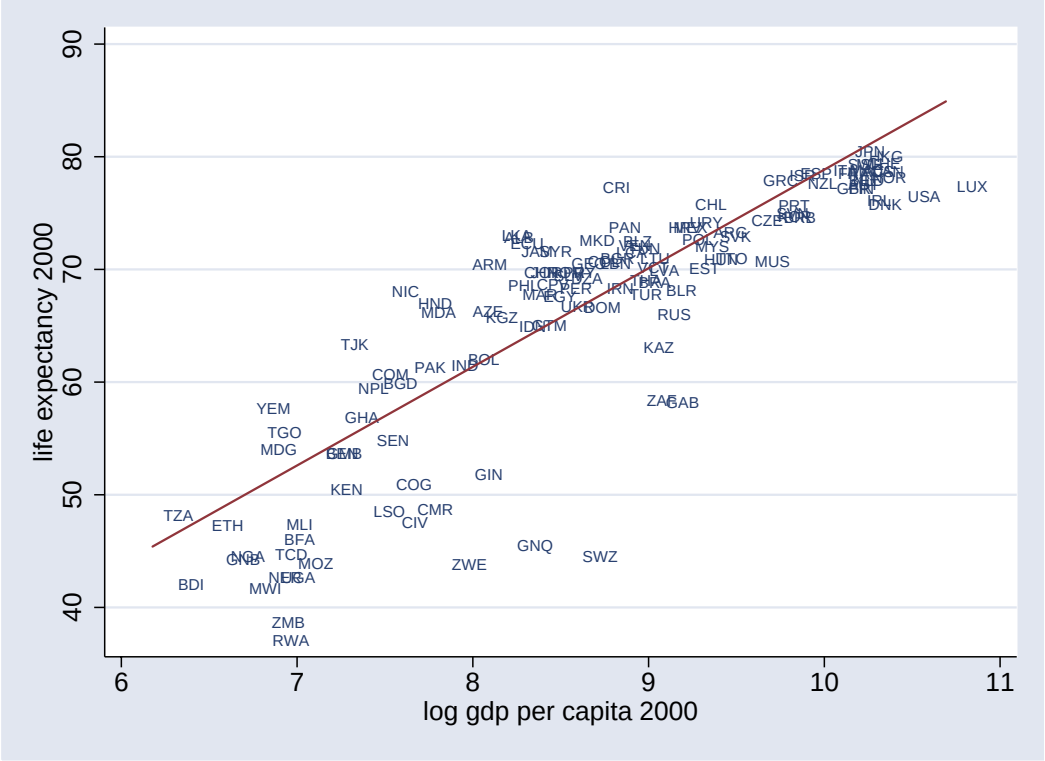

Figure 1.6. The association between income per capita and life expectancy at birth in 2000.

the 20th century until the fall of the Apartheid regime, GDP per capita grew considerably but the real wages of black South Africans, who make up the majority of the population, likely fell during this period. This of course does not imply that economic growth in South Africa was not beneficial. South Africa is still one of the richest countries in sub-Saharan Africa. Nevertheless, this observation alerts us to other aspects of the economy and also underlines the potential conflicts inherent in the growth process. Similarly, most existing evidence suggests that during the early phases of the British Industrial Revolution, which started the process of modern economic growth, the living standards of most workers may have fallen or at best remained stagnant. This pattern of potential divergence between GDP per capita and the economic fortunes of large numbers of individuals and society is not only interesting in and of itself, but it may also inform us about why certain segments of the society may be in favor of policies and institutions that do not encourage growth.

1.3. Economic Growth and Income Differences

How could one country be more than thirty times richer than another? The answer lies in differences in growth rates. Take two countries, A and B, with the same initial level of income at some date. Imagine that country A has 0% growth per capita, so its income per capita remains constant, while country B grows at 2% per capita. In 200 years' time country B will

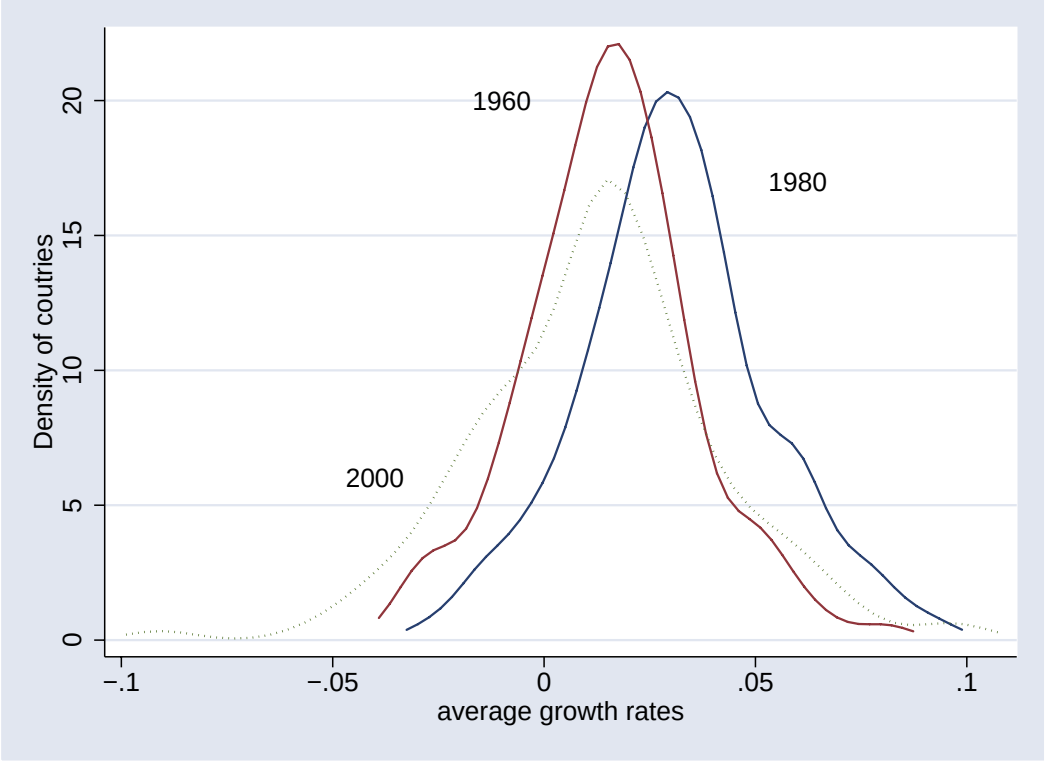

Figure 1.7. Estimates of the distribution of countries according to the growth rate of GDP per worker (PPP-adjusted) in 1960, 1980 and 2000.

be more than 52 times richer than country A. Therefore, the United States is considerably richer than Nigeria because it has grown steadily over an extended period of time, while Nigeria has not (and we will see that there is a lot of truth to this simple calculation; see Figures 1.8, 1.10 and 1.12).

In fact, even in the historically-brief postwar era, we see tremendous differences in growth rates across countries. This is shown in Figure 1.7 for the postwar era, which plots the density of growth rates across countries in 1960, 1980 and 2000. The growth rate in 1960 refers to the (geometric) average of the growth rate between 1950 and 1969, the growth rate in 1980 refers to the average growth rate between 1970 and 1989 and 2000 refers to the average between 1990 and 2000 (in all cases subject to data availability; all data from Penn World tables). Figure 1.7 shows that in each time interval, there is considerable variability in growth rates; the cross-country distribution stretches from negative growth rates to average growth rates as high as 10% a year.

Figure 1.8 provides another look at these patterns by plotting log GDP per capita for a number of countries between 1960 and 2000 (in this case, we look at GDP per capita instead of GDP per worker both for data coverage and also to make the figures more comparable to the historical figures below). At the top of the figure, we see US and UK GDP per capita increasing at a steady pace, with a slightly faster growth in the United States, so that the

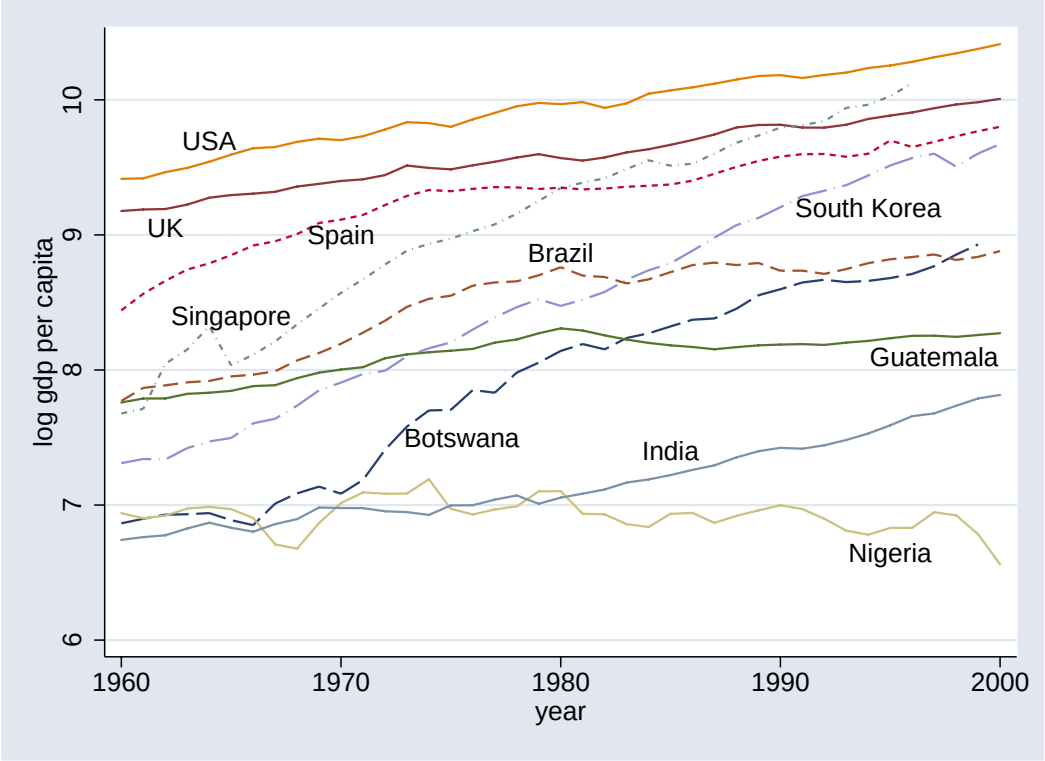

Figure 1.8. The evolution of income per capita in the United States, United Kingdom, Spain, Singapore, Brazil, Guatemala, South Korea, Botswana, Nigeria and India, 1960-2000.

log ("proportional") gap between the two countries is larger in 2000 than it is in 1960. Spain starts much poorer than the United States and the UK in 1960, but grows very rapidly between 1960 and the mid-1970s, thus closing the gap between itself and the United States and the UK. The three countries that show very rapid growth in this figure are Singapore, South Korea and Botswana. Singapore starts much poorer than the UK and Spain in 1960, but grows very rapidly and by the mid-1990s, it has become richer than both. South Korea has a similar trajectory, though it starts out poorer than Singapore and grows slightly less rapidly, so that by the end of the sample it is still a little poorer than Spain. The other country that has grown very rapidly is the "African success story" Botswana, which was extremely poor at the beginning of the sample. Its rapid growth, especially after 1970, has taken Botswana to the ranks of the middle-income countries by 2000.

The two Latin American countries in this picture, Brazil and Guatemala, illustrate the often-discussed Latin American economic malaise of the postwar era. Brazil starts out richer than Singapore, South Korea and Botswana and has a relatively rapid growth rate between 1960 and 1980. But it experiences stagnation from 1980 onwards, so that by the end of the sample Singapore, South Korea and Botswana have all become richer than Brazil.

Guatemala's experience is similar but even more bleak. Contrary to Brazil, there is little growth in Guatemala between 1960 and 1980 and no growth between 1980 and 2000.

Finally, Nigeria and India start out at similar levels of income per capita as Botswana but experience little growth until the 1980s. Starting in 1980, the Indian economy experiences relatively rapid growth, though this has not been sufficient for its income per capita to catch up with the other nations in the figure. Finally, Nigeria, in a pattern that is unfortunately all-too-familiar in sub-Saharan Africa, experiences a contraction of its GDP per capita, so that in 2000 it is in fact poorer than it was in 1960.

The patterns shown in Figure 1.8 are what we would like to understand and explain. Why is the United States richer in 1960 than other nations and able to grow at a steady pace thereafter? How did Singapore, South Korea and Botswana manage to grow at a relatively rapid pace for 40 years? Why did Spain grow relatively rapidly for about 20 years, but then slow down? Why did Brazil and Guatemala stagnate during the 1980s? What is responsible for the disastrous growth performance of Nigeria?

1.4. Origins of Today's Income Differences and World Economic Growth

The growth rate differences shown in Figures 1.7 and 1.8 are interesting in their own right and could also be, in principle, responsible for the large differences in income per capita we observe today. But are they? The answer is *no*. Figure 1.8 shows that in 1960 there was already a very large gap between the United States on the one hand and India and Nigeria on the other.

This can be seen more easily in Figure 1.9, which plots log GDP per worker in 2000 versus log GDP per capita in 1960 (in both cases relative to the US value) superimposed over the 45° line. Most observations are around the 45° line, indicating that the relative ranking of countries has changed little between 1960 and 2000. Thus the origins of the very large income differences across nations are not to be found in the postwar era. There are striking growth differences during the postwar era, but the evidence presented so far suggests that the "world income distribution" has been more or less stable, with a slight tendency towards becoming more unequal.

If not in the postwar era, when did this growth gap emerge? The answer is that much of the divergence took place during the 19th and early 20th centuries. Figures 1.10-1.12 give a glimpse of these 19th-century developments by using the data compiled by Angus Maddison for GDP per capita differences across nations going back to 1820 (or sometimes earlier). These data are less reliable than Summers-Heston's Penn World tables, since they do not come from standardized national accounts. Moreover, the sample is more limited and does not include observations for all countries going back to 1820. Finally, while these data include a correction for PPP, this is less reliable than the price comparisons used to construct the price indices in the Penn World tables. Nevertheless, these are the best available estimates for differences in prosperity across a large number of nations going back to the 19th century.

Figure 1.9. Log GDP per worker in 2000 versus log GDP per worker in 1960, together with the 45° line.

Figure 1.10 illustrates the divergence; it depicts the evolution of average income between five groups of countries, Western Offshoots of Europe (the United States, Canada, Australia and New Zealand), Western Europe, Latin America, Asia and Africa. It shows the relatively rapid growth of the Western Offshoots and West European countries during the 19th century, while Asia and Africa remained stagnant and Latin America showed little growth. The relatively small (proportional) income gap in 1820 had become much larger by 1960.

Another major macroeconomic fact is visible in Figure 1.10: Western Offshoots and West European nations experience a noticeable dip in GDP per capita around 1929. This is because of the famous Great Depression. Western offshoots, in particular the United States, only recovered fully from this large recession in the wake of WWII. How an economy can experience such a sharp decline in output and how it recovers from such a shock are among the major questions of macroeconomics. While the Great Depression falls outside the scope of the current book, we will later discuss the relationship between economic crises and growth as well as potential sources of volatility in economic growth.

A variety of other evidence suggests that differences in income per capita were even smaller once we go back further than 1820. Maddison also has estimates for average income for the same groups of countries going back to 1000 AD or even earlier. Figure 1.10 can be extended back using these data; the results are shown in Figure 1.11. Although these numbers

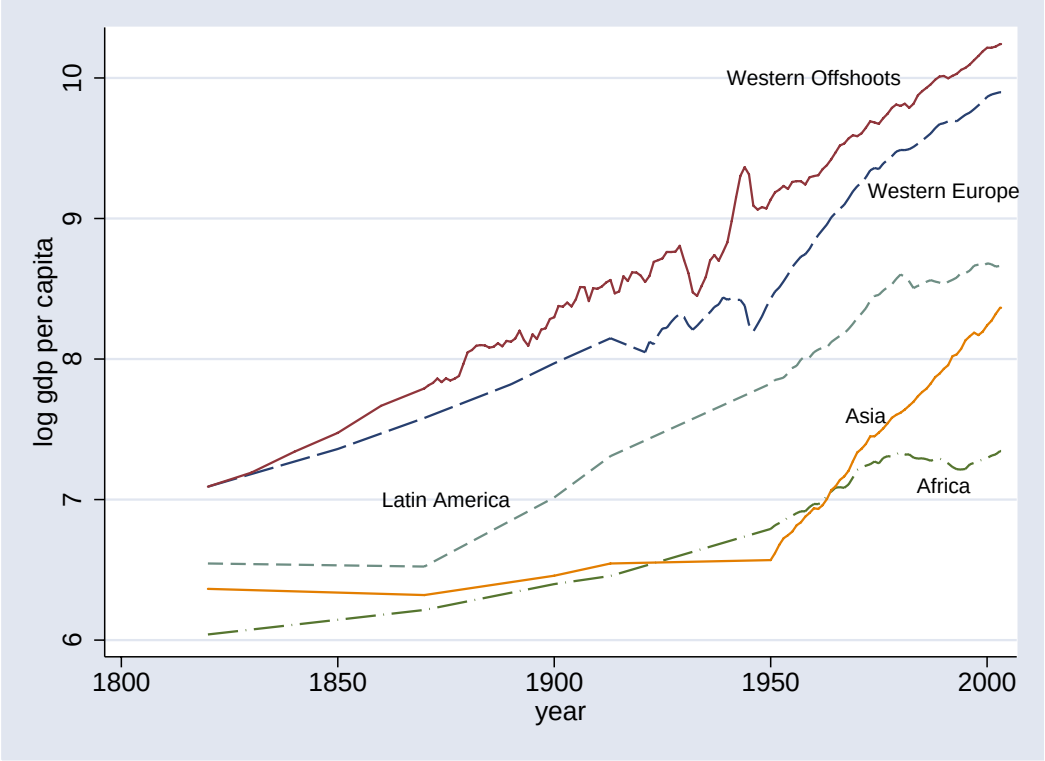

Figure 1.10. The evolution of average GDP per capita in Western Offshoots, Western Europe, Latin America, Asia and Africa, 1820-2000.

are based on scattered evidence and informed guesses, the general pattern is consistent with qualitative historical evidence and the fact that income per capita in any country cannot have been much less than $500 in terms of 2000 US dollars, since individuals could not survive with real incomes much less than this level. Figure 1.11 shows that as we go further back, the gap among countries becomes much smaller. This further emphasizes that the big divergence among countries has taken place over the past 200 years or so. Another noteworthy feature that becomes apparent from this figure is the remarkable nature of world economic growth. Much evidence suggests that there was only limited economic growth before the 18th century and certainly before the 15th century. While certain civilizations, including Ancient Greece, Rome, China and Venice, managed to grow, their growth was either not sustained (thus ending with collapses and crises) or progressed only at a slow pace. No society before 19th-century Western Europe and the United States achieved steady growth at comparable rates.

Notice also that Maddison's estimates show a slow but steady increase in West European GDP per capita even earlier, starting in 1000. This assessment is not shared by all economic historians, many of whom estimate that there was little increase in income per capita before 1500 or even before 1800. For our purposes this is not central, however. What is important is that, using Walter Rostow's terminology, Figure 1.11 shows a pattern of *takeoff* into sustained growth; the economic growth experience of Western Europe and Western Offshoots

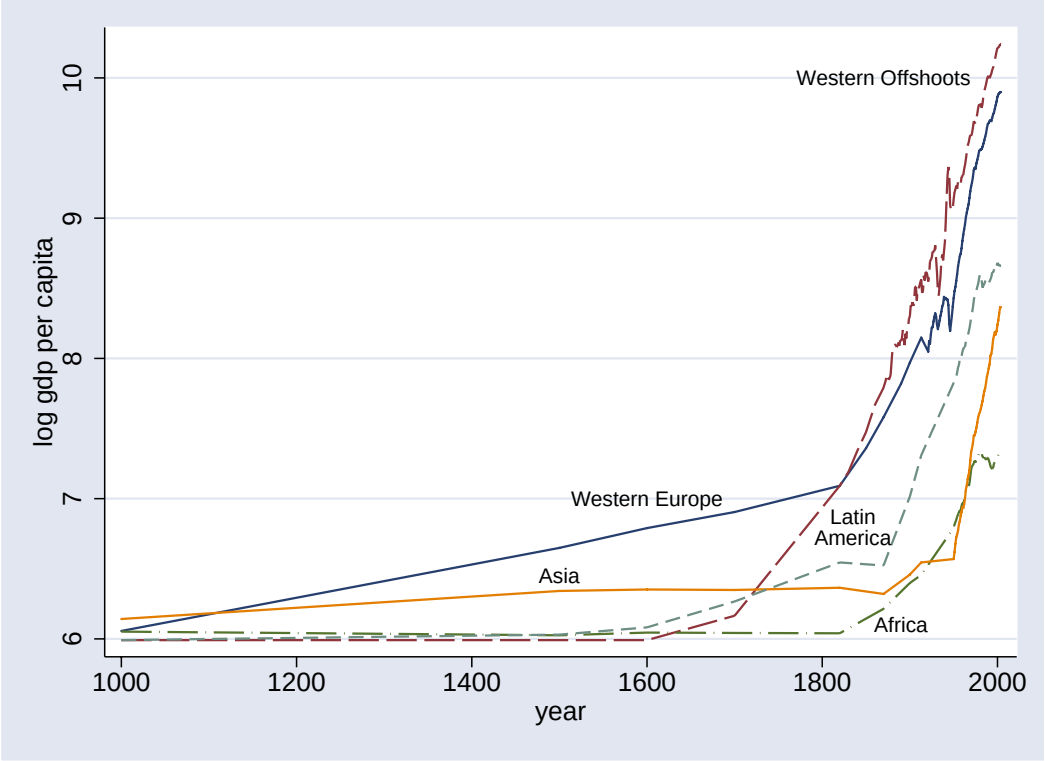

Figure 1.11. The evolution of average GDP per capita in Western Offshoots, Western Europe, Latin America, Asia and Africa, 1000-2000.

appears to have changed dramatically about 200 years or so ago. Economic historians debate whether there was a discontinuous change in economic activity that deserves the terms takeoff or Industrial Revolution. This debate is besides the point for our purposes. Whether or not the change was discontinuous, it was present and transformed the functioning of many economies. As a result of this transformation, the stagnant or slowly-growing economies of Europe embarked upon a path of sustained growth. The origins of today's riches and also of today's differences in prosperity are to be found in this pattern of takeoff during the 19th century. In the same time that much of Western Europe and its Offshoots grew rapidly, much of the rest of the world did not experience a comparable takeoff (or did so much later). Therefore, an understanding of modern economic growth and current cross-country income differences ultimately necessitates an inquiry into the causes of why the takeoff occurred, why it did so about 200 years ago and why it took place only in some areas and not in others.

Figure 1.12 shows the evolution of income per capita for the United States, Britain, Spain, Brazil, China, India and Ghana. This figure confirms the patterns shown in Figure 1.10 for averages, with the United States Britain and Spain growing much faster than India and Ghana throughout, and also much faster than Brazil and China except during the growth spurts experienced by these two countries.

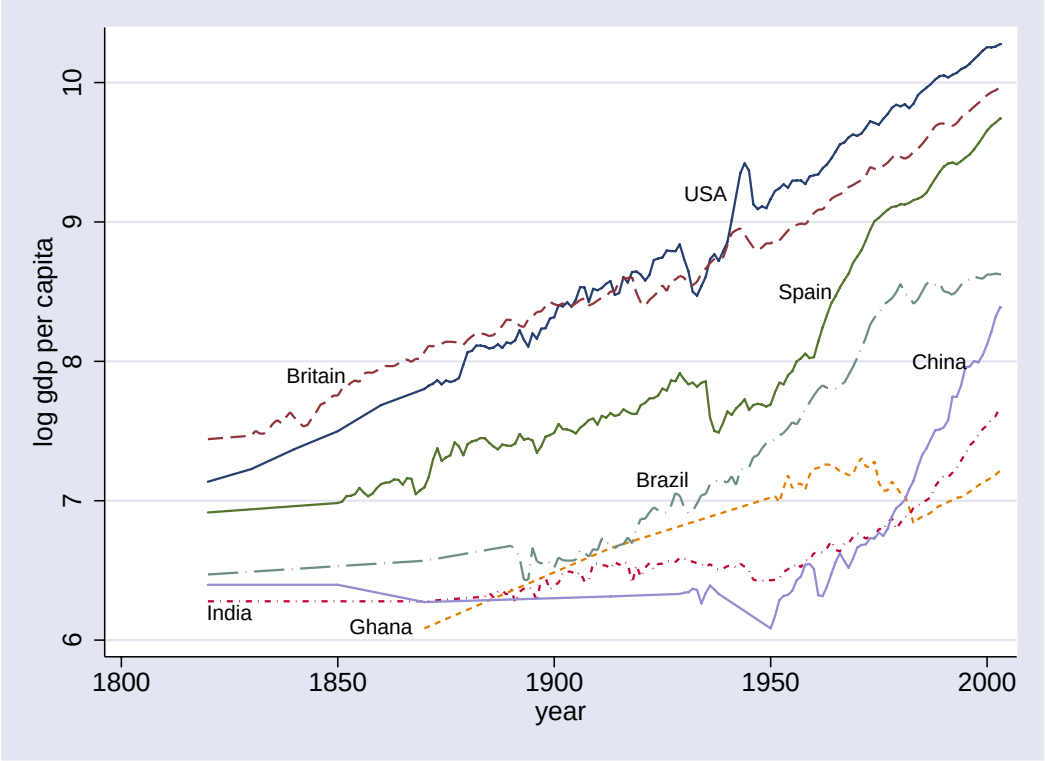

Figure 1.12. The evolution of income per capita in the United States, Britain, Spain, Brazil, China, India and Ghana, 1820-2000.

Overall, on the basis of the available information we can conclude that the origins of the current cross-country differences in income per capita are in the 19th and early 20th centuries (or perhaps even during the late 18th century). This cross-country divergence took place at the same time as a number of countries in the world "took off" and achieved sustained economic growth. Therefore understanding modern economic growth is not only interesting and important in its own right, but also holds the key to understanding the causes of cross-country differences in income per capita today.

1.5. Conditional Convergence

We have so far documented the large differences in income per capita across nations, the slight divergence in economic fortunes over the postwar era and the much larger divergence since the early 1800s. The analysis focused on the "unconditional" distribution of income per capita (or per worker). In particular, we looked at whether the income gap between two countries increases or decreases regardless of these countries' "characteristics" (e.g., institutions, policies, technology or even investments). Barro and Sala-i-Martin argue that it is instead more informative to look at the "conditional" distribution. Here the question is whether the income gap between two countries that are similar in observable characteristics is becoming narrower or wider over time. In this case, the picture is one of conditional convergence: in

the postwar period, the income gap between countries that share the same characteristics typically closes over time (though it does so quite slowly). This is important both for understanding the statistical properties of the world income distribution and also as an input into the types of theories that we would like to develop.

How do we capture conditional convergence? Consider a typical "Barro growth regression":

$$g_{t,t-1} = \beta \ln y_{t-1} + \mathbf{X}_{t-1}^T \boldsymbol{\alpha} + \varepsilon_t \tag{1.1}$$

where $g_{t,t-1}$ is the *annual* growth rate between dates $t-1$ and t, y_{t-1} is output per worker (or income per capita) at date $t-1$, and $\mathbf{X}_{t-1}$ is a vector of variables that the regression is conditioning on with coefficient vector $\boldsymbol{\alpha}$ (and $\mathbf{X}^T$ denotes the transpose of this vector, see Appendix Chapters A and B). These variables are included because they are potential determinants of steady state income and/or growth. First note that without covariates eq. (1.1) is quite similar to the relationship shown in Figure 1.9 above. In particular, since $g_{t,t-1} \simeq \ln y_t - \ln y_{t-1}$, eq. (1.1) can be written as

$$\ln y_t \simeq (1+\beta) \ln y_{t-1} + \varepsilon_t.$$

Figure 1.9 showed that the relationship between log GDP per worker in 2000 and log GDP per worker in 1960 can be approximated by the 45° line, so that in terms of this equation, β should be approximately equal to 0. This is confirmed by Figure 1.13, which depicts the relationship between the (geometric) average growth rate between 1960 and 2000 and log GDP per worker in 1960. This figure reiterates that there is no "unconditional" convergence for the entire world over the postwar period.

While there is no convergence for the entire world, when we look among the "OECD" nations,[2] we see a different pattern. Figure 1.14 shows that there is a strong negative relationship between log GDP per worker in 1960 and the annual growth rate between 1960 and 2000 among the OECD countries. What distinguishes this sample from the entire world sample is the relative homogeneity of the OECD countries, which have much more similar institutions, policies and initial conditions than the entire world. This suggests that there might be a type of conditional convergence when we control for certain country characteristics potentially affecting economic growth.

This is what the vector $\mathbf{X}_{t-1}$ captures in eq. (1.1). In particular, when this vector includes variables such as years of schooling or life expectancy, using cross-sectional regressions Barro and Sala-i-Martin estimate β to be approximately -0.02, indicating that the income gap between countries that have the same human capital endowment has been narrowing over the postwar period on average at about 2 percent a year. When this equation is estimated using a panel data and the vector $\mathbf{X}$ includes a full set of country fixed effects, then the estimates of β become more negative, indicating faster convergence.

[2] "OECD" here refers to the initial members of the OECD club and excludes the more recent OECD members such as Turkey, Mexico and Korea.

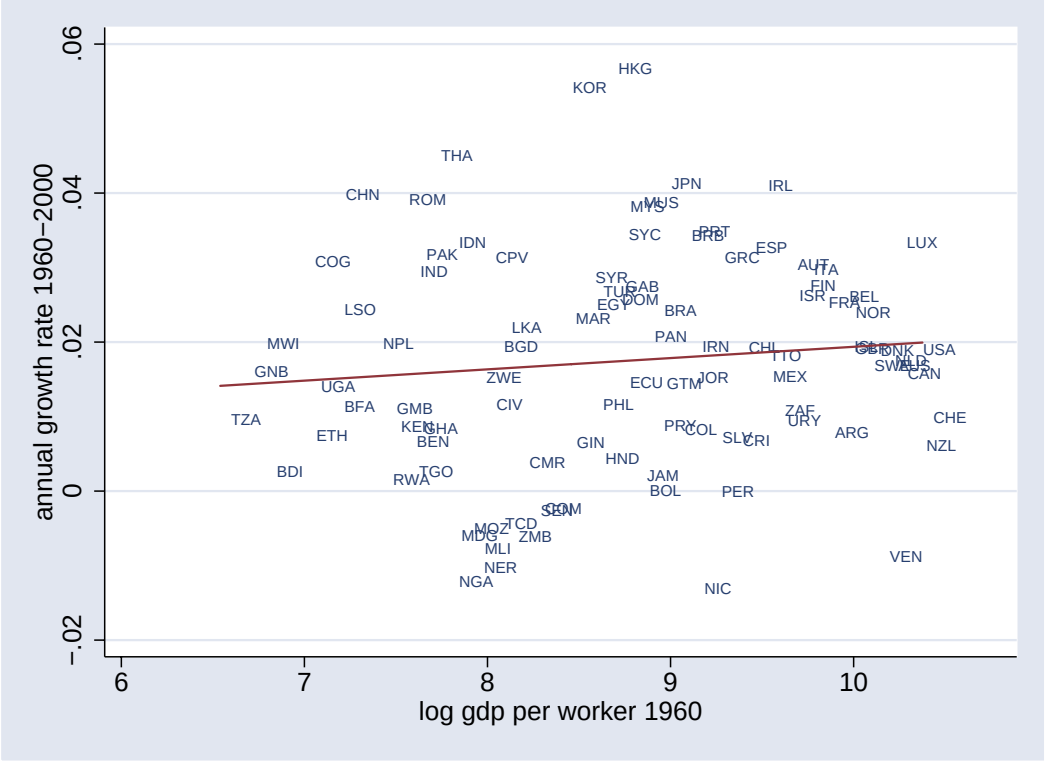

FIGURE 1.13. Annual growth rate of GDP per worker between 1960 and 2000 versus log GDP per worker in 1960 for the entire world.

In summary, there is no evidence of (unconditional) convergence in the world income distribution over the postwar era (in fact, the evidence suggests some amount of divergence in incomes across nations). But, there is some evidence for conditional convergence, meaning that the income gap between countries that are similar in observable characteristics appears to narrow over time. This last observation is relevant both for understanding among which countries the economic divergence has occurred and for determining what types of models we should consider for understanding the process of economic growth and the differences in economic performance across nations. For example, we will see that many growth models, including the basic Solow and the neoclassical growth models, suggest that there should be "transitional dynamics" as economies below their steady-state (target) level of income per capita grow towards that level. Conditional convergence is consistent with this type of transitional dynamics.

1.6. Correlates of Economic Growth

The discussion of conditional convergence in the previous section emphasized the importance of certain country characteristics that might be related to the process of economic growth. What types of countries grow more rapidly? Ideally, we would like to answer this

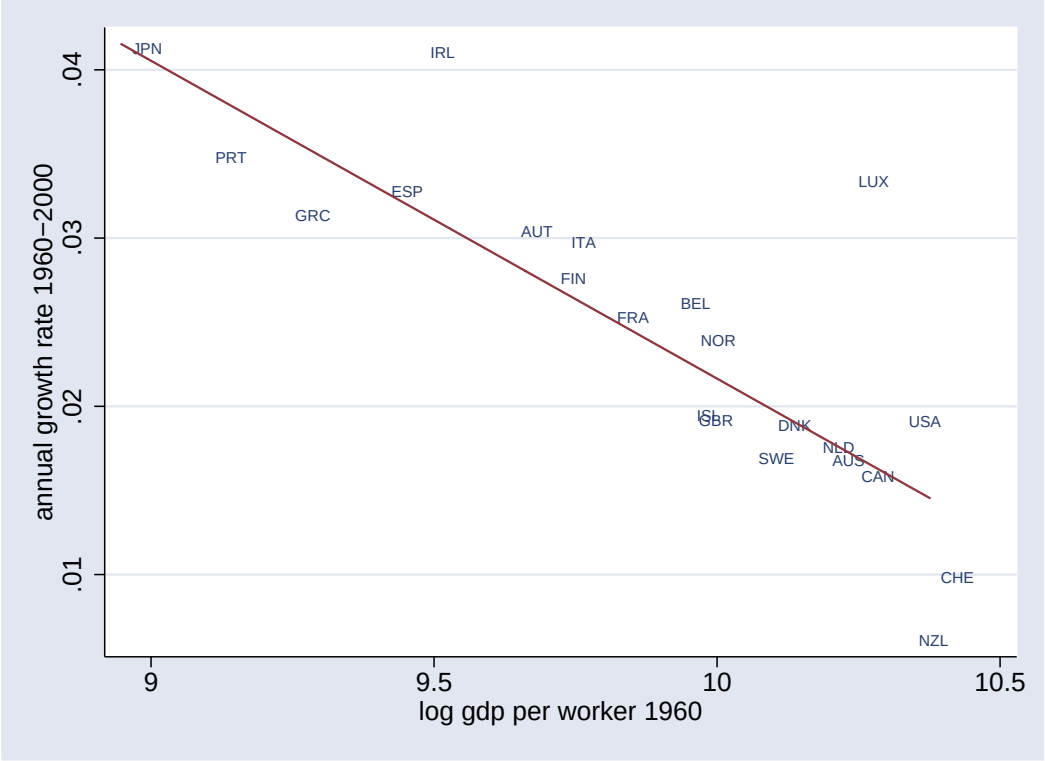

Figure 1.14. Annual growth rate of GDP per worker between 1960 and 2000 versus log GDP per worker in 1960 for core OECD countries.

question at a "causal" level. In other words, we would like to know which specific characteristics of countries (including their policies and institutions) have a causal effect on growth. A causal effect here refers to the answer to the following counterfactual thought experiment: if, all else equal, a particular characteristic of the country were changed "exogenously" (that is, not as part of equilibrium dynamics or in response to a change in other observable or unobservable variables), what would be the effect on equilibrium growth? Answering such causal questions is quite challenging, however, precisely because it is difficult to isolate changes in endogenous variables that are not driven by equilibrium dynamics or by some other potentially omitted factors.

For this reason, let us start with the more modest question of what factors correlate with post-war economic growth. With an eye to the theories that will come in the next two chapters, the two obvious candidates to look at are investments in physical capital and in human capital.

Figure 1.15 shows a strong positive association between the average growth of investment to GDP ratio and economic growth. Figure 1.16 shows a positive correlation between average years of schooling and economic growth. These figures therefore suggest that the countries that have grown faster are typically those that have invested more in physical capital and those that had greater human capital at the beginning of the postwar era. It has to be

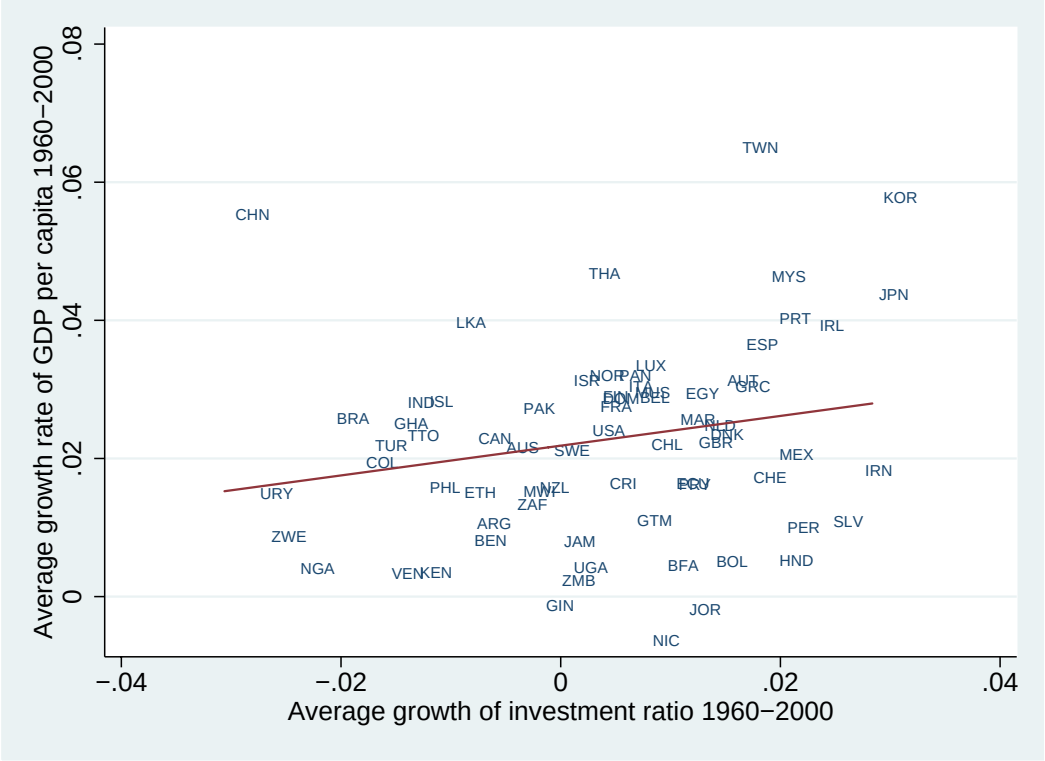

Figure 1.15. The relationship between average growth of GDP per capita and average growth of investments to GDP ratio, 1960-2000.

stressed that these figures do not imply that physical or human capital investment are the causes of economic growth (even though we expect from basic economic theory that they should contribute to output growth). So far these are simply correlations, and they are likely driven, at least in part, by omitted factors affecting both investment and schooling on the one hand and economic growth on the other.

We will investigate the role of physical and human capital in economic growth further in Chapter 3. One of the major points that will emerge from our analysis there is that focusing only on physical and human capital is not sufficient. Both to understand the process of sustained economic growth and to account for large cross-country differences in income, we also need to understand why societies differ in the efficiency with which they use their physical and human capital. Economists normally use the shorthand expression "technology" to capture factors other than physical and human capital that affect economic growth and performance. It is therefore important to remember that technology differences across countries include not only genuine differences in production techniques and in the quality of machines used in production, but also differences in productive efficiency resulting from differences in the organization of production, from differences in the way that markets are organized and from potential market failures (see in particular Chapter 21 on differences in productive efficiency resulting from the organization of markets and from market failures). A

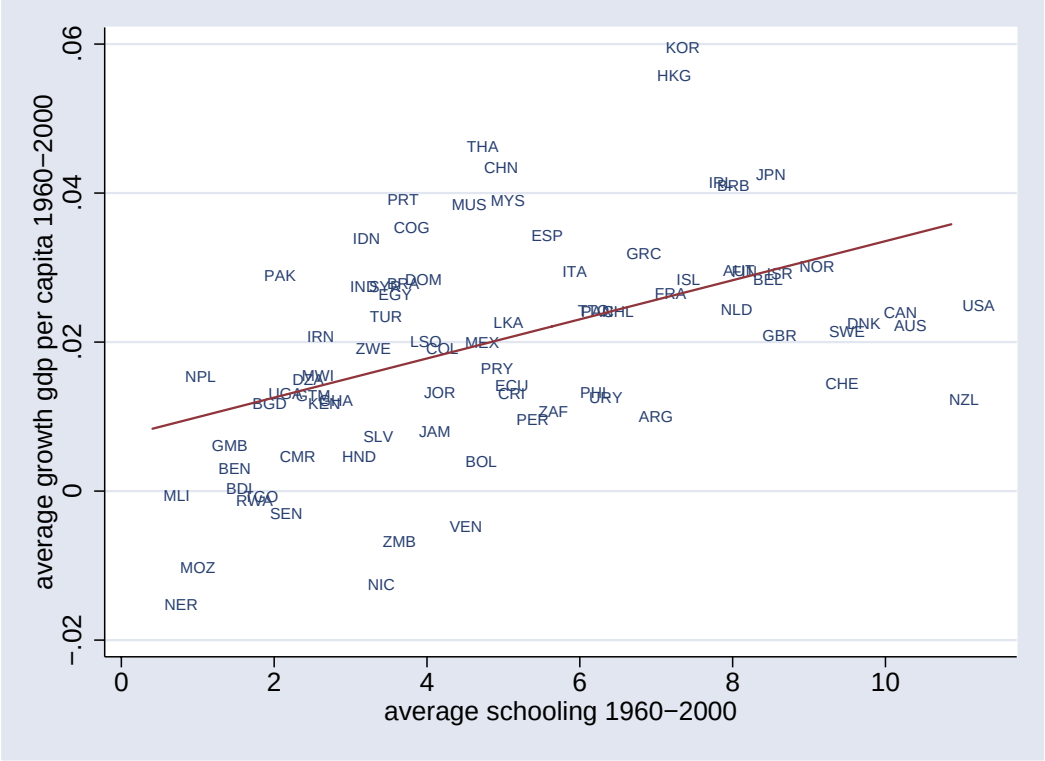

Figure 1.16. The relationship between average growth of GDP per capita and average years of schooling, 1960-2000.

detailed study of "technology" (broadly construed) is necessary for understanding both the world-wide process of economic growth and cross-country differences. The role of technology in economic growth will be investigated in Chapter 3 and in later chapters.

1.7. From Correlates to Fundamental Causes

The correlates of economic growth, such as physical capital, human capital and technology, will be our first topic of study. But these are only *proximate causes* of economic growth and economic success (even if we convince ourselves that there is an element of causality in the correlations shown above). It would not be entirely satisfactory to explain the process of economic growth and cross-country differences with technology, physical capital and human capital, since there are, presumably, reasons why technology, physical capital and human capital differ across countries. In particular, if these factors are so important in generating large cross country income differences and causing the takeoff into modern economic growth, why do certain societies fail to improve their technologies, invest more in physical capital, and accumulate more human capital?

Let us return to Figure 1.8 to illustrate this point further. This figure shows that South Korea and Singapore have managed to grow at very rapid rates over the past 50 years, while Nigeria has failed to do so. We can try to explain the successful performances of South Korea

and Singapore by looking at the correlates of economic growth—or at the proximate causes of economic growth. We can conclude, as many have done, that rapid capital accumulation has been a major cause of these growth miracles, and debate the relative roles of human capital and technology. We can simply blame the failure of Nigeria to grow on its inability to accumulate capital and to improve its technology. These answers are undoubtedly informative for understanding the mechanics of economic successes and failures of the postwar era. But at some level they will also not have answered the central questions: how did South Korea and Singapore manage to grow, while Nigeria failed to take advantage of the growth opportunities? If physical capital accumulation is so important, why did Nigeria not invest more in physical capital? If education is so important, why are education levels in Nigeria still so low and why is existing human capital not being used more effectively? The answer to these questions is related to the *fundamental causes* of economic growth.

I will refer to potential factors affecting why societies end up with different technology and accumulation choices as the fundamental causes of economic growth. At some level, fundamental causes are the factors that enable us to link the questions of economic growth to the concerns of the rest of social sciences and ask questions about the roles of policies, institutions, culture and exogenous environmental factors. At the risk of oversimplifying complex phenomena, we can think of the following list of potential fundamental causes: (i) luck (or multiple equilibria) that lead to divergent paths among societies with identical opportunities, preferences and market structures; (ii) geographic differences that affect the environment in which individuals live and that influence the productivity of agriculture, the availability of natural resources, certain constraints on individual behavior or even individual attitudes; (iii) institutional differences that affect the laws and regulations under which individuals and firms function and thus shape the incentives they have for accumulation, investment and trade; and (iv) cultural differences that determine individuals' values, preferences and beliefs. Chapter 4 will present a detailed discussion of the distinction between proximate and fundamental causes and what types of fundamental causes are more promising in explaining the process of economic growth and cross-country income differences.

For now, it is useful to briefly return to the contrast between South Korea and Singapore versus Nigeria, and ask the questions (even if we are not in a position to fully answer them yet): can we say that South Korea and Singapore owe their rapid growth to luck, while Nigeria was unlucky? Can we relate the rapid growth of South Korea and Singapore to geographic factors? Can we relate them to institutions and policies? Can we find a major role for culture? Most detailed accounts of post-war economics and politics in these countries emphasize the role of growth-promoting policies in South Korea and Singapore—including the relative security of property rights and investment incentives provided to firms. In contrast, Nigeria's postwar history is one of civil war, military coups, extreme corruption and an overall environment failing to provide incentives to businesses to invest and upgrade their technologies. It therefore seems necessary to look for fundamental causes of economic growth that make contact with these facts. Jumping ahead a little, it already appears implausible that luck can be the

major explanation for the differences in postwar economic performance; there were already significant economic differences between South Korea, Singapore and Nigeria at the beginning of the postwar era. It is also equally implausible to link the divergent fortunes of these countries to geographic factors. After all, their geographies did not change, but the growth spurts of South Korea and Singapore started in the postwar era. Moreover, even if we can say that Singapore benefited from being an island, without hindsight one might have concluded that Nigeria had the best environment for growth because of its rich oil reserves.[3] Cultural differences across countries are likely to be important in many respects, and the rapid growth of many Asian countries is often linked to certain "Asian values". Nevertheless, cultural explanations are also unlikely to provide the whole story when it comes to fundamental causes, since South Korean or Singaporean culture did not change much after the end of WWII, while their rapid growth performances are distinctly post-war phenomena. Moreover, while South Korea grew rapidly, North Korea, whose inhabitants share the same culture and Asian values, had one of the most disastrous economic performances of the past 50 years.

This admittedly quick (and partial) account suggests that we have to look at institutions and policies that affect incentives to accumulate physical and human capital and improve technology to develop a better understanding of the fundamental causes of economic growth. Institutions and policies were favorable to economic growth in South Korea and Singapore, but not in Nigeria. Understanding the fundamental causes of economic growth is, in large part, about understanding the impact of these institutions and policies on economic incentives and why, for example, they have been growth-enhancing in South Korea and Singapore, but not in Nigeria. The intimate link between fundamental causes and institutions highlighted by this discussion motivates the last part of the book, which is devoted to the political economy of growth, that is, to the study of how institutions affect growth and why they differ across countries.

An important caveat should be noted at this point. Discussions of geography, institutions and culture can sometimes be carried out without explicit reference to growth models or even to growth empirics. After all, this is what many non-economist social scientists do. However, fundamental causes can only have a big impact on economic growth if they affect parameters and policies that have a first-order influence on physical and human capital and technology. Therefore, an understanding of the mechanics of economic growth is essential for evaluating whether candidate fundamental causes of economic growth could indeed play the role that they are sometimes ascribed. Growth empirics plays an equally important role in distinguishing among competing fundamental causes of cross-country income differences. It is only by formulating parsimonious models of economic growth and confronting them with

[3]One can then turn this around and argue that Nigeria is poor because of a "natural resource curse," i.e., precisely because it has abundant and valuable natural resources. But this is not an entirely compelling empirical argument, since there are other countries, such as Botswana, with abundant natural resources that have grown rapidly over the past 50 years. More important, the only plausible channel through which abundance of natural resources may lead to worse economic outcomes is related to institutional and political economy factors. This then takes us to the realm of institutional fundamental causes.

data that we can gain a better understanding of both the proximate and the fundamental causes of economic growth.

1.8. The Agenda

The three major questions that have emerged from the brief discussion so far are:

(1) Why are there such large differences in income per capita and worker productivity across countries?
(2) Why do some countries grow rapidly while other countries stagnate?
(3) What sustains economic growth over long periods of time and why did sustained growth start 200 years or so ago?

In each case, a satisfactory answer requires a set of well-formulated models that illustrate the mechanics of economic growth and cross-country income differences, together with an investigation of the fundamental causes of the different trajectories which these nations have embarked upon. In other words, in each case we need a combination of theoretical models and empirical work.

The traditional growth models—in particular, the basic Solow and the neoclassical models—provide a good starting point, and the emphasis they place on investment and human capital seems consistent with the patterns shown in Figures 1.15 and 1.16. However, we will also see that technological differences across countries (either because of their differential access to technological opportunities or because of differences in the efficiency of production) are equally important. Traditional models treat technology (and market structure) as given or at best as evolving exogenously like a blackbox. But if technology is so important, we ought to understand why and how it progresses and why it differs across countries. This motivates our detailed study of endogenous technological progress and technology adoption. Specifically, we will try to understand how differences in technology may arise, persist and contribute to differences in income per capita. Models of technological change will also be useful in thinking about the sources of sustained growth of the world economy over the past 200 years and why the growth process took off 200 years or so ago and has proceeded relatively steadily since then.

Some of the other patterns we encountered in this chapter will inform us about the types of models that have the most promise in explaining economic growth and cross-country differences in income. For example, we have seen that cross-country income differences can only be accounted for by understanding why some countries have grown rapidly over the past 200 years, while others have not. Therefore, we need models that can explain how some countries can go through periods of sustained growth, while others stagnate.

Nevertheless, we have also seen that the postwar world income distribution is relatively stable (at most spreading out slightly from 1960 to 2000). This pattern has suggested to many economists that we should focus on models that generate large "permanent" cross-country differences in income per capita, but not necessarily large "permanent" differences in growth rates (at least not in the recent decades). This is based on the following reasoning:

with substantially different long-run growth rates (as in models of endogenous growth, where countries that invest at different rates grow at permanently different rates), we should expect significant divergence. We saw above that despite some widening between the top and the bottom, the cross-country distribution of income across the world is relatively stable over the postwar era.

Combining the post-war patterns with the origins of income differences related to the economic growth over the past two centuries suggests that we should look for models that can account both for long periods of significant growth differences and also for a "stationary" world income distribution, with large differences across countries. The latter is particularly challenging in view of the nature of the global economy today, which allows for free-flow of technologies and large flows of money and commodities across borders. We therefore need to understand how the poor countries fell behind and what prevents them today from adopting and imitating the technologies and the organizations (and importing the capital) of the richer nations.

And as the discussion in the previous section suggests, all of these questions can be (and perhaps should be) answered at two distinct, but related levels (and in two corresponding steps). The first step is to use theoretical models and data to understand the mechanics of economic growth. This step will shed light on the proximate causes of growth and explain differences in income per capita in terms of differences in physical capital, human capital and technology, and these in turn will be related to some other variables such as preferences, technology, market structure, openness to international trade and perhaps some distortions or policy variables.

The second step is to look at the fundamental causes underlying these proximate factors and to investigate why some societies are organized differently than others. Why do societies have different market structures? Why do some societies adopt policies that encourage economic growth while others put up barriers against technological change? These questions are central to a study of economic growth and can only be answered by developing systematic models of the political economy of development and looking at the historical process of economic growth to generate data that can shed light on these fundamental causes.

Our next task is to systematically develop a series of models to understand the mechanics of economic growth. I will present a detailed exposition of the mathematical structure of a number of dynamic general equilibrium models that are useful for thinking about economic growth and related macroeconomic phenomena, and I will emphasize the implications of these models for the sources of differences in economic performance across societies. Only by understanding these mechanics can we develop a useful framework for thinking about the causes of why some countries are growing and some others are not, and why some countries are rich and others are not.

1.9. References and Literature

The empirical material presented in this chapter is largely standard and parts of it can be found in many books, though interpretations and exact emphases differ. Excellent introductions, with slightly different emphases, are provided in Jones's (1998, Chapter 1) and Weil's (2005, Chapter 1) undergraduate economic growth textbooks. Barro and Sala-i-Martin (2004) also present a brief discussion of the stylized facts of economic growth, though their focus is on postwar growth and conditional convergence rather than the very large cross-country income differences and the long-run perspective emphasized here. Excellent and very readable accounts of the key questions of economic growth, with a similar perspective to the one here, are provided in Helpman (2005) and in Aghion and Howitt's new book (2008). Aghion and Howitt also provide a very useful introduction to many of the same topics discussed in the current book.

Much of the data used in this chapter come from Summers-Heston's Penn World tables (latest version, Summers, Heston and Aten, 2005). These tables are the result of a very careful study by Robert Summers and Alan Heston to construct internationally comparable price indices and internationally comparable estimates of income per capita and consumption. PPP adjustment is made possible by these data. Summers and Heston (1991) give a very lucid discussion of the methodology for PPP adjustment and its use in the Penn World tables. PPP adjustment enables us to construct measures of income per capita that are comparable across countries. Without PPP adjustment, differences in income per capita across countries can be computed using the current exchange rate or some fundamental exchange-rate. There are many problems with such exchange-rate-based measures. The most important one is that they do not make an allowance for the fact that relative prices and even the overall price level differ markedly across countries. PPP-adjustment brings us much closer to differences in "real income" and "real consumption". Information on "workers" (active population), consumption and investment are also from this dataset. GDP, consumption and investment data from the Penn World tables are expressed in 1996 constant US dollars. Life expectancy data are from the World Bank's World Development Indicators CD-ROM, and refer to the average life expectancy of males and females at birth. This dataset also contains a range of other useful information. Schooling data are from Barro and Lee's (2002) dataset, which contains internationally comparable information on years of schooling.

In all figures and regressions, growth rates are computed as geometric averages. In particular, the geometric average growth rate of output per capita y between date t and $t+T$ is

$$g_{t,t+T} \equiv \left(\frac{y_{t+T}}{y_t}\right)^{1/T} - 1.$$

The geometric average growth rate is more appropriate to use in the context of income per capita than the arithmetic average, since the growth rate refers to "proportional growth". It can be easily verified from this formula that if $y_{t+1} = (1+g)\, y_t$ for all t, then $g_{t+T} = g$.

Historical data are from various works by Angus Maddison, in particular, Maddison (2001, 2005). While these data are not as reliable as the estimates from the Penn World tables, the general patterns they show are typically consistent with evidence from a variety of different sources. Nevertheless, there are points of contention. For example, as Figure 1.11 shows, Maddison's estimates show a slow but relatively steady growth of income per capita in Western Europe starting in 1000. This is disputed by some historians and economic historians. A relatively readable account, which strongly disagrees with this conclusion, is provided in Pomeranz (2001), who argues that income per capita in Western Europe and the Yangtze Valley in China were broadly comparable as late as 1800. This view also receives support from recent research by Allen (2004), which documents that the levels of agricultural productivity in 1800 were comparable in Western Europe and China. Acemoglu, Johnson and Robinson (2002 and 2005) use urbanization rates as a proxy for income per capita and obtain results that are intermediate between those of Maddison and Pomeranz. The data in Acemoglu, Johnson and Robinson (2002) also confirms the fact that there were very limited income differences across countries as late as the 1500s and that the process of rapid economic growth started sometime in the 19th century (or perhaps in the late 18th century). Recent research by Broadberry and Gupta (2006) also disputes Pomeranz's arguments and gives more support to a pattern in which there was already an income gap between Western Europe and China by the end of the 18th century.

The term *takeoff* I used in Section 1.4 is introduced in Walter Rostow's famous book *Stages of Economic Growth* (1960) and has a broader connotation than the term "Industrial Revolution," which economic historians typically use to refer to the process that started in Britain at the end of the 18th century (e.g., Ashton, 1968). Mokyr (1990) contains an excellent discussion of the debate on whether the beginning of industrial growth was due to a continuous or discontinuous change. Consistent with my emphasis here, Mokyr concludes that this is secondary to the more important fact that the modern process of growth *did* start around this time.

There is a large literature on the "correlates of economic growth," starting with Barro (1991). This work is surveyed in Barro and Sala-i-Martin (2004) and Barro (1999). Much of this literature, however, interprets these correlations as causal effects, even when this is not warranted (see the further discussion in Chapters 3 and 4).

Note that while Figure 1.15 looks at the relationship between the average growth of investment to GDP ratio and economic growth, Figure 1.16 shows the relationship between average schooling (not its growth) and economic growth. There is a much weaker relationship between growth of schooling and economic growth, which may be for a number of reasons. First, there is considerable measurement error in schooling estimates (see Krueger and Lindahl, 2000). Second, as shown in some of the models that will be discussed later, the main role of human capital may be to facilitate technology adoption, thus we may expect a stronger relationship between the level of schooling and economic growth than the change in schooling

and economic growth (see Chapter 10). Finally, the relationship between the level of schooling and economic growth may be partly spurious, in the sense that it may be capturing the influence of some other omitted factors also correlated with the level of schooling; if this is the case, these omitted factors may be removed when we look at changes. While we cannot reach a firm conclusion on these alternative explanations, the strong correlation between the level of average schooling and economic growth documented in Figure 1.16 is interesting in itself.

The narrowing of income per capita differences in the world economy when countries are weighted by population is explored in Sala-i-Martin (2005). Deaton (2005) contains a critique of Sala-i-Martin's approach. The point that incomes must have been relatively equal around 1800 or before, because there is a lower bound on real incomes necessary for the survival of an individual, was first made by Maddison (1992), and was later popularized by Pritchett (1996). Maddison's estimates of GDP per capita and Acemoglu, Johnson and Robinson's estimates based on urbanization confirm this conclusion.

The estimates of the density of income per capita reported above are similar to those used by Quah (1994, 1995) and Jones (1996). These estimates use a nonparametric Gaussian kernel. The specific details of the kernel estimates do not change the general shape of the densities. Quah was also the first to emphasize the stratification in the world income distribution and the possible shift towards a "bi-modal" distribution, which is visible in Figure 1.3. He dubbed this the "Twin Peaks" phenomenon (see also Durlauf and Quah, 1994). Barro (1991) and Barro and Sala-i-Martin (1992) emphasize the presence and importance of conditional convergence and argue against the relevance of the stratification pattern emphasized by Quah and others. The estimate of conditional convergence of about 2% the year is from Barro and Sala-i-Martin (1992). Caselli, Esquivel and Lefort (1996) show that panel data regressions lead to considerably higher rates of conditional convergence.

The first economist to emphasize the importance of conditional convergence and conduct a cross-country study of convergence was Baumol (1986). However, Baumol used the available data at the time, which were of lower quality than the Summers-Heston data. This also made him conduct his empirical analysis on a selected sample of countries, potentially biasing his results (see De Long, 1991). Barro's (1991) and Barro and Sala-i-Martin's (1992) work using the Summers-Heston data has been instrumental in generating renewed interest in cross-country growth regressions.

The data on GDP growth and black real wages in South Africa are from Wilson (1972). Wages refer to real wages in gold mines. Feinstein (2004) provides an excellent economic history of South Africa. The implications of the British Industrial Revolution for real wages and living standards of workers are discussed in Mokyr (1993). Another example of rapid economic growth with falling real wages is provided by the experience of the Mexican economy in the early 20th century (see Gómez-Galvarriato, 1998). There is also evidence that during this period, the average height of the population might have been declining as well, which is often associated with falling living standards, see López Alonso and Porras Condy (2003).

There is a major debate on the role of technology and capital accumulation in the growth experiences of East Asian nations, particularly South Korea and Singapore. See Young (1994) for the argument that increases in physical capital and labor inputs explain almost all of the rapid growth in these two countries. See Klenow and Rodriguez-Clare (1996) and Hsieh (2001) for the opposite point of view.

The difference between proximate and fundamental causes will be discussed further in later chapters. This distinction is emphasized in a different context by Diamond (1996), though it is also implicitly present in North and Thomas's (1973) classic book. It is discussed in detail in the context of long-run economic development and economic growth in Acemoglu, Johnson and Robinson (2006). I will revisit these issues in greater detail in Chapter 4.

CHAPTER 2

The Solow Growth Model

The previous chapter introduced a number of basic facts and posed the main questions concerning the sources of economic growth over time and the causes of differences in economic performance across countries. These questions are central not only to growth theory but also to macroeconomics and social sciences more generally. Our next task is to develop a simple framework that can help us think about the *proximate* causes and the mechanics of the process of economic growth and cross-country income differences. We will use this framework both to study potential sources of economic growth and also to perform simple comparative statics to gain an understanding of what features of societies are conducive to higher levels of income per capita and more rapid economic growth.

Our starting point will be the so-called Solow-Swan model named after Robert (Bob) Solow and Trevor Swan, or simply the *Solow model* for the more famous of the two economists. These two economists published two pathbreaking articles in the same year, 1956 (Solow, 1956, and Swan, 1956) introducing the Solow model. Bob Solow later developed many implications and applications of this model and was awarded the Nobel prize in economics for these contributions. This model has shaped the way we approach not only economic growth but also the entire field of macroeconomics. Consequently, a byproduct of our analysis of this chapter will be a detailed exposition of a workhorse model of macroeconomics.

The Solow model is remarkable in its simplicity. Looking at it today, one may fail to appreciate how much of an intellectual breakthrough it was relative to what came before. Before the advent of the Solow growth model, the most common approach to economic growth built on the model developed by Roy Harrod and Evsey Domar (Harrod, 1939, Domar, 1946). The Harrod-Domar model emphasized potential dysfunctional aspects of economic growth, for example, how economic growth could go hand-in-hand with increasing unemployment (see Exercise 2.18 on this model). The Solow model demonstrated why the Harrod-Domar model was not an attractive place to start. At the center of the Solow growth model, distinguishing it from the Harrod-Domar model, is the *neoclassical* aggregate production function. This function not only enables the Solow model to make contact with microeconomics, but it also serves as a bridge between the model and the data as we will see in the next chapter.

An important feature of the Solow model, which will be shared by many models we will see in this book, is that it is a simple and *abstract* representation of a complex economy. At first, it may appear too simple or too abstract. After all, to do justice to the process of growth or macroeconomic equilibrium, we have to think of many different individuals with different tastes, abilities, incomes and roles in society, many different sectors and multiple

social interactions. Instead, the Solow model cuts through these complications by constructing a simple one-good economy, with little reference to individual decisions. Therefore, for us the Solow model will be both a starting point and a springboard for richer models.

In this chapter, I present the basic Solow model. The closely related neoclassical growth model will be presented in Chapter 8.

2.1. The Economic Environment of the Basic Solow Model

Economic growth and development are dynamic processes, focusing on how and why output, capital, consumption and population change over time. The study of economic growth and development therefore necessitates dynamic models. Despite its simplicity, the Solow growth model is a dynamic general equilibrium model (though, importantly, many key features of dynamic general equilibrium models emphasized in Chapter 5, such as preferences and dynamic optimization, are missing in this model).

The Solow model can be formulated either in discrete or in continuous time. I start with the discrete-time version, both because it is conceptually simpler and the cause it is more commonly used in macroeconomic applications. However, many growth models are formulated in continuous time and I will also provide a detailed exposition of the continuous-time version of the Solow model and show that it is often more convenient to work with.

2.1.1. Households and Production. Consider a closed economy, with a unique final good. The economy is in discrete time running to an infinite horizon, so that time is indexed by $t = 0, 1, 2, ...$ Time periods here may correspond to days, weeks, or years. So far there is no need to take a position on this.

The economy is inhabited by a large number of households. Throughout the book I use the terms households, individuals and agents. The Solow model makes relatively few assumptions on households because their optimization problem is not explicitly modeled. This is the main difference between the Solow and the *neoclassical growth* models. The latter is the Solow model plus dynamic consumer (household) optimization. To fix ideas, you may want to assume that all households are identical, so that the economy trivially admits *a representative household*—meaning that the demand and labor supply side of the economy can be represented as if it resulted from the behavior of a single household. The representative household assumption will be discussed in detail in Chapter 5.

What do we need to know about households in this economy? The answer is: not much. We have not yet endowed households with preferences (utility functions). Instead, for now, households are assumed to save a constant exogenous fraction $s \in (0, 1)$ of their disposable income—regardless of what else is happening in the economy. This is the same assumption used in basic Keynesian models and in the Harrod-Domar model mentioned above. It is also at odds with reality. Individuals do not save a constant fraction of their incomes; for example, if they did, then the announcement by the government that there will be a large tax increase next year should have no effect on their saving decisions, which seems both unreasonable and empirically incorrect. Nevertheless, the exogenous constant saving rate is a convenient

starting point and we will spend a lot of time in the rest of the book analyzing how consumers behave and make intertemporal choices.

The other key agents in the economy are firms. Firms, like consumers, are highly heterogeneous in practice. Even within a narrowly-defined sector of an economy (such as shoe manufacturing), no two firms are identical. But again for simplicity, let us start with an assumption similar to the representative household assumption, but now applied to firms: suppose that all firms in this economy have access to the same production function for the final good, or that the economy admits a *representative firm*, with a representative (or aggregate) production function. When this representive firm assumption is reasonable is discussed in Chapter 5. Suppose also that this aggregate production function exhibits *constant returns to scale* (see below for a definition). More explicitly, the aggregate production function for the unique final good is

$$Y(t) = F[K(t), L(t), A(t)] \tag{2.1}$$

where $Y(t)$ is the total amount of production of the final good at time t, $K(t)$ is the capital stock, $L(t)$ is total employment, and $A(t)$ is technology at time t. Employment can be measured in different ways. For example, we may want to think of $L(t)$ as corresponding to hours of employment or number of employees. The capital stock $K(t)$ corresponds to the quantity of "machines" (or more explicitly, equipment and structures) used in production, and it is typically measured in terms of the value of the machines. There are also multiple ways of thinking of capital (and equally many ways of specifying how capital comes into existence). Since the objective here is to start out with a simple workable model, I make the rather sharp simplifying assumption that capital is the same as the final good of the economy. However, instead of being consumed, capital is used in the production process of more goods. To take a concrete example, think of the final good as "corn". Corn can be used both for consumption and as an input, as "seed", for the production of more corn tomorrow. Capital then corresponds to the amount of corn used as seeds for further production.

Technology, on the other hand, has no natural unit and $A(t)$ is simply a *shifter* of the production function (2.1). For mathematical convenience, I will often represent $A(t)$ in terms of a number, but it is useful to bear in mind that, at the end of the day, it is a representation of a more abstract concept. Later I will discuss models in which $A(t)$ can be multidimensional, representing different types of technologies. As noted in Chapter 1, we may often want to think of a broad notion of technology, incorporating the effects of the organization of production and of markets on the efficiency with which the factors of production are utilized. In the current model, $A(t)$ represents all these effects.

A major assumption of the Solow growth model (and of the neoclassical growth model we will study in Chapter 8) is that technology is *free*; it is publicly available as a non-excludable, non-rival good. Recall that a good is *non-rival* if its consumption or use by others does not preclude an individual's consumption or use. It is *non-excludable*, if it is impossible to prevent any agents from using it or from consuming it. Technology is a good candidate for

a non-excludable, non-rival good; once the society has some knowledge useful for increasing the efficiency of production, this knowledge can be used by any firm without impinging on the use of it by others. Moreover, it is typically difficult to prevent firms from using this knowledge (at least once it is in the public domain and it is not protected by patents). For example, once the society knows how to make wheels, everybody can use that knowledge to make wheels without diminishing the ability of others to do the same (thus making the knowledge to produce wheels non-rival). Moreover, unless somebody has a well-enforced patent on wheels, anybody can decide to produce wheels (thus making the know-how to produce wheels non-excludable). The implication of the assumptions that technology is non-rival and non-excludable is that $A(t)$ is freely available to all potential firms in the economy and firms do not have to pay for making use of this technology. Departing from models in which technology is freely available will be a major step towards developing models of endogenous technological progress in Part 4 and towards understanding why there may be significant technology differences across countries in Part 6 below.

As an aside, you might want to note that some authors use x_t or K_t when working with discrete time and reserve the notation $x(t)$ or $K(t)$ for continuous time. Since I go back and forth between continuous time and discrete time, I use the latter notation throughout. When there is no risk of confusion, I drop time arguments, but whenever there is the slightest risk of confusion, I will err on the side of caution and include the time arguments.

Let us next impose the following standard assumptions on the aggregate production function.

Assumption 1. ***(Continuity, Differentiability, Positive and Diminishing Marginal Products, and Constant Returns to Scale)*** *The production function $F : \mathbb{R}^3_+ \to \mathbb{R}_+$ is twice differentiable in K and L, and satisfies*

$$F_K(K,L,A) \equiv \frac{\partial F(K,L,A)}{\partial K} > 0, \quad F_L(K,L,A) \equiv \frac{\partial F(K,L,A)}{\partial L} > 0,$$
$$F_{KK}(K,L,A) \equiv \frac{\partial^2 F(K,L,A)}{\partial K^2} < 0, \quad F_{LL}(K,L,A) \equiv \frac{\partial^2 F(K,L,A)}{\partial L^2} < 0.$$

Moreover, F exhibits constant returns to scale in K and L.

All of the components of Assumption 1 are important. First, the notation $F : \mathbb{R}^3_+ \to \mathbb{R}_+$ implies that the production function takes nonnegative arguments (that is, $K, L \in \mathbb{R}_+$) and maps to nonnegative levels of output ($Y \in \mathbb{R}_+$). It is natural that the level of capital and the level of employment should be positive. Since A has no natural units, it could have been negative. But there is no loss of generality in restricting it to be positive. The second important aspect of Assumption 1 is that F is a continuous function in its arguments and is also differentiable. There are many interesting production functions which are not differentiable and some interesting ones that are not even continuous. But working with differentiable functions makes it possible for us to use differential calculus, and the loss of some generality is a small price to pay for this convenience. Assumption 1 also specifies that marginal products are positive (so that the level of production increases with the amount of

inputs); this also rules out some potential production functions and can be relaxed without much complication (see Exercise 2.8). More importantly, Assumption 1 imposes that the marginal product of both capital and labor are diminishing, that is, $F_{KK} < 0$ and $F_{LL} < 0$, so that more capital, holding everything else constant, increases output by less and less. And the same applies to labor. This property is sometimes also referred to as "diminishing returns" to capital and labor. The degree of diminishing returns to capital plays a very important role in many of the results of the basic growth model. In fact, these features distinguish the Solow growth model from its antecedent, the Harrod-Domar model (see Exercise 2.18).

The other important assumption is that of constant returns to scale. Recall that F exhibits *constant returns to scale* in K and L if it is *linearly homogeneous* (homogeneous of degree 1) in these two variables. More specifically:

Definition 2.1. *Let $K \in \mathbb{N}$. The function $g : \mathbb{R}^{K+2} \to \mathbb{R}$ is* ***homogeneous of degree*** *m in $x \in \mathbb{R}$ and $y \in \mathbb{R}$ if*

$$g(\lambda x, \lambda y, z) = \lambda^m g(x, y, z) \text{ for all } \lambda \in \mathbb{R}_+ \text{ and } z \in \mathbb{R}^K.$$

It can be easily verified that linear homogeneity implies that the production function F is concave, though not strictly so (see Exercise 2.2).

Linearly homogeneous (constant returns to scale) production functions are particularly useful because of the following theorem:

Theorem 2.1. ***(Euler's Theorem)*** *Suppose that $g : \mathbb{R}^{K+2} \to \mathbb{R}$ is differentiable in $x \in \mathbb{R}$ and $y \in \mathbb{R}$, with partial derivatives denoted by g_x and g_y and is homogeneous of degree m in x and y. Then*

$$mg(x, y, z) = g_x(x, y, z)\, x + g_y(x, y, z)\, y \text{ for all } x \in \mathbb{R},\ y \in \mathbb{R} \text{ and } z \in \mathbb{R}^K.$$

Moreover, $g_x(x, y, z)$ and $g_y(x, y, z)$ are themselves homogeneous of degree $m-1$ in x and y.

Proof. We have that g is differentiable and

$$\lambda^m g(x, y, z) = g(\lambda x, \lambda y, z). \tag{2.2}$$

Differentiate both sides of eq. (2.2) with respect to λ, which gives

$$m\lambda^{m-1} g(x, y, z) = g_x(\lambda x, \lambda y, z)\, x + g_y(\lambda x, \lambda y, z)\, y$$

for any λ. Setting $\lambda = 1$ yields the first result. To obtain the second result, differentiate both sides of eq. (2.2) with respect to x:

$$\lambda g_x(\lambda x, \lambda y, z) = \lambda^m g_x(x, y, z).$$

Dividing both sides by λ establishes the desired result. □

2.1.2. Market Structure, Endowments and Market Clearing. For most of the book, I assume that all factor markets are competitive. This is yet another assumption that is not totally innocuous. Both labor markets and capital markets have imperfections, with certain important implications for economic growth. But these implications can be best appreciated by starting out with the competitive benchmark.

As in standard competitive general equilibrium models, the next step is to specify endowments, that is, the amounts of labor and capital that the economy starts with and who owns these endowments. It is natural to suppose that all factors of production are owned by households. In particular, households own all of the labor, which they supply inelastically. Inelastic supply means that there is some endowment of labor in the economy, for example equal to the population, $\bar{L}(t)$, and all of this will be supplied regardless of the price (as long as this price is nonnegative). The *labor market clearing* condition can then be expressed as:

$$L(t) = \bar{L}(t) \tag{2.3}$$

for all t, where $L(t)$ denotes the demand for labor (and also the level of employment). More generally, this equation should be written in complementary slackness form. In particular, let the *wage rate* (or the rental price of labor) at time t be $w(t)$, then the labor market clearing condition takes the form

$$L(t) \leq \bar{L}(t), \; w(t) \geq 0 \text{ and } \left(L(t) - \bar{L}(t)\right) w(t) = 0. \tag{2.4}$$

The complementary slackness formulation makes sure that labor market clearing does not happen at a negative wage—or that if labor demand happens to be low enough, employment could be below $\bar{L}(t)$ at zero wage. However, this will not be an issue in most of the models studied in this book, because Assumption 1 and competitive labor markets ensure that wages have to be strictly positive (see Exercise 2.1). In view of this result, I will use the simpler condition (2.3) throughout.

The households also own the capital stock of the economy and rent it to firms. Let us denote the *rental price of capital* at time t by $R(t)$. The capital market clearing condition is similar to (2.3) and requires the demand for capital by firms to be equal to the supply of capital by households:

$$K^s(t) = K^d(t),$$

where $K^s(t)$ is the supply of capital by households and $K^d(t)$ is the demand by firms. Capital market clearing is straightforward to ensure in the class of models analyzed in this book. In particular, it is sufficient that the amount of capital $K(t)$ used in production at time t (from firms' optimization behavior) the consistent with households' endowments and saving behavior.

Let us take households' initial holdings of capital, $K(0) \geq 0$, as given (as part of the description of the environment), and this will determine the initial condition of the dynamical system that will represent the equilibrium behavior in this economy. For now how this initial capital stock is distributed among the households is not important, since households' optimization decisions are not modeled explicitly and the economy is simply assumed to save

a fraction s of its income. When we turn to models with household optimization below, an important part of the description of the environment will be to specify the preferences and the budget constraints of households.

At this point, I could also introduce the price of the final good at time t, say $P(t)$. But there is no need to do this, since there is a choice of a numeraire commodity in this economy, whose price will be normalized to 1. In particular, as discussed in greater detail in Chapter 5, Walras's Law implies that the price of one of the commodities, the numeraire, should be normalized to 1. In fact, throughout I will do something stronger and normalize the price of the final good to 1 *in all periods*. Ordinarily, one cannot choose more than one numeraire—otherwise, one would be fixing the relative price between two numeraires. But again as will be explained in Chapter 5, we can build on an insight by Kenneth Arrow (Arrow, 1964) that it is sufficient to price *securities* (assets) that transfer one unit of consumption from one date (or state of the world) to another. In the context of dynamic economies, this implies that we need to keep track of an *interest rate* across periods, denoted by $r(t)$, and this will enable us to normalize the price of the final good to 1 in every period (and naturally, we need to keep track of the wage rate $w(t)$, which will determine the price of labor relative to final goods at any date t).

This discussion should already alert you to a central fact: you should think of all of the models in this book as *general equilibrium economies*, where different commodities correspond to the same good at different dates. Recall from basic general equilibrium theory that the same good at different dates (or in different states or localities) is a different commodity. Therefore, in almost all of the models in this book, there will be *an infinite number of commodities*, since time runs to infinity. This raises a number of special issues, which are discussed in Chapter 5 and later.

Returning to the basic Solow model, the next assumption is that capital depreciates, meaning that machines that are used in production lose some of their value because of wear and tear. In terms of the corn example above, some of the corn that is used as seeds is no longer available for consumption or for use as seeds in the following period. Let us assume that this depreciation takes an "exponential form," which is mathematically very tractable. This means that capital depreciates (exponentially) at the rate $\delta \in (0,1)$, so that out of 1 unit of capital this period, only $1-\delta$ is left for next period. As noted above, depreciation here stands for the wear and tear of the machinery, but it can also represent the replacement of old machines by new machines in more realistic models (see Chapter 14).

The loss of part of the capital stock affects the interest rate (rate of return to savings) faced by the household. Given the assumption of exponential depreciation at the rate δ and the normalization of the price of the final good to 1, this implies that the *interest rate* faced by the households will be

$$r(t) = R(t) - \delta,$$

where recall that $R(t)$ is the rental price of capital at time t. A unit of final good can be consumed now or it can be used as capital and rented to firms. In the latter case, the

household will receive $R(t)$ units of good in the next period as the rental price, but will lose δ units of the capital, since δ fraction of capital depreciates over time. This implies that the individual has given up one unit of commodity dated $t-1$ for $r(t) = R(t) - \delta$ units of commodity dated t. The relationship between $r(t)$ and $R(t)$ explains the similarity between the symbols for the interest rate and the rental rate of capital. The interest rate faced by households will play a central role in the dynamic optimization decisions of households below. In the Solow model, this interest rate does not directly affect the allocation of resources.

2.1.3. Firm Optimization and Equilibrium. We are now in a position to look at the optimization problem of firms and to the competitive equilibrium of this economy. Throughout the book I assume that the objective of firms is to maximize profits. Since I have assumed the existence of an aggregate production function, it is sufficient to consider the problem of a *representative firm*. Throughout, unless otherwise stated, I assume that capital markets are functioning, so firms can rent capital in spot markets. For a given technology level $A(t)$, and given factor prices $R(t)$ and $w(t)$, the profit maximization problem of the representative firm at time t can be represented by the following static problem:

$$\max_{K \geq 0, L \geq 0} F[K, L, A(t)] - R(t) K - w(t) L. \tag{2.5}$$

When there are irreversible investments or costs of adjustments, as discussed, for example, in Section 7.8 in Chapter 7, the maximization problem of firms becomes dynamic. But in the absence of these features, maximizing profits separately at each date t is equivalent to maximizing the net present discounted value of profits. This feature simplifies the analysis considerably.

A couple of additional features are worth noting:

(1) The maximization problem is set up in terms of aggregate variables. This is without loss of any generality given the representative firm.
(2) There is nothing multiplying the F term, since the price of the final good has been normalized to 1. Thus the first term in (2.5) is the revenues of the representative firm (or the revenues of all of the firms in the economy).
(3) This way of writing the problem already imposes competitive factor markets, since the firm is taking as given the rental prices of labor and capital, $w(t)$ and $R(t)$ (which are in terms of the numeraire, the final good).
(4) This is a concave problem, since F is concave (see Exercise 2.2).

One important feature to note is that, because F exhibits constant returns to scale (Assumption 1), the maximization problem (2.5) does not have a well-defined solution (see Exercise 2.3); either there does not exist any (K, L) that achieves the maximum value of this program (which is infinity) or $K = L = 0$ or multiple values of (K, L) will achieve the maximum value of this program (when this value happens to be 0). This problem is related to the fact that in a world with constant returns to scale, the size of each individual firm is not indeterminate (only aggregates are determined). The same problem arises here because (2.5) is written without imposing that factor markets should clear. In particular, a

competitive equilibrium requires that all firms (and thus the representative firm) maximize profits and factor markets clear (the demand for labor and capital are equal to the supplies of these factor at all points unless the prices of these factors are equal to zero, which is ruled out by Assumption 1). This immediately implies that the representative firm should make zero profits (since otherwise it would wish to hire arbitrarily large amounts of capital and labor exceeding the supplies, which are fixed). It also implies that total demand for labor, L, must be equal to the available supply of labor, $L(t)$ (and similarly, the total demand for capital, K, should equal the total supply, $K(t)$). If this were not the case and $L < L(t)$, then there would be excess supply of labor and the wage would be equal to zero. But this is not consistent with firm maximization, since given Assumption 1 the representative firm would then wish to hire an arbitrarily large amount of labor, again exceeding the supply. This argument, combined with the fact that F is differentiable (Assumption 1), implies that given the supplies of capital and labor at time t, $K(t)$ and $L(t)$, factor prices must satisfy the following familiar conditions equating factor prices to marginal products:[1]

$$w(t) = F_L[K(t), L(t), A(t)], \tag{2.6}$$

and

$$R(t) = F_K[K(t), L(t), A(t)]. \tag{2.7}$$

Next, using Euler's Theorem, Theorem 2.1, verifies that at the prices (2.6) and (2.7) firms (order represents the firm) make zero profits.

PROPOSITION 2.1. *Suppose Assumption 1 holds. Then, in the equilibrium of the Solow growth model, firms make no profits, and in particular,*

$$Y(t) = w(t) L(t) + R(t) K(t).$$

PROOF. This follows immediately from Theorem 2.1 for the case of constant returns to scale ($m = 1$). □

The fact that firms make no profits implies that the ownership of firms does not need to be specified. All we need to know is that firms are profit-maximizing entities.

In addition to these standard assumptions on the production function, the following boundary conditions, referred to as *the Inada conditions*, are often imposed in the analysis of economic growth and macroeconomic equilibria.

ASSUMPTION 2. ***(Inada Conditions)*** *F satisfies the Inada conditions*

$$\begin{aligned} \lim_{K\to 0} F_K(K, L, A) &= \infty \text{ and } \lim_{K\to\infty} F_K(K, L, A) = 0 \text{ for all } L > 0 \text{ and all } A \\ \lim_{L\to 0} F_L(K, L, A) &= \infty \text{ and } \lim_{L\to\infty} F_L(K, L, A) = 0 \text{ for all } K > 0 \text{ and all } A. \end{aligned}$$

[1] An alternative way to derive (2.6) and (2.7) is to consider the cost minimization problem of the representative firm, which takes the form of minimizing $rK + wL$ with respect to K and L subject to the constraint that $F(K, L, A(t)) = Y(t)$ for some level of output $Y(t)$. This problem has a unique solution for any given level of $Y(t)$. Then imposing market clearing, that is, $Y(t) = F(K(t), L(t), A(t))$ with $K(t)$ and $L(t)$ corresponding to the supplies of capital and labor, yields (2.6) and (2.7).

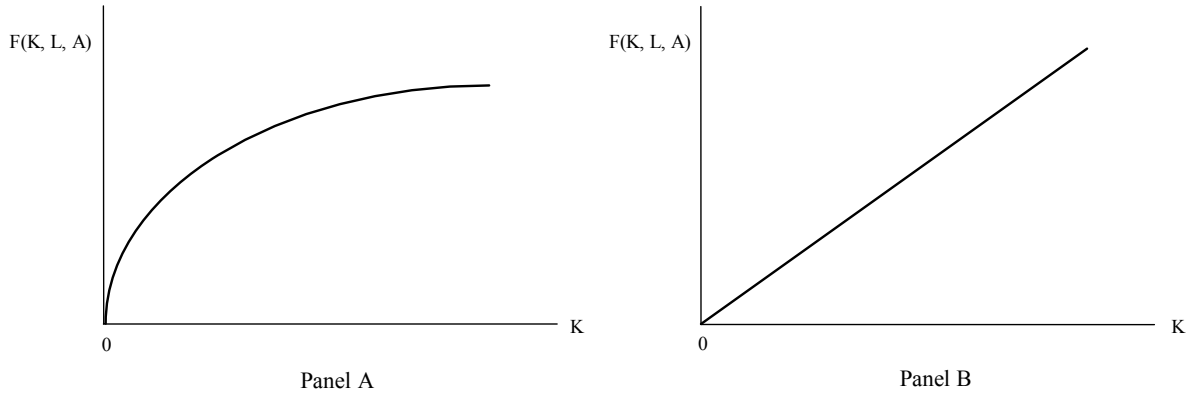

Figure 2.1. Production functions and the marginal product of capital. The example in Panel A satisfies the Inada conditions in Assumption 2, while the example in Panel B does not.

Moreover, $F(0, L, A) = 0$ *for all* L *and* A.

The role of these conditions—especially in ensuring the existence of *interior equilibria*—will become clear in a little while. They imply that the "first units" of capital and labor are highly productive and that when capital or labor are sufficiently abundant, their marginal products are close to zero. The condition that $F(0, L, A) = 0$ for all L and A makes capital an essential input, and is particularly useful in the neoclassical growth model studied in Chapter 8. Figure 2.1 draws the production function $F(K, L, A)$ as a function of K, for given L and A, in two different cases; in Panel A, the Inada conditions are satisfied, while in Panel B, they are not.

I will refer to Assumptions 1 and 2 throughout much of the book.

2.2. The Solow Model in Discrete Time

I next present the dynamics of economic growth in the discrete-time Solow model.

2.2.1. Fundamental Law of Motion of the Solow Model. Recall that K depreciates exponentially at the rate δ, so that the law of motion of the capital stock is given by

$$K(t+1) = (1-\delta) K(t) + I(t), \tag{2.8}$$

where $I(t)$ is investment at time t.

From national income accounting for a closed economy, the total amount of final good in the economy must be either consumed or invested, thus

$$Y(t) = C(t) + I(t), \tag{2.9}$$

where $C(t)$ is consumption.[2] Using (2.1), (2.8) and (2.9), any *feasible* dynamic allocation in this economy must satisfy

$$K(t+1) \leq F[K(t), L(t), A(t)] + (1-\delta)K(t) - C(t)$$

for $t = 0, 1, ...$ The question is to determine the equilibrium dynamic allocation among the set of feasible dynamic allocations. Here the *behavioral rule* that households save a constant fraction of their income simplifies the structure of equilibrium considerably (this is a "behavioral rule" since it is not derived from the maximization of a well-defined utility function). One implication of this assumption is that any welfare comparisons based on the Solow model have to be taken with a grain of salt, since we do not know what the preferences of the individuals are.

Since the economy is closed (and there is no government spending), aggregate investment is equal to savings,

$$S(t) = I(t) = Y(t) - C(t).$$

The assumption that households save a constant fraction $s \in (0,1)$ of their income can be expressed as

$$S(t) = sY(t), \tag{2.10}$$

which, in turn, implies that they consume the remaining $1-s$ fraction of their income, and thus

$$C(t) = (1-s)Y(t). \tag{2.11}$$

In terms of capital market clearing, this implies that the supply of capital for time $t+1$ resulting from households' behavior can be expressed as $K^s(t+1) = (1-\delta)K(t) + S(t) = (1-\delta)K(t) + sY(t)$. Setting supply and demand equal to each other, this implies $K^s(t+1) = K(t+1)$. Moreover, from (2.3), $L(t) = \bar{L}(t)$. Combining these market clearing conditions with (2.1) and (2.8) yields *the fundamental law of motion* of the Solow growth model:

$$K(t+1) = sF[K(t), L(t), A(t)] + (1-\delta)K(t). \tag{2.12}$$

This is a nonlinear *difference equation*. The equilibrium of the Solow growth model is described by this equation together with laws of motion for $L(t)$ (or $\bar{L}(t)$) and $A(t)$.

2.2.2. Definition of Equilibrium. The Solow model is a mixture of an old-style Keynesian model and a modern dynamic macroeconomic model. Households do not optimize when it comes to their savings/consumption decisions. Instead, their behavior is captured by eq.'s (2.10) and (2.11). Nevertheless, firms still maximize and factor markets do clear. Thus it is useful to start defining equilibria in the way that is customary in modern dynamic macro models. Since $L(t) = \bar{L}(t)$ from (2.3), throughout I write the exogenous evolution of labor endowments in terms of $L(t)$ to simplify notation.

[2] In addition, we can introduce government spending $G(t)$ on the right-hand side of (2.9). Government spending does not play a major role in the Solow growth model, thus its introduction is relegated to Exercise 2.7.

DEFINITION 2.2. *In the basic Solow model for a given sequence of* $\{L(t), A(t)\}_{t=0}^{\infty}$ *and an initial capital stock* $K(0)$*, an equilibrium path is a sequence of capital stocks, output levels, consumption levels, wages and rental rates* $\{K(t), Y(t), C(t), w(t), R(t)\}_{t=0}^{\infty}$ *such that* $K(t)$ *satisfies (2.12),* $Y(t)$ *is given by (2.1),* $C(t)$ *is given by (2.11), and* $w(t)$ *and* $R(t)$ *are given by (2.6) and (2.7).*

The most important point to note about Definition 2.2 is that an equilibrium is defined as an entire path of allocations and prices. An economic equilibrium does *not* refer to a static object; it specifies the entire path of behavior of the economy.

2.2.3. Equilibrium Without Population Growth and Technological Progress. It is useful to start with the following assumptions, which will be relaxed later in this chapter:

(1) There is no population growth; total population is constant at some level $L > 0$. Moreover, since individuals supply labor inelastically, this implies $L(t) = L$.
(2) There is no technological progress, so that $A(t) = A$.

Let us define the capital-labor ratio of the economy as

$$k(t) \equiv \frac{K(t)}{L}, \tag{2.13}$$

which is a key object for the analysis. Now using the constant returns to scale assumption, output (income) per capita, $y(t) \equiv Y(t)/L$, can be expressed as

$$\begin{aligned} y(t) &= F\left[\frac{K(t)}{L}, 1, A\right] \\ &\equiv f(k(t)). \end{aligned} \tag{2.14}$$

In other words, with constant returns to scale output per capita is simply a function of the capital-labor ratio. Note that $f(k)$ here depends on A, so I could have written $f(k, A)$. I do not do this to simplify the notation and also because until Section 2.7, there will be no technological progress. Thus, for now A is constant and can be normalized to $A = 1$.[3] From Theorem 2.1, the marginal products of capital and labor (and thus their rental prices) can be reexpressed as:

$$\begin{aligned} R(t) &= f'(k(t)) > 0 \text{ and} \\ w(t) &= f(k(t)) - k(t) f'(k(t)) > 0. \end{aligned} \tag{2.15}$$

The fact that both of these factor prices are positive follows from Assumption 1, which imposed that the first derivatives of F with respect to capital and labor are always positive.

EXAMPLE 2.1. **(The Cobb-Douglas Production Function)** Let us consider the most common example of production function used in macroeconomics, the Cobb-Douglas production function. Let us immediately add the caveat that even though the Cobb-Douglas form is convenient and widely used, it is also very special, and many interesting phenomena

[3]Later, when technological change is taken to be "labor-augmenting," the term A can also be taken out and the per capita production function can be written as $y = Af(k)$, with a slightly different definition of k as effective capital-labor ratio (see, for example, equation (2.46) in Section 2.7).

discussed later in this book are ruled out by this production function. The Cobb-Douglas production function can be written as

$$\begin{aligned} Y(t) &= F[K(t), L(t), A(t)] \\ &= AK(t)^{\alpha} L(t)^{1-\alpha}, 0 < \alpha < 1. \end{aligned} \tag{2.16}$$

It can easily be verified that this production function satisfies Assumptions 1 and 2, including the constant returns to scale feature imposed in Assumption 1. Dividing both sides by $L(t)$, the per capita production function in (2.14) becomes:

$$y(t) = Ak(t)^{\alpha},$$

where $y(t)$ again denotes output per worker and $k(t)$ is capital-labor ratio as defined in (2.13). The representation of factor prices as in (2.15) can also be verified. From the per capita production function representation, in particular eq. (2.15), the rental price of capital can be expressed as

$$\begin{aligned} R(t) &= \frac{\partial Ak(t)^{\alpha}}{\partial k(t)}, \\ &= \alpha Ak(t)^{-(1-\alpha)}. \end{aligned}$$

Alternatively, in terms of the original production function (2.16), the rental price of capital in (2.7) is given by

$$\begin{aligned} R(t) &= \alpha AK(t)^{\alpha-1} L(t)^{1-\alpha} \\ &= \alpha Ak(t)^{-(1-\alpha)}, \end{aligned}$$

which is equal to the previous expression and thus verifies the form of the marginal product given in eq. (2.15). Similarly, from (2.15),

$$\begin{aligned} w(t) &= Ak(t)^{\alpha} - \alpha Ak(t)^{-(1-\alpha)} \times k(t) \\ &= (1-\alpha) AK(t)^{\alpha} L(t)^{-\alpha}, \end{aligned}$$

which verifies the alternative expression for the wage rate in (2.6).

Returning to the analysis with the general production function, the per capita representation of the aggregate production function enables us to divide both sides of (2.12) by L to obtain the following simple difference equation for the evolution of the capital-labor ratio:

$$k(t+1) = sf(k(t)) + (1-\delta) k(t). \tag{2.17}$$

Since this difference equation is derived from (2.12), it also can be referred to as the *equilibrium difference equation* of the Solow model: it describes the equilibrium behavior of the key object of the model, the capital-labor ratio. The other equilibrium quantities can all be obtained from the capital-labor ratio $k(t)$.

At this point, let us also define a *steady-state equilibrium* for this model.

DEFINITION 2.3. *A steady-state equilibrium without technological progress and population growth is an equilibrium path in which $k(t) = k^*$ for all t.*

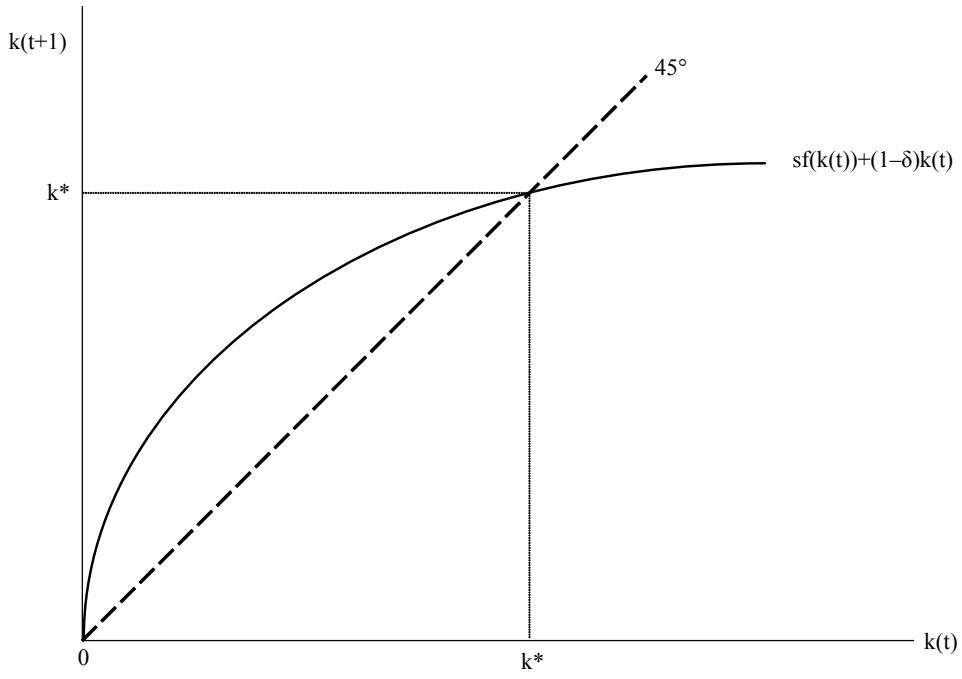

FIGURE 2.2. Determination of the steady-state capital-labor ratio in the Solow model without population growth and technological change.

In a steady-state equilibrium the capital-labor ratio remains constant. Since there is no population growth, this implies that the level of the capital stock will also remain constant. Mathematically, a "steady-state equilibrium" corresponds to a "stationary point" of the equilibrium difference equation (2.17). Most of the models in this book will admit a steady-state equilibrium and the economy will typically tend to this steady-state equilibrium over time. This is also the case for this simple model.

This can be seen by plotting the difference equation that governs the equilibrium behavior of this economy, (2.17), which is done in Figure 2.2. The thick curve represents the right-hand side of eq. (2.17) and the dashed line corresponds to the 45° line. Their (positive) intersection gives the steady-state value of the capital-labor ratio k^*, which satisfies

$$\frac{f(k^*)}{k^*} = \frac{\delta}{s}. \tag{2.18}$$

Notice that in Figure 2.2 there is another intersection between (2.17) and the 45^0 line at $k = 0$. This is because, from Assumption 2, capital is an essential input, and thus $f(0) = 0$. With no production, there are no savings, and the system remains at $k = 0$, making $k = 0$ a steady-state equilibrium. I will ignore this intersection throughout. This is for a number of reasons. First, $k = 0$ is a steady-state equilibrium only when capital is an essential input and $f(0) = 0$, but this assumption can be relaxed without any affect on the rest of the results. If capital were not essential, $f(0)$ will be positive and $k = 0$ will cease to be a

steady-state equilibrium. This is illustrated in Figure 2.3, which draws (2.17) for the case where $f(0) = \varepsilon$ for any $\varepsilon > 0$. Second, as we will see below, this intersection, even when it exists, is an *unstable point*, thus the economy would never travel towards this point starting with $K(0) > 0$. Finally, and most importantly, this intersection has no economic interest for us.[4]

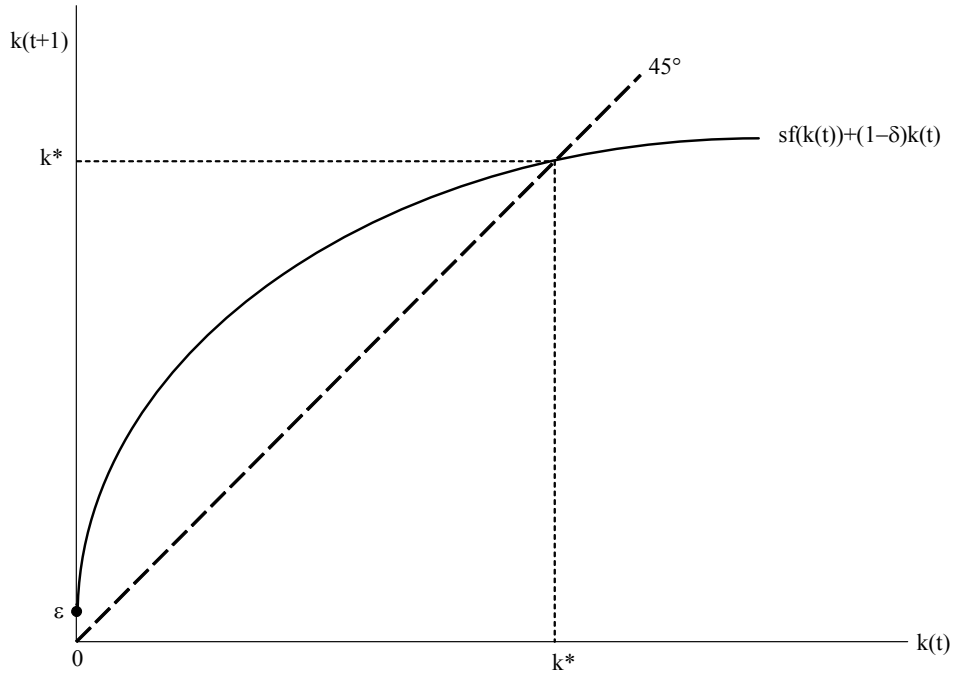

FIGURE 2.3. Unique steady state in the basic Solow model when $f(0) = \varepsilon > 0$.

An alternative visual representation shows the steady state as the intersection between a ray through the origin with slope δ (representing the function δk) and the function $sf(k)$. Figure 2.4 illustrates this and is also useful for two other purposes. First, it depicts the levels of consumption and investment in a single figure. The vertical distance between the horizontal axis and the δk line at the steady-state equilibrium gives the amount of investment per capita at the steady-state equilibrium (equal to δk^*), while the vertical distance between the function $f(k)$ and the δk line at k^* gives the level of consumption per capita. Clearly, the sum of these two terms make up $f(k^*)$. Second, Figure 2.4 also emphasizes that the steady-state equilibrium in the Solow model essentially sets investment, $sf(k)$, equal to the amount of capital that needs to be "replenished", δk. This interpretation will be particularly useful when population growth and technological change are incorporated.

[4]Hakenes and Irmen (2006) show that even with $f(0) = 0$, the Inada conditions imply that in the continuous-time version of the Solow model $k = 0$ may not be a steady-state equilibrium.

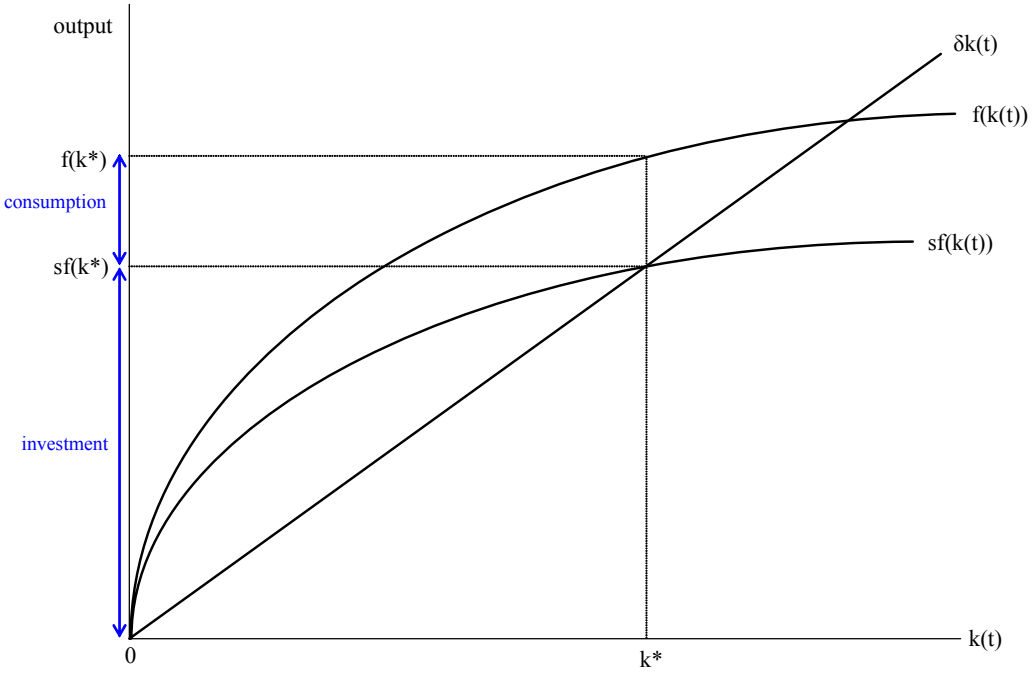

FIGURE 2.4. Investment and consumption in the steady-state equilibrium.

This analysis therefore leads to the following proposition (with the convention that the intersection at $k = 0$ is being ignored even when $f(0) = 0$):

PROPOSITION 2.2. *Consider the basic Solow growth model and suppose that Assumptions 1 and 2 hold. Then, there exists a unique steady state equilibrium where the capital-labor ratio $k^* \in (0, \infty)$ is given by (2.18), per capita output is given by*

$$y^* = f(k^*) \tag{2.19}$$

and per capita consumption is given by

$$c^* = (1 - s) f(k^*). \tag{2.20}$$

PROOF. The preceding argument establishes that any k^* that satisfies (2.18) is a steady state. To establish existence, note that from Assumption 2 (and from L'Hospital's rule, see Theorem A.22 in Appendix Chapter A), $\lim_{k \to 0} f(k)/k = \infty$ and $\lim_{k \to \infty} f(k)/k = 0$. Moreover, $f(k)/k$ is continuous from Assumption 1, so by the Intermediate Value Theorem (Theorem A.3 in Appendix Chapter A) there exists k^* such that (2.18) is satisfied. To see uniqueness, differentiate $f(k)/k$ with respect to k, which gives

$$\frac{\partial [f(k)/k]}{\partial k} = \frac{f'(k) k - f(k)}{k^2} = -\frac{w}{k^2} < 0, \tag{2.21}$$

where the last equality uses (2.15). Since $f(k)/k$ is everywhere (strictly) decreasing, there can only exist a unique value k^* that satisfies (2.18).

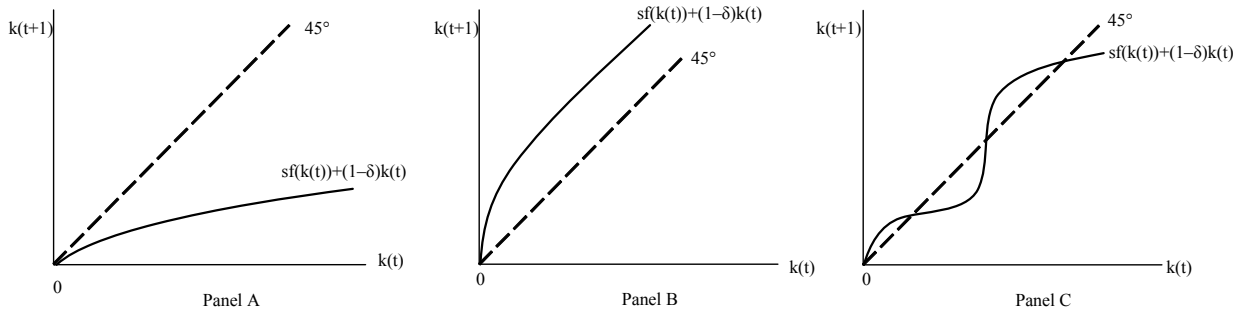

FIGURE 2.5. Examples of nonexistence and nonuniqueness of interior steady states when Assumptions 1 and 2 are not satisfied.

Equations (2.19) and (2.20) then follow by definition. □

Figure 2.5 shows through a series of examples why Assumptions 1 and 2 cannot be dispensed with for the existence and uniqueness results in Proposition 2.2. In the first two panels, the failure of Assumption 2 leads to a situation in which there is no steady state equilibrium with positive activity, while in the third panel, the failure of Assumption 1 leads to non-uniqueness of steady states.

So far the model is very parsimonious: it does not have many parameters and abstracts from many features of the real world in order to focus on the question of interest. An understanding of how cross-country differences in certain parameters translate into differences in growth rates or output levels is essential for our focus. This will be done in the next proposition. But before doing so, let us generalize the production function in one simple way, and assume that

$$f(k) = A\tilde{f}(k),$$

where $A > 0$, so that A is a shift parameter, with greater values corresponding to greater productivity of factors. This type of productivity is referred to as "Hicks-neutral" (see below). For now, it is simply a convenient way of looking at the impact of productivity differences across countries. Since $f(k)$ satisfies the regularity conditions imposed above, so does $\tilde{f}(k)$.

PROPOSITION 2.3. *Suppose Assumptions 1 and 2 hold and $f(k) = A\tilde{f}(k)$. Denote the steady-state level of the capital-labor ratio by $k^*(A, s, \delta)$ and the steady-state level of output by $y^*(A, s, \delta)$ when the underlying parameters are A, s and δ. Then, we have*

$$\frac{\partial k^*(A,s,\delta)}{\partial A} > 0, \frac{\partial k^*(A,s,\delta)}{\partial s} > 0 \text{ and } \frac{\partial k^*(A,s,\delta)}{\partial \delta} < 0$$
$$\frac{\partial y^*(A,s,\delta)}{\partial A} > 0, \frac{\partial y^*(A,s,\delta)}{\partial s} > 0 \text{ and } \frac{\partial y^*(A,s,\delta)}{\partial \delta} < 0.$$

PROOF. The proof follows immediately by writing

$$\frac{\tilde{f}(k^*)}{k^*} = \frac{\delta}{As},$$

which holds for an open set of values of k^*, A, s and δ. Now apply the Implicit Function Theorem (Theorem A.26 in Appendix Chapter A) to obtain the results. For example,

$$\frac{\partial k^*}{\partial s} = \frac{\delta (k^*)^2}{As^2 w^*} > 0$$

where $w^* = f(k^*) - k^* f'(k^*) > 0$. The other results follow similarly. □

Therefore, countries with higher saving rates and better technologies will have higher capital-labor ratios and will be richer. Those with greater (technological) depreciation, will tend to have lower capital-labor ratios and will be poorer. All of the results in Proposition 2.3 are intuitive and they provide us with a first glimpse of the potential determinants of the capital-labor ratios and output levels across countries.

The same comparative statics with respect to A and δ immediately apply to c^* as well. However, it is straightforward to see that c^* will not be monotone in the saving rate (think, for example, of the extreme case where $s = 1$). In fact, there will exist a specific level of the saving rate, s_{gold}, referred to as the "golden rule" saving rate, which maximizes the steady-state level of consumption. Since we are treating the saving rate as an exogenous parameter and have not specified the objective function of households yet, we cannot say whether the golden rule saving rate is "better" than some other saving rate. It is nevertheless interesting to characterize what this golden rule saving rate corresponds to.

To do this, let us first write the steady state relationship between c^* and s and suppress the other parameters:

$$\begin{aligned} c^*(s) &= (1-s) f(k^*(s)), \\ &= f(k^*(s)) - \delta k^*(s), \end{aligned}$$

where the second equality exploits the fact that in steady state $sf(k) = \delta k$. Now differentiating this second line with respect to s (again using the Implicit Function Theorem), we have

$$\frac{\partial c^*(s)}{\partial s} = \left[f'(k^*(s)) - \delta \right] \frac{\partial k^*}{\partial s}. \tag{2.22}$$

Let us define the golden rule saving rate s_{gold} to be such that $\partial c^*(s_{gold})/\partial s = 0$. The corresponding steady-state golden rule capital stock is defined as k^*_{gold}. These quantities and the relationship between consumption and the saving rate are plotted in Figure 2.6.

The next proposition shows that s_{gold} and k^*_{gold} are uniquely defined and the latter satisfies (2.23).

PROPOSITION 2.4. *In the basic Solow growth model, the highest level of steady-state consumption is reached for s_{gold}, with the corresponding steady-state capital level k^*_{gold} such that*

$$f'\left(k^*_{gold}\right) = \delta. \tag{2.23}$$

PROOF. By definition $\partial c^*(s_{gold})/\partial s = 0$. From Proposition 2.3, $\partial k^*/\partial s > 0$, thus (2.22) can be equal to zero only when $f'(k^*(s_{gold})) = \delta$. Moreover, when $f'(k^*(s_{gold})) = \delta$, it can be verified that $\partial^2 c^*(s_{gold})/\partial s^2 < 0$, so $f'(k^*(s_{gold})) = \delta$ indeed corresponds to a local

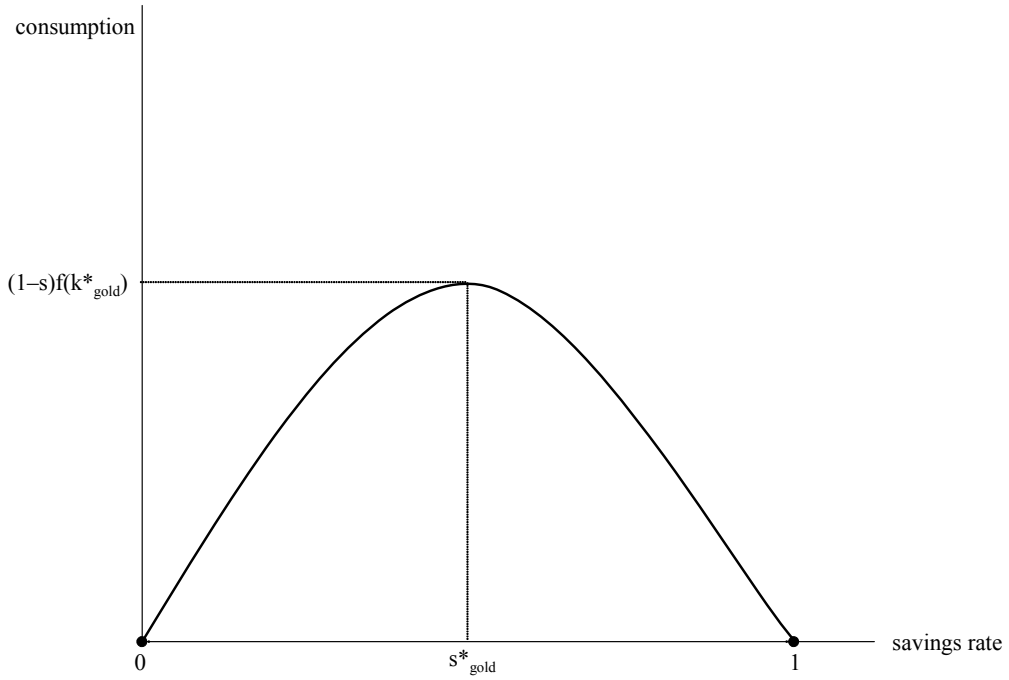

Figure 2.6. The "golden rule" level of savings rate, which maximizes steady-state consumption.

maximum. That $f'(k^*(s_{gold})) = \delta$ also yields the global maximum is a consequence of the following observations: for all $s \in [0,1]$, we have $\partial k^*/\partial s > 0$ and moreover, when $s < s_{gold}$, $f'(k^*(s)) - \delta > 0$ by the concavity of f, so $\partial c^*(s)/\partial s > 0$ for all $s < s_{gold}$, and by the converse argument, $\partial c^*(s)/\partial s < 0$ for all $s > s_{gold}$. Therefore, only s_{gold} satisfies $f'(k^*(s)) = \delta$ and gives the unique global maximum of consumption per capita. □

In other words, there exists a unique saving rate, s_{gold}, and also a unique corresponding capital-labor ratio, k^*_{gold}, maximizing the level of steady-state consumption. When the economy is below k^*_{gold}, the higher saving rate will increase consumption, whereas when the economy is above k^*_{gold}, steady-state consumption can be increased by saving less. In the latter case, lower savings translate into higher consumption because the capital-labor ratio of the economy is too high; individuals are investing too much and not consuming enough. This is the essence of what is referred to as *dynamic inefficiency*, which will be discussed in greater detail in Chapter 9. For now, recall that there is no explicit utility function here, so statements about "inefficiency" must the considered with caution. In fact, the reason why this type of dynamic inefficiency will not arise when consumption-saving decisions are endogenized may already be apparent to many of you.

2.3. Transitional Dynamics in the Discrete Time Solow Model

Proposition 2.2 establishes the existence of a unique steady-state equilibrium (with positive activity). Recall that an *equilibrium path* does not refer simply to the steady state, but to the entire path of capital stock, output, consumption and factor prices. This is an important point to bear in mind, especially since the term "equilibrium" is used differently in economics than in physical sciences. Typically, in engineering and physical sciences, an equilibrium refers to a point of rest of a dynamical system, thus to what I have so far referred to as *the steady-state equilibrium*. One may then be tempted to say that the system is in "disequilibrium" when it is away from the steady state. However, in economics, the non-steady-state behavior of an economy is also governed by market clearing and optimizing behavior of households and firms. Most economies spend much of their time in non-steady-state situations. Thus we are typically interested in the entire dynamic equilibrium path of the economy, not just its steady state.

To determine what the equilibrium path of our simple economy looks like we need to study the "transitional dynamics" of the equilibrium difference equation (2.17) starting from an arbitrary capital-labor ratio, $k(0) > 0$. Of special interest is the answer to the question of whether the economy will tend to this steady state starting from an arbitrary capital-labor ratio, and how it will behave along the transition path. It is important to consider an arbitrary capital-labor ratio, since, as noted above, the total amount of capital at the beginning of the economy, $K(0) > 0$, is taken as a state variable, while for now, the supply of labor L is fixed. Therefore, at time $t = 0$, the economy starts with $k(0) = K(0)/L$ as its initial value and then follows the law of motion given by the difference equation (2.17). Thus the question is whether the difference equation (2.17) will take us to the unique steady state starting from an arbitrary initial capital-labor ratio.

Before doing this, recall some definitions and key results from the theory of dynamical systems. Appendix Chapter B provides more details and a number of further results. Consider the nonlinear system of autonomous difference equations,

$$\mathbf{x}(t+1) = \mathbf{G}(\mathbf{x}(t)), \tag{2.24}$$

where $\mathbf{x}(t) \in \mathbb{R}^n$ and $\mathbf{G} : \mathbb{R}^n \to \mathbb{R}^n$ (where $n \in \mathbb{R}$). Let $\mathbf{x}^*$ be a *fixed point* of the mapping $\mathbf{G}(\cdot)$, that is,

$$\mathbf{x}^* = \mathbf{G}(\mathbf{x}^*).$$

I will refer to $\mathbf{x}^*$ as a *steady state* of the difference equation (2.24).[5] The relevant notion of stability is introduced in the next definition.

Definition 2.4. *A steady state $\mathbf{x}^*$ is (locally)* ***asymptotically stable*** *if there exists an open set $B(\mathbf{x}^*)$ containing $\mathbf{x}^*$ such that for any solution $\{\mathbf{x}(t)\}_{t=0}^{\infty}$ to (2.24) with $\mathbf{x}(0) \in B(\mathbf{x}^*)$, $\mathbf{x}(t) \to \mathbf{x}^*$. Moreover, $\mathbf{x}^*$ is* ***globally asymptotically stable*** *if for all $\mathbf{x}(0) \in \mathbb{R}^n$, for any solution $\{\mathbf{x}(t)\}_{t=0}^{\infty}$, $\mathbf{x}(t) \to \mathbf{x}^*$.*

[5]Various other terms are used to describe $\mathbf{x}^*$, for example, "equilibrium point" or "critical point". Since these other terms have different meanings in economics, I will refer to $\mathbf{x}^*$ as a steady state throughout.

The next theorem provides the main results on the stability properties of systems of linear difference equations. The following theorems are special cases of the results presented in Appendix Chapter B.

Theorem 2.2. ***(Stability for Systems of Linear Difference Equations)*** *Consider the following linear difference equation system*

$$\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \mathbf{b} \tag{2.25}$$

with initial value $\mathbf{x}(0)$*, where* $\mathbf{x}(t) \in \mathbb{R}^n$ *for all* t*,* $\mathbf{A}$ *is an* $n \times n$ *matrix and* $\mathbf{b}$ *is a* $n \times 1$ *column vector. Let* $\mathbf{x}^*$ *be the steady state of the difference equation given by* $\mathbf{A}\mathbf{x}^* + \mathbf{b} = \mathbf{x}^*$*. Suppose that all of the eigenvalues of* $\mathbf{A}$ *are strictly inside the unit circle in the complex plane. Then, the steady state of the difference equation (2.25),* $\mathbf{x}^*$*, is* globally asymptotically stable*, in the sense that starting from any* $\mathbf{x}(0) \in \mathbb{R}^n$*, the unique solution* $\{\mathbf{x}(t)\}_{t=0}^{\infty}$ *satisfies* $\mathbf{x}(t) \to \mathbf{x}^*$*.*

Unfortunately, much less can be said about nonlinear systems, but the following is a standard *local* stability result.

Theorem 2.3. ***(Local Stability for Systems of Nonlinear Difference Equations)*** *Consider the following nonlinear autonomous system*

$$\mathbf{x}(t+1) = \mathbf{G}(\mathbf{x}(t)) \tag{2.26}$$

with initial value $\mathbf{x}(0)$*, where* $\mathbf{G}: \mathbb{R}^n \to \mathbb{R}^n$*. Let* $\mathbf{x}^*$ *be a steady state of this system, i.e.,* $\mathbf{G}(\mathbf{x}^*) = \mathbf{x}^*$*, and suppose that* $\mathbf{G}$ *is differentiable at* $\mathbf{x}^*$*. Define*

$$\mathbf{A} \equiv D\mathbf{G}(\mathbf{x}^*),$$

where $D\mathbf{G}$ *denotes the matrix of partial derivatives (Jacobian) of* $\mathbf{G}$*. Suppose that all of the eigenvalues of* $\mathbf{A}$ *are strictly inside the unit circle. Then, the steady state of the difference equation (2.26),* $\mathbf{x}^*$*, is* locally asymptotically stable*, in the sense that there exists an open neighborhood of* $\mathbf{x}^*$*,* $\mathbf{B}(\mathbf{x}^*) \subset \mathbb{R}^n$ *such that starting from any* $\mathbf{x}(0) \in \mathbf{B}(\mathbf{x}^*)$*,* $\mathbf{x}(t) \to \mathbf{x}^*$*.*

An immediate corollary of Theorem 2.3 is the following useful result:

Corollary 2.1.
(1) *Let* $x(t), a, b \in \mathbb{R}$*, then the unique steady state of the linear difference equation* $x(t+1) = ax(t) + b$ *is globally asymptotically stable (in the sense that* $x(t) \to x^* = b/(1-a)$*) if* $|a| < 1$*.*
(2) *Let* $g: \mathbb{R} \to \mathbb{R}$ *be differentiable in the neighborhood of the steady state* x^**, defined by* $g(x^*) = x^*$*, and suppose that* $|g'(x^*)| < 1$*. Then, the steady state* x^* *of the nonlinear difference equation* $x(t+1) = g(x(t))$ *is locally asymptotically stable. Moreover, if* $|g'(x)| < 1$ *for all* $x \in \mathbb{R}$*, then* x^* *is globally asymptotically stable.*

PROOF. The first part follows immediately from Theorem 2.2. The local stability of g in the second part follows from Theorem 2.3. Global stability follows since

$$\begin{aligned} |x(t+1)-x^*| &= |g(x(t))-g(x^*)| \\ &= \left|\int_{x^*}^{x(t)} g'(x)dx\right| \\ &< |x(t)-x^*|, \end{aligned}$$

where the last inequality uses the hypothesis that $|g'(x)|<1$ for all $x\in\mathbb{R}$. □

We can now apply Corollary 2.1 to the equilibrium difference equation of the Solow model, (2.17):

PROPOSITION 2.5. *Suppose that Assumptions 1 and 2 hold, then the steady-state equilibrium of the Solow growth model described by the difference equation (2.17) is globally asymptotically stable, and starting from any $k(0)>0$, $k(t)$ monotonically converges to k^*.*

PROOF. Let $g(k)\equiv sf(k)+(1-\delta)k$. First observe that $g'(k)$ exists and is always strictly positive, i.e., $g'(k)>0$ for all k. Next, from (2.17),

$$k(t+1)=g(k(t)), \tag{2.27}$$

with a unique steady state at k^*. From (2.18), the steady-state capital k^* satisfies $\delta k^*=sf(k^*)$, or

$$k^*=g(k^*). \tag{2.28}$$

Now recall that $f(\cdot)$ is concave and differentiable from Assumption 1 and satisfies $f(0)\geq 0$ from Assumption 2. For any strictly concave differentiable function, we have

$$f(k)>f(0)+kf'(k)\geq kf'(k), \tag{2.29}$$

where the second inequality uses the fact that $f(0)\geq 0$. Since (2.29) implies that $\delta=sf(k^*)/k^*>sf'(k^*)$, we have $g'(k^*)=sf'(k^*)+1-\delta<1$. Therefore,

$$g'(k^*)\in(0,1).$$

Corollary 2.1 then establishes local asymptotic stability.

To prove global stability, note that for all $k(t)\in(0,k^*)$,

$$\begin{aligned} k(t+1)-k^* &= g(k(t))-g(k^*) \\ &= -\int_{k(t)}^{k^*} g'(k)\,dk, \\ &< 0 \end{aligned}$$

where the first line follows by subtracting (2.28) from (2.27), the second line uses the Fundamental Theorem of Calculus (Theorem B.2), and the third line follows from the observation

that $g'(k) > 0$ for all k. Next, (2.17) also implies

$$\begin{aligned}\frac{k(t+1)-k(t)}{k(t)} &= s\frac{f(k(t))}{k(t)}-\delta \\ &> s\frac{f(k^*)}{k^*}-\delta \\ &= 0,\end{aligned}$$

where the second line uses the fact that $f(k)/k$ is decreasing in k (from (2.29) above) and the last line uses the definition of k^*. These two arguments together establish that for all $k(t) \in (0, k^*)$, $k(t+1) \in (k(t), k^*)$. An identical argument implies that for all $k(t) > k^*$, $k(t+1) \in (k^*, k(t))$. Therefore, $\{k(t)\}_{t=0}^{\infty}$ monotonically converges to k^* and is globally stable. □

This stability result can be seen diagrammatically in Figure 2.7. Starting from initial capital stock $k(0) > 0$, which is below the steady-state level k^*, the economy grows towards k^* and experiences *capital deepening*—meaning that the capital-labor ratio increases. Together with capital deepening comes growth of per capita income. If, instead, the economy were to start with $k'(0) > k^*$, it would reach the steady state by decumulating capital and contracting (that is, by experiencing negative growth).

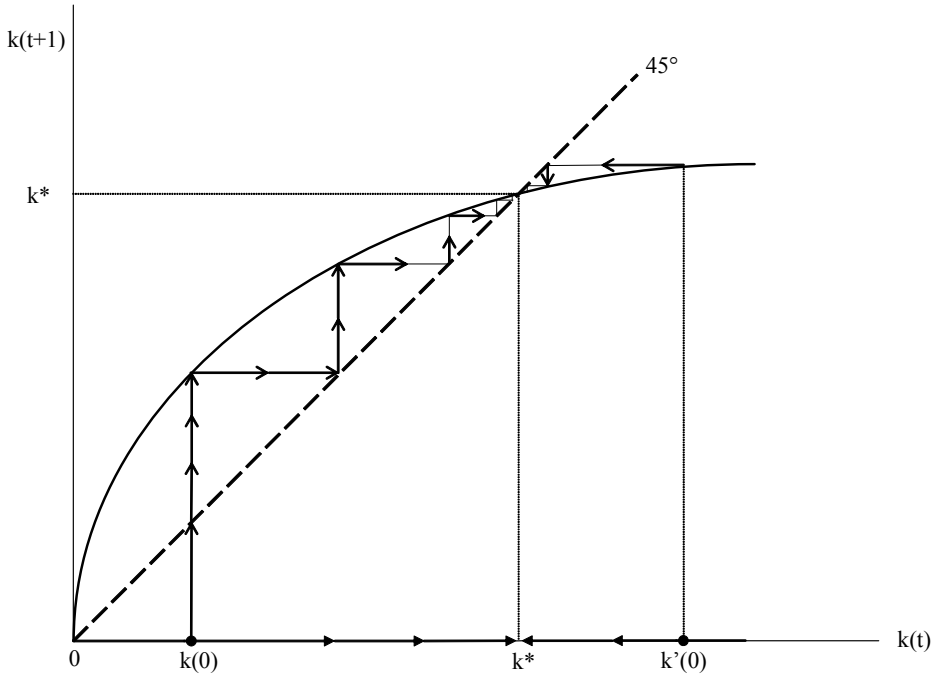

FIGURE 2.7. Transitional dynamics in the basic Solow model.

The following proposition is an immediate corollary of Proposition 2.5:

Proposition 2.6. *Suppose that Assumptions 1 and 2 hold, and* $k(0) < k^*$*, then* $\{w(t)\}_{t=0}^{\infty}$ *is an increasing sequence and* $\{R(t)\}_{t=0}^{\infty}$ *is a decreasing sequence. If* $k(0) > k^*$*, the opposite results apply.*

Proof. See Exercise 2.9. □

Recall that when the economy starts with too little capital relative to its labor supply, the capital-labor ratio will increase. This implies that the marginal product of capital will fall due to diminishing returns to capital and the wage rate will increase. Conversely, if it starts with too much capital, it will decumulate capital, and in the process the wage rate will decline and the rate of return to capital will increase.

The analysis has established that the Solow growth model has a number of nice properties; unique steady state, asymptotic stability, and finally, simple and intuitive comparative statics. Yet, so far, it has no growth. The steady state is the point at which there is no growth in the capital-labor ratio, no more capital deepening and no growth in output per capita. Consequently, the basic Solow model (without technological progress) can only generate economic growth along the transition path when the economy starts with $k(0) < k^*$. However, this growth is *not sustained*: it slows down over time and eventually comes to an end. Section 2.7 will show that the Solow model can incorporate economic growth by allowing *exogenous* technological change. Before doing this, it is useful to look at the relationship between the discrete-time and continuous-time formulations.

2.4. The Solow Model in Continuous Time

2.4.1. From Difference to Differential Equations. Recall that the time periods could refer to days, weeks, months or years. In some sense, the time unit is not important. This suggests that perhaps it may be more convenient to look at dynamics by making the time unit as small as possible, that is, by going to continuous time. While much of modern macroeconomics (outside of growth theory) uses discrete-time models, many growth models are formulated in continuous time. The continuous-time setup has a number of advantages, since some pathological results of discrete-time models disappear in continuous time (see Exercise 2.16). Moreover, continuous-time models have more flexibility in the analysis of dynamics and allow explicit-form solutions in a wider set of circumstances. These considerations motivate a detailed study of both the discrete-time and the continuous-time versions of the basic models.

Let us start with a simple difference equation

$$x(t+1) - x(t) = g(x(t)). \tag{2.30}$$

This equation states that between time t and $t+1$, the absolute growth in x is given by $g(x(t))$. Let us now consider the following approximation

$$x(t+\Delta t) - x(t) \simeq \Delta t \cdot g(x(t)),$$

for any $\Delta t \in [0,1]$. When $\Delta t = 0$, this equation is just an identity. When $\Delta t = 1$, it gives (2.30). In-between it is a linear approximation, which should not be too bad if the distance between t and $t+1$ is not very large, so that $g(x) \simeq g(x(t))$ for all $x \in [x(t), x(t+1)]$ (however, you should also convince yourself that this approximation could in fact be quite bad if you take a very nonlinear function g, for which the behavior changes significantly between $x(t)$ and $x(t+1)$). Now divide both sides of this equation by Δt, and take limits to obtain

$$\lim_{\Delta t \to 0} \frac{x(t+\Delta t) - x(t)}{\Delta t} = \dot{x}(t) \simeq g(x(t)), \tag{2.31}$$

where throughout the book I use the "dot" notation

$$\dot{x}(t) \equiv \frac{\partial x(t)}{\partial t}$$

to denote time derivatives. Equation (2.31) is a differential equation representing the same dynamics as the difference equation (2.30) for the case in which the distance between t and $t+1$ is "small".

2.4.2. The Fundamental Equation of the Solow Model in Continuous Time. We can now repeat all of the analysis so far using the continuous-time representation. Nothing has changed on the production side, so we continue to have (2.6) and (2.7) as the factor prices, but now these refer to instantaneous rental rates. For example, $w(t)$ is the flow of wages that workers receive at time t. Savings are again given by

$$S(t) = sY(t),$$

while consumption is still given by (2.11) above.

Let us also introduce population growth into this model, and assume that the labor force $L(t)$ grows proportionally, that is,

$$L(t) = \exp(nt) L(0). \tag{2.32}$$

The purpose of doing so is that in many of the classical analyses of economic growth, population growth plays an important role, so it is useful to see how it affects things here. There is still no technological progress.

Recall that

$$k(t) \equiv \frac{K(t)}{L(t)},$$

which implies that

$$\begin{aligned} \frac{\dot{k}(t)}{k(t)} &= \frac{\dot{K}(t)}{K(t)} - \frac{\dot{L}(t)}{L(t)}, \\ &= \frac{\dot{K}(t)}{K(t)} - n, \end{aligned}$$

where I used the fact that, from (2.32), $\dot{L}(t)/L(t) = n$. From the limiting argument leading to eq. (2.31) in the previous subsection, the law of motion of the capital stock is given by

$$\dot{K}(t) = sF[K(t), L(t), A(t)] - \delta K(t).$$

Using the definition of $k(t)$ as the capital-labor ratio and the constant returns to scale properties of the production function, the fundamental law of motion of the Solow model in continuous time is obtained as

$$\frac{\dot{k}(t)}{k(t)} = s\frac{f(k(t))}{k(t)} - (n+\delta), \tag{2.33}$$

where, following usual practice, I have transformed the left-hand side to the proportional change in the capital-labor ratio by dividing both sides by $k(t)$.[6]

Definition 2.5. *In the basic Solow model in continuous time with population growth at the rate n, no technological progress and an initial capital stock $K(0)$, an equilibrium path is given by processes (sequences) of capital stocks, labor, output levels, consumption levels, wages and rental rates $[K(t), L(t), Y(t), C(t), w(t), R(t)]_{t=0}^{\infty}$ such that $L(t)$ satisfies (2.32), $k(t) \equiv K(t)/L(t)$ satisfies (2.33), $Y(t)$ is given by (2.1), $C(t)$ is given by (2.11), and $w(t)$ and $R(t)$ are given by (2.6) and (2.7).*

As before, a *steady-state* equilibrium involves $k(t)$ remaining constant at some level k^*.

It is easy to verify that the equilibrium differential equation (2.33) has a unique steady state at k^*, which is given by a slight modification of (2.18) above to incorporate population growth:

$$\frac{f(k^*)}{k^*} = \frac{n+\delta}{s}. \tag{2.34}$$

In other words, going from discrete to continuous time has not changed any of the basic economic features of the model, and again the steady state can be plotted in a diagram similar to the one used above (now with the population growth rate featuring in there as well). This is done in Figure 2.8, which also highlights that the logic of the steady state is the same with population growth as it was without population growth. The amount of investment, $sf(k)$, is used to replenish the capital-labor ratio, but now there are two reasons for replenishments. The capital stock depreciates, again exponentially, at the flow rate δ. In addition, the capital stock must also increase as population grows in order to maintain the capital-labor ratio at a constant level. The amount of capital that needs to be replenished is therefore $(n+\delta)k$.

Proposition 2.7. *Consider the basic Solow growth model in continuous time and suppose that Assumptions 1 and 2 hold. Then, there exists a unique steady state equilibrium where the capital-labor ratio is equal to $k^* \in (0, \infty)$ and is given by (2.34), per capita output is given by*

$$y^* = f(k^*),$$

[6]Throughout I adopt the notation $[x(t)]_{t=0}^{\infty}$ to denote the continuous-time path of variable $x(t)$. An alternative notation often used in the literature is $(x(t); t \geq 0)$. I prefer the former both because it is slightly more compact and also because it is more similar to the discrete-time notation for the time path of a variable, $\{x(t)\}_{t=0}^{\infty}$. When referring to $[x(t)]_{t=0}^{\infty}$, I use the terms process and sequence interchangeably.

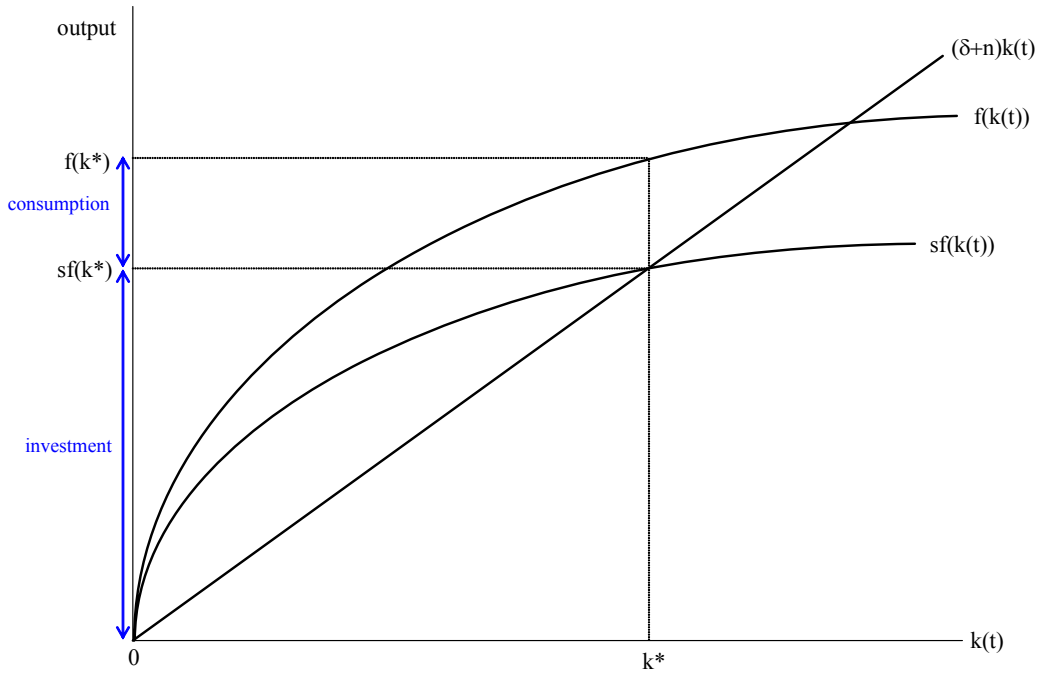

FIGURE 2.8. Investment and consumption in the state-state equilibrium with population growth.

and per capita consumption is given by

$$c^* = (1-s) f(k^*).$$

PROOF. See Exercise 2.5. □

Moreover, again defining $f(k) = A\tilde{f}(k)$:

PROPOSITION 2.8. *Suppose Assumptions 1 and 2 hold and $f(k) = A\tilde{f}(k)$. Denote the steady-state equilibrium level of the capital-labor ratio by $k^*(A,s,\delta,n)$ and the steady-state level of output by $y^*(A,s,\delta,n)$ when the underlying parameters are given by A, s and δ. Then, we have*

$$\frac{\partial k^*(A,s,\delta,n)}{\partial A} > 0, \frac{\partial k^*(A,s,\delta,n)}{\partial s} > 0, \frac{\partial k^*(A,s,\delta,n)}{\partial \delta} < 0 \text{ and } \frac{\partial k^*(A,s,\delta,n)}{\partial n} < 0$$
$$\frac{\partial y^*(A,s,\delta,n)}{\partial A} > 0, \frac{\partial y^*(A,s,\delta,n)}{\partial s} > 0, \frac{\partial y^*(A,s,\delta,n)}{\partial \delta} < 0 \text{and } \frac{\partial y^*(A,s,\delta,n)}{\partial n} < 0.$$

PROOF. See Exercise 2.6. □

The new result relative to the earlier comparative static proposition is that now a higher population growth rate, n, also reduces the capital-labor ratio and output per capita. The reason for this is simple: a higher population growth rate means there is more labor to use the existing amount of capital, which only accumulates slowly, and consequently the equilibrium capital-labor ratio ends up lower. This result implies that countries with higher population growth rates will have lower incomes per person (or per worker).

2.5. Transitional Dynamics in the Continuous Time Solow Model

The analysis of transitional dynamics and stability with continuous time yields similar results to those in Section 2.3, but the analysis will be slightly simpler. Let us first recall the basic results on the stability of systems of differential equations. Once again, further details are contained in Appendix Chapter B.

Theorem 2.4. ***(Stability of Linear Differential Equations)*** *Consider the following linear differential equation system*

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{b} \tag{2.35}$$

with initial value $\mathbf{x}(0)$*, where* $\mathbf{x}(t) \in \mathbb{R}^n$ *for all* t*,* $\mathbf{A}$ *is an* $n \times n$ *matrix and* $\mathbf{b}$ *is a* $n \times 1$ *column vector. Let* $\mathbf{x}^*$ *be the steady state of the system given by* $\mathbf{A}\mathbf{x}^* + \mathbf{b} = 0$*. Suppose that all of the eigenvalues of* $\mathbf{A}$ *have negative real parts. Then, the steady state of the differential equation (2.35)* $\mathbf{x}^*$ *is globally asymptotically stable, in the sense that starting from any* $\mathbf{x}(0) \in \mathbb{R}^n$*,* $\mathbf{x}(t) \to \mathbf{x}^*$*.*

Theorem 2.5. ***(Local Stability of Nonlinear Differential Equations)*** *Consider the following nonlinear autonomous differential equation*

$$\dot{\mathbf{x}}(t) = \mathbf{G}[\mathbf{x}(t)] \tag{2.36}$$

with initial value $\mathbf{x}(0)$*, where* $\mathbf{G} : \mathbb{R}^n \to \mathbb{R}^n$*. Let* $\mathbf{x}^*$ *be a steady state of this system, i.e.,* $\mathbf{G}(\mathbf{x}^*) = 0$*, and suppose that* $\mathbf{G}$ *is differentiable at* $\mathbf{x}^*$*. Define*

$$\mathbf{A} \equiv D\mathbf{G}(\mathbf{x}^*),$$

and suppose that all of the eigenvalues of $\mathbf{A}$ *have negative real parts. Then, the steady state of the differential equation (2.36),* $\mathbf{x}^*$*, is locally asymptotically stable, in the sense that there exists an open neighborhood of* $\mathbf{x}^*$*,* $\mathbf{B}(\mathbf{x}^*) \subset \mathbb{R}^n$ *such that starting from any* $\mathbf{x}(0) \in \mathbf{B}(\mathbf{x}^*)$*, we have* $\mathbf{x}(t) \to \mathbf{x}^*$*.*

Once again an immediate corollary is:

Corollary 2.2.
(1) *Let* $x(t) \in \mathbb{R}$*. Then, the steady state of the linear difference equation* $\dot{x}(t) = ax(t)$ *is globally asymptotically stable (in the sense that* $x(t) \to 0$*) if* $a < 0$*.*
(2) *Let* $g : \mathbb{R} \to \mathbb{R}$ *be differentiable in the neighborhood of the steady state* x^* *defined by* $g(x^*) = 0$ *and suppose that* $g'(x^*) < 0$*. Then, the steady state of the nonlinear differential equation* $\dot{x}(t) = g(x(t))$*,* x^**, is locally asymptotically stable.*

(3) *Let $g : \mathbb{R} \to \mathbb{R}$ and suppose that there exists a unique x^* such that $g(x^*) = 0$. Moreover, suppose $g(x) < 0$ for all $x > x^*$ and $g(x) > 0$ for all $x < x^*$. Then, the steady state of the nonlinear differential equation $\dot{x}(t) = g(x(t))$, x^*, is globally asymptotically stable, that is, starting with any $x(0)$, $x(t) \to x^*$.*

PROOF. See Exercise 2.10. □

Notice that the equivalent of the Part 3 of Corollary 2.2 is not true in discrete time, and this will be illustrated in Exercise 2.16.

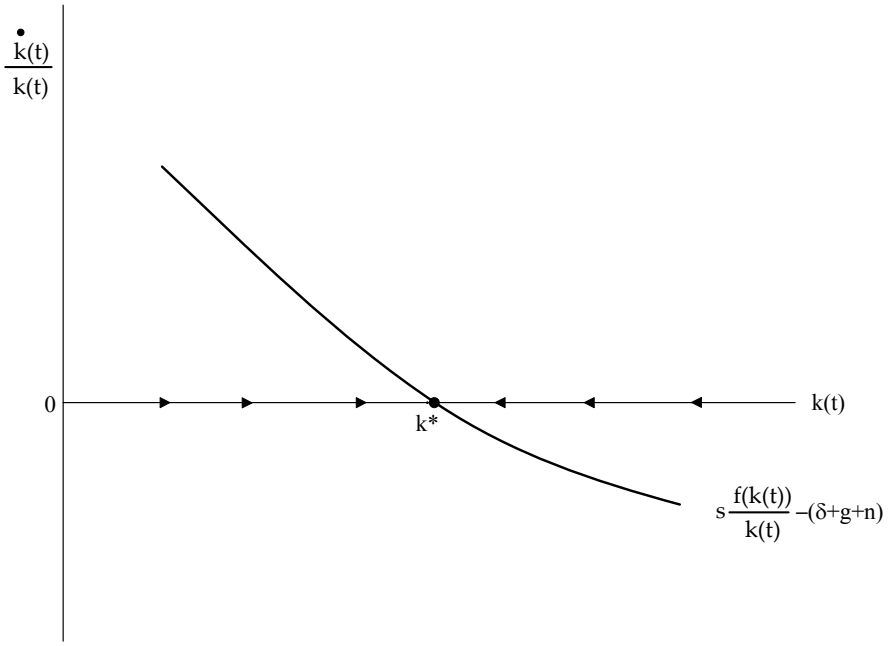

FIGURE 2.9. Dynamics of the capital-labor ratio in the basic Solow model.

In view of these results, Proposition 2.5 immediately generalizes:

PROPOSITION 2.9. *Suppose that Assumptions 1 and 2 hold, then the basic Solow growth model in continuous time with constant population growth and no technological change is globally asymptotically stable, and starting from any $k(0) > 0$, $k(t)$ monotonically converges to k^*.*

PROOF. The proof of stability is now simpler and follows immediately from Part 3 of Corollary 2.2 Part (3) by noting that whenever $k < k^*$, $sf(k) - (n+\delta)k > 0$ and whenever $k > k^*$, $sf(k) - (n+\delta)k < 0$. □

Figure 2.9 shows the analysis of stability diagrammatically. The figure plots the right-hand side of (2.33) and makes it clear that whenever $k < k^*$, $\dot{k} > 0$ and whenever $k > k^*$, $\dot{k} < 0$, so that the capital-labor ratio monotonically converges to the steady-state value k^*.

Example 2.2. **(Dynamics with the Cobb-Douglas Production Function)** Let us return to the Cobb-Douglas production function introduced in Example 2.1

$$F[K, L, A] = AK^{\alpha}L^{1-\alpha} \text{ with } 0 < \alpha < 1.$$

As noted above, the Cobb-Douglas production function is special, mainly because it has an elasticity of substitution between capital and labor equal to 1. Recall that for a homothetic production function $F(K, L)$, the elasticity of substitution is defined by

$$\sigma \equiv -\left[\frac{\partial \ln (F_K/F_L)}{\partial \ln (K/L)}\right]^{-1}, \tag{2.37}$$

where F_K and F_L denote the marginal products of capital and labor. In addition, F is required to be homothetic, so that F_K/F_L is only a function of K/L. For the Cobb-Douglas production function $F_K/F_L = \alpha L/((1-\alpha)K)$, thus $\sigma = 1$. This feature implies that when the production function is Cobb-Douglas and factor markets are competitive, equilibrium factor shares will be constant regardless of the capital-labor ratio. In particular, the share of capital in national income is

$$\begin{aligned}
\alpha_K(t) &= \frac{R(t)K(t)}{Y(t)} \\
&= \frac{F_K(K(t), L(t))K(t)}{Y(t)} \\
&= \frac{\alpha A K(t)^{\alpha-1} L(t)^{1-\alpha} K(t)}{AK(t)^{\alpha} L(t)^{1-\alpha}} \\
&= \alpha.
\end{aligned}$$

Similarly, the share of labor is $\alpha_L(t) = 1 - \alpha$. The reason for this is that with an elasticity of substitution equal to 1, as capital increases, its marginal product decreases proportionally, leaving the capital share (the amount of capital times its marginal product) constant.

Recall that with the Cobb-Douglas technology, the per capita production function takes the form $f(k) = Ak^{\alpha}$, so the steady state is given again from (2.34) (with population growth at the rate n) as

$$A(k^*)^{\alpha-1} = \frac{n+\delta}{s}$$

or

$$k^* = \left(\frac{sA}{n+\delta}\right)^{\frac{1}{1-\alpha}},$$

which is a simple expression for the steady-state capital-labor ratio. It immediately follows that k^* is increasing in s and A, and decreasing in n and δ (and these results are naturally consistent with those in Proposition 2.8). In addition, k^* is increasing in α. This is because a higher α implies less diminishing returns to capital, thus a higher capital-labor ratio is necessary to reduce the average return to capital to a level consistent with steady state as given in eq. (2.34).

Transitional dynamics are also straightforward in this case. In particular,

$$\dot{k}(t) = sA[k(t)]^{\alpha} - (n+\delta)k(t)$$

with initial condition $k(0)$. To solve this equation, let $x(t) \equiv k(t)^{1-\alpha}$, so the equilibrium law of motion of the capital labor ratio can be written in terms of $x(t)$ as

$$\dot{x}(t) = (1-\alpha) sA - (1-\alpha)(n+\delta) x(t),$$

which is a linear differential equation, with a general solution

$$x(t) = \frac{sA}{n+\delta} + \left[x(0) - \frac{sA}{n+\delta}\right] \exp\left(-(1-\alpha)(n+\delta) t\right).$$

(see Appendix Chapter B). Expressing this solution in terms of the capital-labor ratio

$$k(t) = \left\{\frac{sA}{n+\delta} + \left[[k(0)]^{1-\alpha} - \frac{sA}{\delta}\right] \exp\left(-(1-\alpha)(n+\delta) t\right)\right\}^{\frac{1}{1-\alpha}}.$$

This solution illustrates that starting from any $k(0)$, the equilibrium $k(t) \to k^* = (sA/(n+\delta))^{1/(1-\alpha)}$, and in fact, the rate of adjustment is related to $(1-\alpha)(n+\delta)$. More specifically, the gap between $k(0)$ and the steady-state value k^* narrows at the exponential rate $(1-\alpha)(n+\delta)$. This is intuitive: a higher α implies less diminishing returns to capital, which slows down the rate at which the marginal and average product of capital decline as capital accumulates, and this reduces the rate of adjustment to steady state. Similarly, a smaller δ means less replacement of depreciated capital and a smaller n means slower population growth, and both of those slow down the adjustment of capital per worker and thus the rate of transitional dynamics.

Example 2.3. **(The Constant Elasticity of Substitution Production Function)** The previous example introduced the Cobb-Douglas production function, which featured an elasticity of substitution equal to 1. The Cobb-Douglas production function is a special case of the constant elasticity of substitution (CES) production function, first introduced by Arrow, Chenery, Minhas and Solow (1961). This production function imposes a constant elasticity, σ, not necessarily equal to 1. Consider a vector-valued index of technology $\mathbf{A}(t) = (A_H(t), A_K(t), A_L(t))$. Then, the CES production function can be written as

$$\begin{aligned} Y(t) &= F[K(t), L(t), \mathbf{A}(t)] \\ &\equiv A_H(t) \left[\gamma (A_K(t) K(t))^{\frac{\sigma-1}{\sigma}} + (1-\gamma)(A_L(t) L(t))^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \end{aligned} \tag{2.38}$$

where $A_H(t) > 0$, $A_K(t) > 0$ and $A_L(t) > 0$ are three different types of technological change which will be discussed further in Section 2.7; $\gamma \in (0,1)$ is a distribution parameter, which determines how important labor and capital services are in determining the production of the final good and $\sigma \in [0, \infty]$ is the elasticity of substitution. To verify that σ is indeed the constant elasticity of substitution, let us use (2.37). In particular, it is easy to verify that the ratio of the marginal product of capital to the marginal productive labor, F_K/F_L, is given by

$$\frac{F_K}{F_L} = \frac{\gamma A_K(t)^{\frac{\sigma-1}{\sigma}} K(t)^{-\frac{1}{\sigma}}}{(1-\gamma) A_L(t)^{\frac{\sigma-1}{\sigma}} L(t)^{-\frac{1}{\sigma}}},$$

so that the elasticity of substitution is indeed given by σ, that is,

$$\sigma = -\left[\frac{\partial \ln (F_K/F_L)}{\partial \ln (K/L)}\right]^{-1}.$$

The CES production function is particularly useful because it is more general and flexible than the Cobb-Douglas form while still being tractable. As $\sigma \to 1$, the CES production function (2.38) converges to the Cobb-Douglas function $Y(t) = A_H(t)(A_K(t))^{\gamma}(A_L(t))^{1-\gamma}(K(t))^{\gamma}(L(t))^{1-\gamma}$. As $\sigma \to \infty$, the CES production function becomes linear, i.e.

$$Y(t) = \gamma A_H(t) A_K(t) K(t) + (1-\gamma) A_H(t) A_L(t) L(t).$$

Finally, as $\sigma \to 0$, the CES production function converges to the Leontief production function with no substitution between factors,

$$Y(t) = A_H(t) \min\{\gamma A_K(t) K(t); (1-\gamma) A_L(t) L(t)\}.$$

The special feature of the Leontief production function is that if $\gamma A_K(t) K(t) \neq (1-\gamma) A_L(t) L(t)$, either capital or labor will be partially "idle" in the sense that a small reduction in capital or labor will have no effect on output or factor prices. Exercise 2.18 illustrates a number of the properties of the CES production function, while Exercise 2.19 provides an alternative derivation of this production function along the lines of the original article by Arrow, Chenery, Minhas and Solow (1961). Finally, an important observation is that the CES production function with $\sigma > 1$ violates Assumption 1 (see Exercise 2.20). Therefore, in the context of aggregate production functions with capital and labor, $\sigma \leq 1$ is the natural assumption.

2.6. A First Look at Sustained Growth

Can the Solow model generate sustained growth *without* technological progress? The answer is yes, but only if some of the assumptions imposed so far are relaxed.

The Cobb-Douglas example above already showed that when α is close to 1, the adjustment of the capital-labor ratio back to its steady-state level can be very slow. A very slow adjustment towards a steady-state has the flavor of "sustained growth" rather than the economy (quickly) settling down to a steady state. In fact, the simplest model of sustained growth essentially takes $\alpha = 1$ in terms of the Cobb-Douglas production function above. To do this, let us relax Assumptions 1 and 2 (which do not allow $\alpha = 1$), and suppose that

$$F[K(t), L(t), A(t)] = AK(t), \tag{2.39}$$

where $A > 0$ is a constant. This is the so-called "AK" model, and in its simplest form output does not even depend on labor. The results here apply with more general constant returns to scale production functions, for example, with

$$F[K(t), L(t), A(t)] = AK(t) + BL(t). \tag{2.40}$$

Nevertheless, it is simpler to illustrate the main insights with (2.39), leaving the analysis of the case when the production function is given by (2.40) to Exercise 2.17.

Let us continue to assume that population grows at a constant rate n as before (cfr. eq. (2.32)). Then, combining this with the production function (2.39), the fundamental law of motion of the capital stock becomes

$$\frac{\dot{k}(t)}{k(t)} = sA - \delta - n.$$

This equation shows that when the parameters satisfy the inequality $sA - \delta - n > 0$, there will be sustained growth in the capital-labor ratio and thus in output per capita. This result is summarized in the next proposition.

Proposition 2.10. *Consider the Solow growth model with the production function (2.39) and suppose that $sA - \delta - n > 0$. Then, in equilibrium, there is sustained growth of output per capita at the rate $sA - \delta - n$. In particular, starting with a capital-labor ratio $k(0) > 0$, the economy has*

$$k(t) = \exp\left((sA - \delta - n)t\right) k(0)$$

and

$$y(t) = \exp\left((sA - \delta - n)t\right) Ak(0).$$

This proposition not only establishes the possibility of endogenous sustained growth, but also shows that in this simplest form, there are no transitional dynamics. The economy always grows at a constant rate $sA - \delta - n$, regardless of the level of capital-labor ratio it starts from. Figure 2.10 shows this equilibrium diagrammatically.

Does the AK model provide an appealing approach to explaining sustained growth? While its simplicity is a plus, the model has a number of unattractive features. First, it is a somewhat knife-edge case, which does not satisfy Assumptions 1 and 2; in particular, it requires the production function to be ultimately linear in the capital stock. Second and relatedly, this feature implies that as time goes by the share of national income accruing to capital will increase towards 1 (if it is not equal to 1 to start with). The next section shows that this does not seem to be borne out by the data. Finally and most importantly, a variety of evidence suggests that technological progress is a major (perhaps the most major) factor in understanding the process of economic growth. A model of sustained growth without technological progress fails to capture this essential aspect of economic growth. Motivated by these considerations, we next turn to the task of introducing technological progress into the baseline Solow growth model.

2.7. Solow Model with Technological Progress

2.7.1. Balanced Growth. The models analyzed so far did not feature technological progress. I now introduce changes in $A(t)$ to capture improvements in the technological know-how of the economy. There is little doubt that today human societies know how to produce many more goods and can do so more efficiently than in the past. In other words, the productive knowledge of the human society has progressed tremendously over the past 200 years, and even more tremendously over the past 1,000 or 10,000 years. This suggests

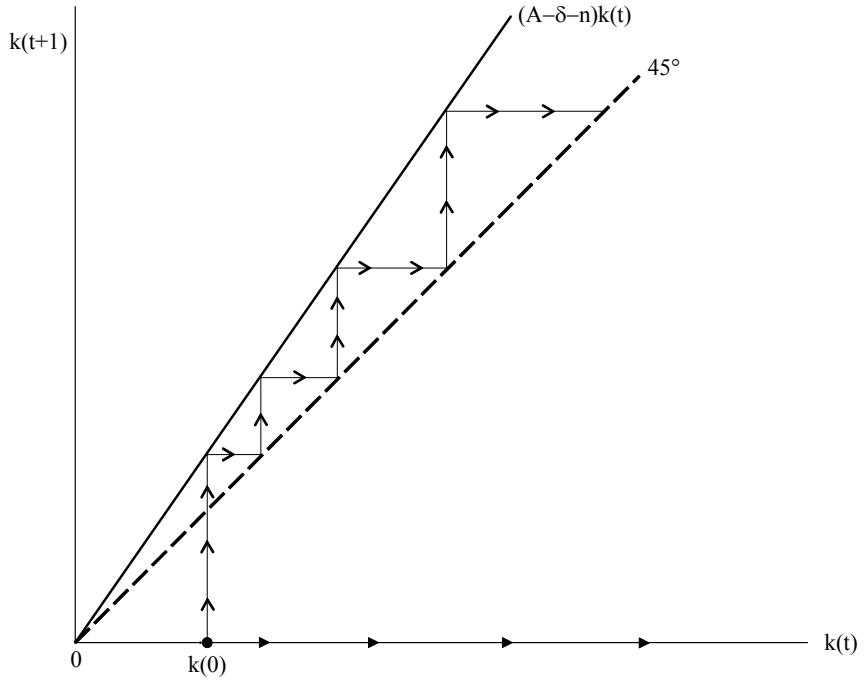

Figure 2.10. Sustained growth with the linear AK technology with $sA-\delta-n>0$.

that an attractive way of introducing economic growth in the framework developed so far is to allow technological progress in the form of changes in $A(t)$. The question is how to do this. It turns out that the production function $F[K(t), L(t), A(t)]$ is too general to achieve our objective and may not lead to *balanced growth*. So our first task is to understand what types of restrictions need to be placed on F.

By balanced growth, we mean a path of the economy consistent with the *Kaldor facts* (Kaldor, 1963), that is, a path where, while output per capita increases, the capital-output ratio, the interest rate, and the distribution of income between capital and labor remain roughly constant. Figure 2.11, for example, shows the evolution of the shares of capital and labor in the US national income.

Despite fairly large fluctuations, there is no trend in these factors shares. Moreover, a range of evidence suggests that there is no apparent trend in interest rates over long time horizons and even in different societies (see, for example, Homer and Sylla, 1991). These facts and the relative constancy of capital-output ratios until the 1970s have made many economists prefer models with balanced growth to those without. The share of capital in national income and the capital-output ratio are not exactly constant. For example, since the 1970s both the share of capital in national income and the capital-output ratio may have increased depending on how one measures them. Nevertheless, constant factor shares and a

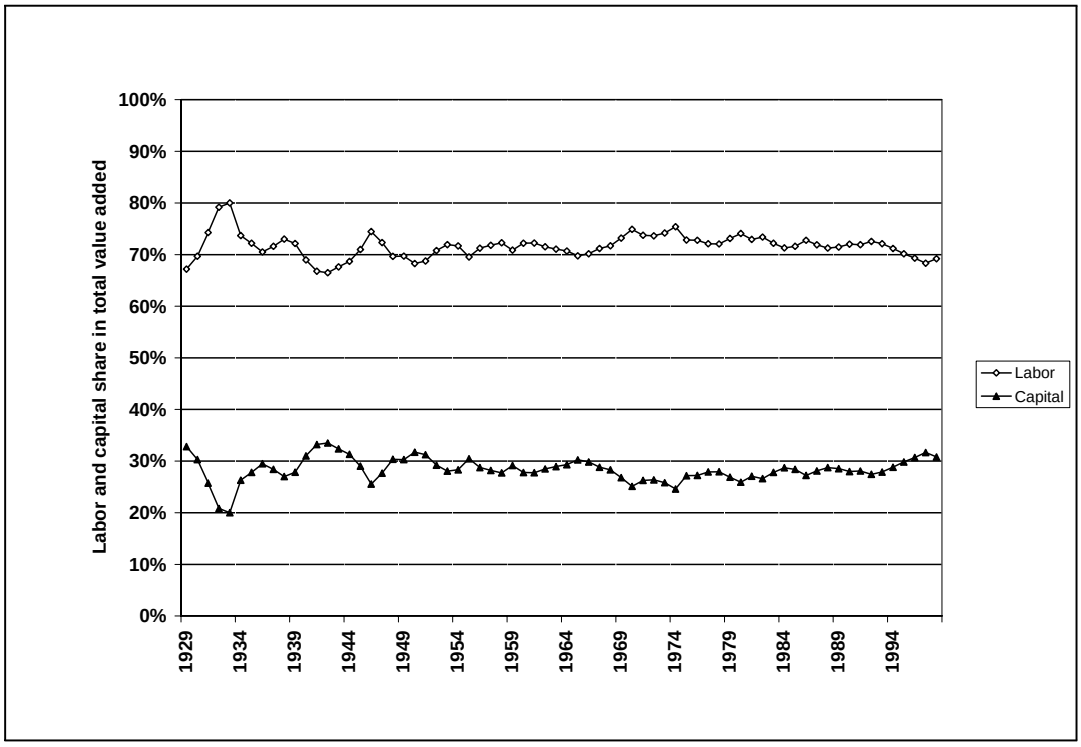

Figure 2.11. Capital and Labor Share in the U.S. GDP.

constant capital-output ratio are a good approximation to reality and a very useful starting point for our models.

Also for future reference, note that the capital share in national income is about 1/3, while the labor share is about 2/3. This ignores the share of land: land is not a major factor of production in modern economies (though this is not true historically or for the less-developed economies of today). Exercise 2.12 illustrates how incorporating land into this framework changes the analysis. Note also that this pattern of the factor distribution of income, combined with economists' desire to work with simple models, often makes them choose a Cobb-Douglas aggregate production function of the form $AK^{1/3}L^{2/3}$ as an approximation to reality (especially since it ensures that factor shares are constant by construction). However, as Theorem 2.6 shows Cobb-Douglas technology is not necessary for constant factor shares or for balanced growth.

Why balanced growth? The most important reason to start with balanced growth is that it is much easier to handle than non-balanced growth, since the equations describing the law of motion of the economy can be represented by difference or differential equations with well-defined steady states. Put more succinctly, the main advantage of balanced growth is that it corresponds to a steady state of a transformed system (thus, balanced growth will imply $\dot{k} = 0$, except that now the definition of k will change). This will enable us to use the same tools used to analyze the steady states of stationary models to study economies with sustained

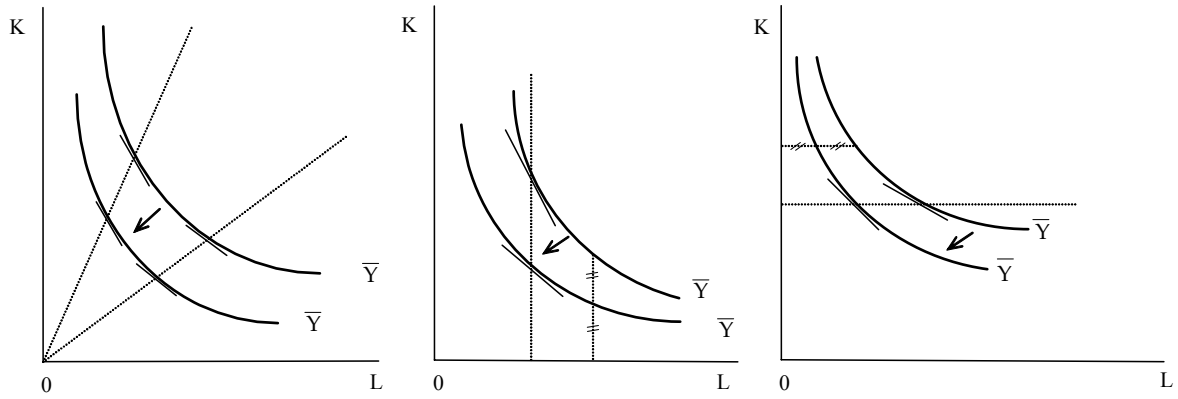

FIGURE 2.12. Hicks-neutral, Solow-neutral and Harrod-neutral shifts in isoquants.

growth. It is nonetheless important to bear in mind that in reality, growth has many non-balanced features. For example, the share of different sectors changes systematically over the growth process, with agriculture shrinking, manufacturing first increasing and then shrinking. Ultimately, we would like to have models that combine certain "balanced features" with these types of structural transformations. I will return to these issues in Part 7 of the book.

2.7.2. Types of Neutral Technological Progress. What are some convenient special forms of the general production function $F[K(t), L(t), A(t)]$? A first possibility is:

$$F[K(t), L(t), A(t)] = A(t)\tilde{F}[K(t), L(t)],$$

for some constant returns to scale function $\tilde{F}$. This functional form implies that the technology term $A(t)$ is simply a multiplicative constant in front of another (quasi-) production function $\tilde{F}$ and is referred to as *Hicks-neutral* after the famous British economist John Hicks. Intuitively, consider the isoquants of the function $F[K(t), L(t), A(t)]$ in the L-K space, which plot combinations of labor and capital for a given technology $A(t)$ such that the level of production is constant. This is shown in Figure 2.12. Hicks-neutral technological progress, in the first panel, corresponds to a relabeling of the isoquants (without any change in their shape).

Another alternative is to have capital-augmenting or *Solow-neutral* technological progress, in the form

$$F[K(t), L(t), A(t)] = \tilde{F}[A(t)K(t), L(t)].$$

This is also referred to as capital-augmenting progress, because a higher $A(t)$ is equivalent to the economy having more capital. This type of technological progress corresponds to the isoquants shifting inwards as if the capital axis is being "shrunk" (since a higher A now corresponds to a greater level of "effective capital"). This is shown in the second panel of Figure 2.12 for a doubling of $A(t)$.

Finally, we can have labor-augmenting or *Harrod-neutral* technological progress, named after an early influential growth theorist Roy Harrod, who we encountered above in the context of the Harrod-Domar model previously:

$$F\left[K\left(t\right),L\left(t\right),A\left(t\right)\right]=\tilde{F}\left[K\left(t\right),A\left(t\right)L\left(t\right)\right].$$

This functional form implies that an increase in technology $A\left(t\right)$ increases output as if the economy had more labor and thus corresponds to an inward shift of the isoquant as if the labor axis is being "shrunk". The approximate form of the shifts in the isoquants are plotted in the third panel of Figure 2.12, again for a doubling of $A\left(t\right)$.

Of course, in practice technological change can be a mixture of these, so we could have a vector valued index of technology $\mathbf{A}\left(t\right)=\left(A_{H}\left(t\right),A_{K}\left(t\right),A_{L}\left(t\right)\right)$ and a production function that looks like

$$F\left[K\left(t\right),L\left(t\right),\mathbf{A}\left(t\right)\right]=A_{H}\left(t\right)\tilde{F}\left[A_{K}\left(t\right)K\left(t\right),A_{L}\left(t\right)L\left(t\right)\right], \tag{2.41}$$

which nests the constant elasticity of substitution production function introduced in Example 2.3 above. Nevertheless, even (2.41) is a restriction on the form of technological progress, since changes in technology, $A\left(t\right)$, could modify the entire production function.

It turns out that, although all of these forms of technological progress look equally plausible ex ante, our desire to focus on balanced growth forces us to one of these types of neutral technological progress. In particular, balanced growth necessitates that all technological progress be labor-augmenting or Harrod-neutral. This is a very surprising result and it is also somewhat troubling, since there are no compelling reasons for why technological progress should take this form. I now state and prove the relevant theorem here and return to the discussion of why long-run technological change might be Harrod-neutral in Chapter 15.

2.7.3. The Steady-State Technological Progress Theorem. A version of the following theorem was first proved by the early growth economist Hirofumi Uzawa (1961). For simplicity and without loss of any generality, let us focus on continuous-time models. The key elements of balanced growth, as suggested by the discussion above, are the constancy of factor shares and the constancy of the capital-output ratio, $K\left(t\right)/Y\left(t\right)$. The shares of capital and labor in national income are

$$\alpha_{K}\left(t\right)\equiv\frac{R\left(t\right)K\left(t\right)}{Y\left(t\right)}\text{ and }\alpha_{L}\left(t\right)\equiv\frac{w\left(t\right)L\left(t\right)}{Y\left(t\right)}.$$

By Assumption 1 and Theorem 2.1, $\alpha_{K}\left(t\right)+\alpha_{L}\left(t\right)=1$.

The following theorem is a stronger version of a result first stated and proved by Uzawa and the proof builds on the argument developed in the recent paper by Schlicht (2006). It shows that asymptotic growth at constant rates (for output, capital and consumption) requires the aggregate production function to have a representation with Harrod-neutral technological progress. Let us define an asymptotic path as a path of output, capital, consumption and labor as $t\to\infty$.

Theorem 2.6. ***(Uzawa's Theorem)*** *Consider a growth model with a constant returns to scale aggregate production function*

$$Y(t) = F\left[K(t), L(t), \tilde{A}(t)\right],$$

with $\tilde{A}(t)$ *representing technology at time* t *and aggregate resource constraint*

$$\dot{K}(t) = Y(t) - C(t) - \delta K(t).$$

Suppose that there is a constant growth rate of population, $L(t) = \exp(nt) L(0)$*, and that there exists an asymptotic path where output, capital and consumption grow at constant rates, that is,* $\dot{Y}(t)/Y(t) = g_Y$*,* $\dot{K}(t)/K(t) = g_K$ *and* $\dot{C}(t)/C(t) = g_C$*. Suppose finally that* $g_K + \delta > 0$*. Then, along the growth path*

(1) $g_Y = g_K = g_C$*; and*

(2) *asymptotically, the aggregate production function can be represented as:*

$$Y(t) = \tilde{F}[K(t), A(t) L(t)],$$

where

$$\frac{\dot{A}(t)}{A(t)} = g = g_Y - n.$$

Proof. **(Part 1)** By hypothesis, as $t \to \infty$, we have $Y(t) = \exp(g_Y(t-\tau)) Y(\tau)$, $K(t) = \exp(g_K(t-\tau)) K(\tau)$ and $L(t) = \exp(n(t-\tau)) L(\tau)$ for some $\tau < \infty$. The aggregate resource constraint at time t implies

$$(g_K + \delta) K(t) = Y(t) - C(t).$$

Since the left-hand side is positive by hypothesis, we can divide both sides by $\exp(g_K(t-\tau))$ and write date t quantities in terms of date τ quantities to obtain

$$(g_K + \delta) K(\tau) = \exp((g_Y - g_K)(t-\tau)) Y(\tau) - \exp((g_C - g_K)(t-\tau)) C(\tau)$$

for all t. Differentiating with respect to time yields

$$(g_Y - g_K) \exp((g_Y - g_K)(t-\tau)) Y(\tau) - (g_C - g_K) \exp((g_C - g_K)(t-\tau)) C(\tau) = 0$$

for all t. This equation can hold for all t either if $g_Y = g_K = g_C$ or if $g_Y = g_C$ and $Y(\tau) = C(\tau)$. However the latter condition contradicts $g_K + \delta > 0$. Therefore, $g_Y = g_K = g_C$ as claimed in the first part of the theorem.

(Part 2) The aggregate production function for time τ can be written as

$$\exp(-g_Y(t-\tau)) Y(t) = F\left[\exp(-g_K(t-\tau)) K(t), \exp(-n(t-\tau)) L(t), \tilde{A}(\tau)\right].$$

Multiplying both sides by $\exp(g_Y(t-\tau))$ and using the constant returns to scale property of F,

$$Y(t) = F\left[\exp((t-\tau)(g_Y - g_K)) K(t), \exp((t-\tau)(g_Y - n)) L(t), \tilde{A}(\tau)\right].$$

From part 1, $g_Y = g_K$ and thus

$$Y(t) = F\left[K(t), \exp((t-\tau)(g_Y - n)) L(t), \tilde{A}(\tau)\right].$$

Moreover, this equation is true for all t regardless of the initial τ, thus

$$\begin{aligned} Y(t) &= \tilde{F}\left[K(t), \exp\left(\left(t-\tau\right)\left(g_Y - n\right)\right) L(t)\right], \\ &= \tilde{F}\left[K(t), A(t) L(t)\right], \end{aligned}$$

with

$$\frac{\dot{A}(t)}{A(t)} = g_Y - n$$

establishing the second part of the theorem. □

A remarkable feature of this theorem is that it was stated and proved without any reference to equilibrium behavior or market clearing. Also, contrary to Uzawa's original theorem, it is not stated for a balanced growth path (meaning an equilibrium path with constant factor shares), but only for an asymptotic path with constant rates of output, capital and consumption growth. The proposition only exploits the definition of asymptotic paths, the constant returns to scale nature of the aggregate production function and the resource constraint. Consequently, the result is very powerful and fairly general.

Before providing a more economic intuition for this result, let us state an immediate implication of this theorem as a corollary, which will be useful both in the discussions below and for the intuition:

Corollary 2.3. *Under the assumptions of Theorem 2.6, if an economy has an asymptotic path with constant growth of output, capital and consumption, then asymptotically technological progress can be represented as Harrod neutral (purely labor-augmenting).*

The intuition for Theorem 2.6 and for the corollary is simple. We have assumed that the economy features capital accumulation in the sense that $g_K + \delta > 0$. From the aggregate resource constraint, this is only possible if output and capital grow at the same rate. Either this growth rate is equal to the rate of population growth, n, in which case there is no technological change (the proposition applies with $g = 0$) or the economy exhibits growth of per capita income and capital-labor ratio. The latter case creates an asymmetry between capital and labor, in the sense that capital is accumulating faster than labor. Constancy of growth then requires technological change to make up for this asymmetry—that is, technology to take a labor-augmenting form.

This intuition does not provide a reason for why technology should take this labor-augmenting (Harrod-neutral) form, however. The theorem and its corollary simply state that if technology did not take this form, an asymptotic path with constant growth rates would not be possible. At some level, this is a distressing result, since it implies that balanced growth (in fact something weaker than balanced growth) is only possible under a very stringent assumption. Chapter 15 will show that when technology is endogenous the same intuition also implies that technology should be endogenously more labor-augmenting than capital-augmenting.

Notice also that this theorem does not state that technological change has to be labor-augmenting all the time. Instead, it requires that technological change ought to be labor-augmenting asymptotically (along the balanced growth path). This is exactly the pattern that certain classes of endogenous-technology models will generate.

Finally, it is important to emphasize that Theorem 2.6 does not require that $Y^*(t) = \tilde{F}[K^*(t), A(t) L(t)]$, but only that it has a representation of the form $Y^*(t) = \tilde{F}[K^*(t), A(t) L(t)]$. This allows one important exception to the statement that "asymptotically technological change has to be Harrod neutral". If the aggregate production function is Cobb-Douglas and takes the form

$$Y(t) = [A_K(t) K(t)]^{\alpha} [A_L(t) L(t)]^{1-\alpha},$$

then both $A_K(t)$ and $A_L(t)$ could grow asymptotically, while maintaining balanced growth. However, in this Cobb-Douglas example we can define $A(t) = [A_K(t)]^{\alpha/(1-\alpha)} A_L(t)$ and the production function can be represented as

$$Y(t) = [K(t)]^{\alpha} [A(t) L(t)]^{1-\alpha}.$$

In other words, technological change can be represented as purely labor-augmenting, which is what Theorem 2.6 requires. Intuitively, the differences between labor-augmenting and capital-augmenting (and other forms) of technological progress matter when the elasticity of substitution between capital and labor is not equal to 1. In the Cobb-Douglas case, as we have seen above, this elasticity of substitution is equal to 1, thus different forms of technological progress are simple transforms of each other.

Another important corollary of Theorem 2.6 is obtained when we also assume that factor markets are competitive.

COROLLARY 2.4. *Under the conditions of Theorem 2.6, if factor markets are competitive, then asymptotic factor shares are constant. That is, as* $t \to \infty$, $\alpha_L(t) \to \alpha_L^*$ *and* $\alpha_K(t) \to \alpha_K^*$.

PROOF. With competitive factor markets, we have that as $t \to \infty$

$$\begin{aligned} \alpha_K(t) &\equiv \frac{R(t) K(t)}{Y(t)} \\ &= \frac{K(t)}{Y(t)} \frac{\partial \tilde{F}[K(t), A(t) L(t)]}{\partial K(t)} \\ &= \alpha_K^*, \end{aligned}$$

where the second line uses the definition of the rental rate of capital in a competitive market, and the third line uses the fact that as $t \to \infty$, $g_Y = g_K$ and $g_K = g + n$ from Theorem 2.6 and that because $\tilde{F}$ exhibits constant returns to scale, its derivative with respect to K must be homogeneous of degree 0. □

This corollary, together with Theorem 2.6, implies that any asymptotic path with constant growth rates for output, capital and consumption must be a balanced growth path and can

only be generated from an aggregate production function asymptotically featuring Harrod-neutral technological change.

In light of this corollary, we can provide further intuition for Theorem 2.6. Suppose the production function takes the special form $F\left[A_K\left(t\right)K\left(t\right),A_L\left(t\right)L\left(t\right)\right]$. The corollary implies that factor shares will be constant. Given constant returns to scale, this can only be the case when total capital inputs, $A_K\left(t\right)K\left(t\right)$, and total labor inputs, $A_L\left(t\right)L\left(t\right)$, grow at the same rate; otherwise, the share of either capital or labor will be increasing over time. But if total labor and capital inputs grow at the same rate, then output $Y\left(t\right)$ must also grow at this rate (because of constant returns to scale). The fact that the capital-output ratio is constant in steady state finally implies that $K\left(t\right)$ must grow at the same rate as output and thus at the same rate as $A_L\left(t\right)L\left(t\right)$. Therefore, balanced growth is only possible if $A_K\left(t\right)$ is asymptotically constant.

2.7.4. The Solow Growth Model with Technological Progress: Continuous Time. I now present an analysis of the Solow growth model with technological progress in continuous time. The discrete-time case can be analyzed analogously and I omit the details to avoid repetition. Theorem 2.6 implies that asymptotically the production function must have a representation of the form

$$Y\left(t\right)=F\left[K\left(t\right),A\left(t\right)L\left(t\right)\right],$$

with purely labor-augmenting technological progress. For simplicity, I assume that it takes this form throughout. Moreover, suppose that there is technological progress at the rate $g>0$, that is,

$$\frac{\dot{A}\left(t\right)}{A\left(t\right)}=g>0, \tag{2.42}$$

and population grows at the rate n as assumed in (2.32) above. Again using the constant saving rate we have

$$\dot{K}\left(t\right)=sF\left[K\left(t\right),A\left(t\right)L\left(t\right)\right]-\delta K\left(t\right). \tag{2.43}$$

The simplest way of analyzing this economy is to express everything in terms of a normalized variable. Since "effective" or efficiency units of labor are given by $A\left(t\right)L\left(t\right)$, and F exhibits constant returns to scale in its two arguments, I now define $k\left(t\right)$ as the *effective capital-labor* ratio (capital divided by efficiency units of labor) so that

$$k\left(t\right)\equiv\frac{K\left(t\right)}{A\left(t\right)L\left(t\right)}. \tag{2.44}$$

Differentiating this expression with respect to time,

$$\frac{\dot{k}\left(t\right)}{k\left(t\right)}=\frac{\dot{K}\left(t\right)}{K\left(t\right)}-g-n. \tag{2.45}$$

The quantity of output per unit of effective labor can be written as

$$\begin{aligned}\hat{y}(t) &\equiv \frac{Y(t)}{A(t)L(t)} \\ &= F\left[\frac{K(t)}{A(t)L(t)}, 1\right] \\ &\equiv f(k(t)).\end{aligned}$$

Income per capita is $y(t) \equiv Y(t)/L(t)$, so that

$$\begin{aligned}y(t) &= A(t)\hat{y}(t) \\ &= A(t)f(k(t)).\end{aligned} \tag{2.46}$$

It should be clear that if $\hat{y}(t)$ is constant, income per capita, $y(t)$, will grow over time, since $A(t)$ is growing. This highlights that in this model, and more generally in models with technological progress, we should not look for "steady states" where income per capita is constant, but for *balanced growth paths*, where income per capita grows at a constant rate, while some transformed variables such as $\hat{y}(t)$ or $k(t)$ in (2.45) remain constant. Since these transformed variables remain constant, balanced growth paths can be thought of as steady states of a transformed model. Motivated by this observation, in models with technological change throughout I will use the terms "steady state" and balanced growth path interchangeably.

Next, substituting for $\dot{K}(t)$ from (2.43) into (2.45),

$$\frac{\dot{k}(t)}{k(t)} = \frac{sF[K(t), A(t)L(t)]}{K(t)} - (\delta + g + n).$$

Using (2.44),

$$\frac{\dot{k}(t)}{k(t)} = \frac{sf(k(t))}{k(t)} - (\delta + g + n), \tag{2.47}$$

which is very similar to the law of motion of the capital-labor ratio in the model without technological progress, (2.33). The only difference is the presence of g, which reflects the fact that now k is no longer the capital-labor ratio but the *effective* capital-labor ratio and thus for k to remain constant in the balanced growth path, the capital-labor ratio needs to increase at the rate g.

An equilibrium in this model is defined similarly to before. A steady state or a balanced growth path is, in turn, defined as an equilibrium in which the effect of capital-labor ratio $k(t)$ is constant. Consequently (proof omitted):

PROPOSITION 2.11. *Consider the basic Solow growth model in continuous time with Harrod-neutral technological progress at the rate g and population growth at the rate n. Suppose that Assumptions 1 and 2 hold, and define the effective capital-labor ratio as in (2.44). Then, there exists a unique steady state (balanced growth path) equilibrium where the effective capital-labor ratio is equal to $k^* \in (0, \infty)$ given by*

$$\frac{f(k^*)}{k^*} = \frac{\delta + g + n}{s}. \tag{2.48}$$

Per capita output and consumption grow at the rate g.

Equation (2.48), which determines the steady-state level of the effective capital-labor ratio, emphasizes that now total savings, $sf(k)$, are used for replenishing the capital stock for three distinct reasons. The first is again the depreciation at the rate δ. The second is population growth at the rate n, which reduces capital per worker. The third is Harrod-neutral technological progress at the rate g, reducing effective capital-labor ratio at this rate when the capital-labor ratio is constant. Thus the replenishment of the effective capital-labor ratio requires investments to be equal to $(\delta + g + n) k$, which is the intuitive explanation for eq. (2.48).

The comparative static results are also similar to before, with the additional comparative static with respect to the initial level of the labor-augmenting technology, $A(0)$ (the level of technology at all points in time, $A(t)$, is completely determined by $A(0)$ given the assumption in (2.42)). We therefore have:

Proposition 2.12. *Suppose Assumptions 1 and 2 hold and let $A(0)$ be the initial level of technology. Denote the balanced growth path level of effective capital-labor ratio by $k^*(A(0), s, \delta, n)$ and the level of output per capita by $y^*(A(0), s, \delta, n, t)$ (the latter is a function of time since it is growing over time). Then,*

$$\begin{aligned} \frac{\partial k^*(A(0), s, \delta, n)}{\partial A(0)} &= 0, \frac{\partial k^*(A(0), s, \delta, n)}{\partial s} > 0, \\ \frac{\partial k^*(A(0), s, \delta, n)}{\partial n} &< 0 \text{ and } \frac{\partial k^*(A(0), s, \delta, n)}{\partial \delta} < 0, \end{aligned}$$

and also,

$$\begin{aligned} \frac{\partial y^*(A(0), s, \delta, n, t)}{\partial A(0)} &> 0, \frac{\partial y^*(A(0), s, \delta, n, t)}{\partial s} > 0, \\ \frac{\partial y^*(A(0), s, \delta, n, t)}{\partial n} &< 0 \text{ and } \frac{\partial y^*(A(0), s, \delta, n, t)}{\partial \delta} < 0, \end{aligned}$$

for each t.

Proof. See Exercise 2.21. □

Finally, the transitional dynamics of the economy with technological progress are similar to the dynamics without technological change.

Proposition 2.13. *Suppose that Assumptions 1 and 2 hold, then the Solow growth model with Harrod-neutral technological progress and population growth in continuous time is asymptotically stable, that is, starting from any $k(0) > 0$, the effective capital-labor ratio converges to a steady-state value k^* ($k(t) \to k^*$).*

Proof. See Exercise 2.22. □

Therefore, the comparative statics and dynamics are very similar to the model without technological progress. The major difference is that now the model generates growth in output per capita, so it can be mapped to the data much better. However, the disadvantage is that growth is driven entirely *exogenously*. The growth rate of the economy is exactly

the same as the exogenous growth rate of the technology stock. The model specifies neither where this technology stock comes from nor how fast it grows.

2.8. Comparative Dynamics

This section briefly undertakes some simple "comparative dynamics" exercises. Comparative dynamics are different from comparative statics in Propositions 2.3, 2.8 or 2.12 in that the focus is now on the entire path of adjustment of the economy following a shock or a change in parameters (rather than steady-state comparisons). The basic Solow model is particularly well suited to such an analysis because of its simplicity. These exercises are also useful because the basic Solow model, and its neoclassical cousin, are often used for analysis of policy changes, medium-run shocks and business cycle dynamics, so an understanding of how the basic model responds to various shocks is useful in a range of applications.

Recall that the law of motion of the effective capital-labor ratio in the continuous-time Solow model is given by (2.47) $\dot{k}(t)/k(t) = sf(k(t))/k(t) - (\delta + g + n)$. The right-hand side of this equation is plotted in Figure 2.13. The intersection with the horizontal axis gives the unique steady-state (balanced growth) equilibrium, with effective capital-labor ratio k^*. This figure is sufficient for the analysis of comparative dynamics. Consider, for example, a one-time, unanticipated, permanent increase in the saving rate from s to s'. This shifts the curve to the right as shown by the dashed line, with a new intersection with the horizontal axis at k^{**}. The dashed arrows under the horizontal axis show how the effective capital-labor ratio adjusts gradually to the new balanced growth effective capital-labor ratio, k^{**}. Immediately after the increase in the saving rate is realized, the capital stock and the effective capital-labor ratio remain unchanged (since they are *state* variables). After this point, k follows the dashed arrows and converges monotonically to k^{**}.

The comparative dynamics following a one-time, unanticipated, permanent decrease in δ or n are identical.

The same diagrammatic analysis can be used for studying the effect of an unanticipated, but transitory change in parameters. For example, imagine that s changes in an unanticipated manner at $t = t'$, but this change will be reversed and the saving rate will return back to its original value at some known future date $t'' > t'$. In this case, starting at t', the economy follows the dashed arrows until t'. After t'', the original steady state of the differential equation applies and together with this the solid arrows above the horizontal axis become effective. Thus from t'' onwards, the economy gradually returns back to its original balanced growth equilibrium, k^*.

We will see that similar comparative dynamics can be carried out in the neoclassical growth model as well, but the response of the economy to some of these changes will be more complex.

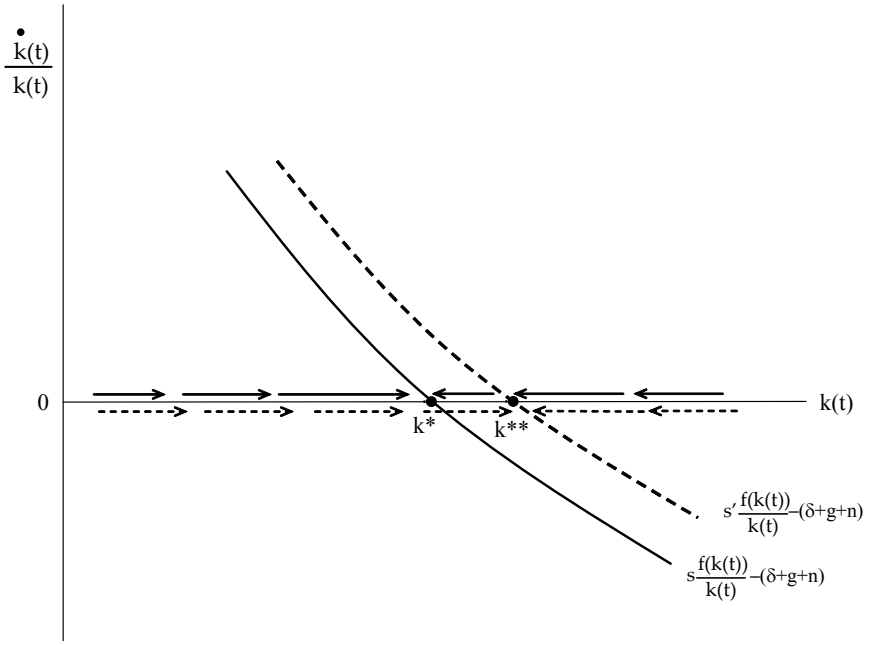

FIGURE 2.13. Dynamics following an increase in the savings rate from s to s'. The solid arrows show the dynamics for the initial steady state, while the dashed arrows show the dynamics for the new steady state.

2.9. Taking Stock

What have we learned from the Solow model? At some level, a lot. We now have a simple and tractable framework, which allows us to discuss capital accumulation and the implications of technological progress. As we will see in the next chapter, this framework is already quite useful in helping us think about the data.

However, at some other level, we have learned relatively little. The questions that Chapter 1 posed are related to why some countries are rich while others are poor, to why some countries grow while others stagnate, and to why the world economy embarked upon the process of steady growth over the past few centuries. The Solow model shows us that if there is no technological progress, and as long as we are not in the AK world, ruled out by Assumption 2, there will be no sustained growth. In this case, we can talk about cross-country output differences, but *not* about growth of countries or growth of the world economy.

The Solow model generates per capita output growth only by introducing exogenous technological progress. But in this case, everything is driven by technological progress, and technological progress itself is exogenous, just a blackbox, outside the model and outside the influence of economic incentives. If technological progress is "where it's at", then we have to study and understand which factors generate technological progress, what makes some firms

and societies invent better technologies, and what induces firms and societies to adopt and use these superior technologies.

Even on the question of capital accumulation, the Solow model is not entirely satisfactory. The rate of capital accumulation is determined by the saving rate, the depreciation rate and the rate of population growth. All of these are taken as exogenous.

In this light, the Solow growth model is most useful for us as a framework laying out the general issues and questions. It emphasizes that to understand growth, we have to understand physical capital accumulation (and human capital accumulation, which will be discussed in the next chapter) and perhaps most importantly, technological progress. All of these are blackboxes in the Solow growth model. Therefore, much of the rest of the book will be about digging deeper, trying to uncover what lies in these blackboxes. I start by introducing consumer optimization in Chapter 8, which enables a more systematic study of capital accumulation. Then, I will turn to models in which human capital accumulation and technological progress are endogenous. A framework in which the rate of accumulation of factors of production and technology are endogenous gives us a framework for posing and answering questions related to the fundamental causes of economic growth.

Nevertheless, even in its bare-bones form, the Solow model is useful in helping us think about the world and bringing useful perspectives, especially related to the proximate causes of economic growth. This is the topic of the next chapter.

2.10. References and Literature

The model analyzed in this chapter was first developed in Solow (1956) and Swan (1956). Solow (1970) gives a nice and accessible treatment, with historical references. Barro and Sala-i-Martin's (2004, Chapter 1) textbook presents a more up-to-date treatment of the basic Solow model at the graduate level, while Jones (1998, Chapter 2) presents an excellent undergraduate treatment.

The treatment in the chapter made frequent references to basic consumer and general equilibrium theory. These are prerequisites for an adequate understanding of the theory of economic growth. Some of the important results from dynamic general equilibrium theory are discussed in Chapter 5. Mas-Colell, Whinston and Green's (1995) graduate microeconomics textbook contains an excellent treatment of most of the necessary material, including producer theory and an accessible presentation of the basic notions of general equilibrium theory, including a discussion of Arrow securities and the definition of Arrow-Debreu commodities.

Properties of homogeneous functions and Euler's Theorem can be found, for example, in Simon and Blume (1994, Chapter 20). The reader should be familiar with the Implicit Function Theorem and properties of concave and convex functions, which will be used throughout the book. A review is given in the Appendix Chapter A. The reader may also want to consult Simon and Blume (1994) and Rudin (1976).

Appendix Chapter B provides an overview of solutions to differential equations and difference equations and a discussion of stability. Theorems 2.2, 2.3, 2.4 and 2.5 follow from

the results presented there. In addition, the reader may want to consult Simon and Blume (1994), Luenberger (1979) or Boyce and DiPrima (1977) for various results on difference and differential equations. Knowledge of solutions to simple differential equations and stability properties of difference and differential equations at the level of Appendix Chapter B is assumed. In addition, the material in Luenberger (1979) is particularly useful since it contains a unified treatment of difference and differential equations. Galor (2005) gives an introduction to difference equations and discrete-time dynamical systems for economists.

The golden rule saving rate was introduced by Edmund Phelps (1961). It is called the "golden rule" rate with reference to the biblical golden rule "do unto others as you would have them do unto you" applied in an intergenerational setting—i.e., presuming that those living and consuming at each different date form a different generation. While the golden rule saving rate is of historical interest and useful for discussions of dynamic efficiency it has no intrinsic optimality property, since it is not derived from well-defined preferences. Optimal saving policies will be discussed in greater detail in Chapter 8.

The balanced growth facts were first noted by Kaldor (1963). Figure 2.11 uses data from Piketty and Saez (2004). Homer and Sylla (1991) discuss the history of interest rates over many centuries and across different societies; they show that there is no notable upward or downward trend in interest rate. Nevertheless, not all aspects of the economic growth process are "balanced," and the non-balanced nature of growth will be discussed in detail in Part 7 of the book, which also contains references to changes in the sectoral composition of output in the course of the growth process.

The steady state theorem, Theorem 2.6, was first proved by Uzawa (1961). Many different proofs are available. The proof given here is adapted from Schlicht (2006), which is also discussed in Jones and Scrimgeour (2006). A similar proof also appears in Wan (1971). Barro and Sala-i-Martin's (2004, Chapter 1) also suggest a proof. Nevertheless, their argument is incomplete, since it assumes that technological change *must be* a combination of Harrod and Solow-neutral technological change, which is rather restrictive and not necessary for the proof. The theorem and the proof provided here are therefore more general and complete.

As noted in the text, the CES production function was first introduced by Arrow, Chenery, Minhas and Solow (1961).

Finally, the interested reader should look at the paper by Hakenes and Irmen (2006) for why Inada conditions can remove the steady state at $k = 0$ in continuous time even when $f(0) = 0$. Here it suffices to say that whether this steady state exists or not is a matter of the order in which limits are taken. In any case, as noted in the text, the steady state at $k = 0$ has no economic content and will be ignored throughout the book.

2.11. Exercises

Exercise 2.1. Show that competitive labor markets and Assumption 1 imply that the wage rate must be strictly positive and thus (2.4) implies (2.3).

Exercise 2.2. Prove that Assumption 1 implies that $F(A, K, L)$ is concave in K and L, but not strictly so.

Exercise 2.3. Show that when F exhibits constant returns to scale and factor markets are competitive, the maximization problem in (2.5) either has no solution (the firm can make infinite profits), a unique solution $K = L = 0$, or a continuum of solutions (that is, any (K, L) with $K/L = \kappa$ for some $\kappa > 0$ is a solution).

Exercise 2.4. Consider the Solow growth model in continuous time with the following per capita production function

$$f(k) = k^4 - 6k^3 + 11k^2 - 6k.$$

(1) Which parts of Assumptions 1 and 2 does the underlying production function $F(K, L)$ violate?
(2) Show that with this production function, there exist three steady-state equilibria.
(3) Prove that two of these steady-state equilibria are locally stable, while one of them is locally unstable. Can any of these steady-state equilibria be globally stable?

Exercise 2.5. Prove Proposition 2.7.

Exercise 2.6. Prove Proposition 2.8.

Exercise 2.7. Let us introduce government spending in the basic Solow model. Consider the basic model without technological change. In particular, suppose that (2.9) takes the form

$$Y(t) = C(t) + I(t) + G(t),$$

with $G(t)$ denoting government spending at time t. Imagine that government spending is given by $G(t) = \sigma Y(t)$.

(1) Discuss how the relationship between income and consumption should be changed. Is it reasonable to assume that $C(t) = sY(t)$?
(2) Suppose that government spending partly comes out of private consumption, so that $C(t) = (s - \lambda\sigma) Y(t)$, where $\lambda \in [0, 1]$. What is the effect of higher government spending (in the form of higher σ) on the equilibrium of the Solow model?
(3) Now suppose that part of government spending is invested in the capital stock of the economy. In particular, let a fraction ϕ of $G(t)$ be invested in the capital stock, so that total investment at time t is given by

$$I(t) = (1 - s - (1 - \lambda)\sigma + \phi\sigma) Y(t).$$

Show that if ϕ is sufficiently high, the steady-state level of capital-labor ratio will increase as a result of higher government spending (corresponding to higher σ). Is this reasonable? How would you alternatively introduce public investments in this model?

Exercise 2.8. Suppose that $F(A, K, L)$ is concave in K and L (though not necessarily strictly so) and satisfies Assumption 2. Prove Propositions 2.2 and 2.5. How do we need to modify Proposition 2.6?

Exercise 2.9. Prove Proposition 2.6.

Exercise 2.10. Prove Corollary 2.2.

Exercise 2.11. Recall the definition of the elasticity of substitution σ in (2.37). Suppose labor markets are competitive and the wage rate is equal to w. Prove that if the aggregate production function $F(K, L, A)$ exhibits constant returns to scale in K and L, then

$$\varepsilon_{y,w} \equiv \frac{\partial y/\partial w}{y/w} = \sigma,$$

where, as usual, $y \equiv F(K, L, A)/L$.

Exercise 2.12. Consider a modified version of the continuous-time Solow growth model where the aggregate production function is

$$F(K, L, Z) = L^{\beta} K^{\alpha} Z^{1-\alpha-\beta},$$

where Z is land, available in fixed inelastic supply. Assume that $\alpha+\beta < 1$, capital depreciates at the rate δ, and there is an exogenous saving rate of s.

(1) First suppose that there is no population growth. Find the steady-state capital-labor ratio in the steady-state output level. Prove that the steady state is unique and globally stable.
(2) Now suppose that there is population growth at the rate n, that is, $\dot{L}/L = n$. What happens to the capital-labor ratio and output level as $t \to \infty$? What happens to returns to land and the wage rate as $t \to \infty$?
(3) Would you expect the population growth rate n or the saving rate s to change over time in this economy? If so, how?

Exercise 2.13. Consider the continuous-time Solow model without technological progress and with constant rate of population growth equal to n. Suppose that the production function satisfies Assumptions 1 and 2. Assume that capital is owned by capitalists and labor is supplied by a different set of agents, the workers. Following a suggestion by Kaldor (1957), suppose that capitalists save a fraction s_K of their income, while workers consume all of their income.

(1) Define and characterize the steady-state equilibrium of this economy and study its stability.
(2) What is the relationship between the steady-state capital-labor ratio k^* and the golden rule capital stock k^*_{gold} defined above?

Exercise 2.14. Consider the Solow growth model with constant saving rate s and depreciation rate of capital equal to δ. Assume that population is constant and the aggregate production function is given by the constant returns to scale production function

$$F[A_K(t) K(t), A_L(t) L(t)]$$

where $\dot{A}_L(t)/A_L(t) = g_L > 0$ and $\dot{A}_K(t)/A_K(t) = g_K > 0$.

(1) Suppose that F is Cobb-Douglas. Determine the steady-state growth rate and the adjustment of the economy to the steady state.

(2) Suppose that F is not Cobb-Douglas (even asymptotically). Prove that there does not exist a steady state. Explain why this is.
(3) For the case in which F is not Cobb-Douglas, determine what happens to the capital-labor ratio and output per capita as $t \to \infty$.

Exercise 2.15. Consider the Solow model with non-competitive labor markets. In particular, suppose that there is no population growth and no technological progress and that output is given by $F(K, L)$. The saving rate is equal to s and the depreciation rate is given by δ.

(1) First suppose that there is a minimum wage $\bar{w}$, such that workers are not allowed to be paid less than $\bar{w}$. If labor demand at this wage falls short of L, employment is equal to the amount of labor demanded by firms, L^d (and the unemployed do not contribute to production and earn zero). Assume that $\bar{w} > f(k^*) - k^* f'(k^*)$, where k^* is the steady-state capital-labor ratio of the basic Solow model given by $f(k^*)/k^* = \delta/s$. Characterize the dynamic equilibrium path of this economy starting with some amount of physical capital $K(0) > 0$.
(2) Next consider a different form of labor market imperfection, whereby workers receive a fraction $\beta > 0$ of output in each firm as their wage income. Characterize a dynamic equilibrium path in this case. [Hint: recall that the saving rate is still equal to s].

Exercise 2.16. Consider the discrete-time Solow growth model with constant population growth at the rate n, no technological change, constant depreciation rate of δ and a constant saving rate s. Assume that the per capita production function is given by the following continuous but non-neoclassical function:

$$f(k) = Ak + b,$$

where $A > 0$ and $b > 0$.

(1) Explain why this production function is non-neoclassical (that is, why does it violate Assumptions 1 and 2 above?).
(2) Suppose that $sA - n - \delta = -1$ (which is possible if n is sufficiently large). Show that for any $k(0) \in (0, sb)$, the economy settles into an asymptotic cycle and continuously fluctuates between $k(0)$ and $sb - k(0)$.
(3) Now consider a more general continuous production function $f(k)$ that does not satisfy Assumptions 1 and 2, such that there exist $k_1, k_2 \in \mathbb{R}_+$ with $k_1 \neq k_2$ and

$$\begin{aligned} k_2 &= f(k_1) - (n + \delta) k_1 \\ k_1 &= f(k_2) - (n + \delta) k_2. \end{aligned}$$

Show that when such (k_1, k_2) exist, there may also exist a stable steady state.
(4) Prove that such cycles are not possible in the continuous-time Solow growth model for any (possibly non-neoclassical) continuous production function $f(k)$. [Hint: consider the equivalent of Figure 2.9 above].
(5) What does the result in parts 2 and 3 imply for the approximations of discrete time by continuous time suggested in Section 2.4?

(6) In light of your answer to part 5, what do you think of the cycles in parts 2 and 3?

Exercise 2.17. Characterize the asymptotic equilibrium of the modified Solow/AK model mentioned above, with a constant saving rate s, depreciation rate δ, no population growth and an aggregate production function of the form

$$F\left[K\left(t\right),L\left(t\right)\right]=A_{K}K\left(t\right)+A_{L}L\left(t\right).$$

Exercise 2.18. Consider the basic Solow growth model with a constant saving rate s, constant population growth at the rate n, aggregate production function given by (2.38), and no technological change.

(1) Determine conditions under which this production function satisfies Assumptions 1 and 2.
(2) Characterize the unique steady-state equilibrium when Assumptions 1 and 2 hold.
(3) Now suppose that σ is sufficiently high so that Assumption 2 does not hold. Show that in this case equilibrium behavior can be similar to that in Exercise 2.17 with sustained growth in the long run. Interpret this result.
(4) Now suppose that $\sigma\rightarrow 0$, so that the production function becomes Leontief,

$$Y\left(t\right)=\min\left\{\gamma A_{K}\left(t\right)K\left(t\right);\left(1-\gamma\right)A_{L}\left(t\right)L\left(t\right)\right\}.$$

The model is then identical to the classical Harrod-Domar growth model developed by Roy Harrod and Evsey Domar (Harrod, 1939, Domar, 1946). Show that in this case there is typically no steady-state equilibrium with full employment and no idle capital. What happens to factor prices in these cases? Explain why this case is "pathological," giving at least two reasons why we may expect equilibria with idle capital or idle labor not to apply in practice.

Exercise 2.19. * We now derive the CES production function following the method in the original article by Arrow, Chenery, Minhas and Solow (1961). These authors noted that a good empirical approximation to the relationship between income per capita and the wage rate was provided by an equation of the form

$$y=\alpha w^{-\sigma},$$

where $y=f\left(k\right)$ is again output per capita and w is the wage rate. With competitive markets, recall that $w=f\left(k\right)-kf'\left(k\right)$. Thus the above equation can be written as

$$y=\alpha\left(y-ky'\right)^{-\sigma},$$

where $y=y\left(k\right)\equiv f\left(k\right)$ and y' denotes $f'\left(k\right)$. This is a nonlinear first-order differential equation.

(1) Using separation of variables (see Appendix Chapter B), show that the solution to this equation satisfies

$$y\left(k\right)=\left[\alpha^{1/\sigma}+c_{0}k^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}},$$

where c_0 is a constant of integration.

(2) How would you put more structure on α and c_0 and derive the exact form of the CES production function in (2.38)?

Exercise 2.20. Show that the see as production function with $\sigma > 1$ violates Assumption 1 and with $\sigma \leq 1$, it satisfies Assumption 1.

Exercise 2.21. Prove Proposition 2.12.

Exercise 2.22. Prove Proposition 2.13.

Exercise 2.23. In this exercise, we work through an alternative conception of technology, which will be useful in the next chapter. Consider the basic Solow model in continuous time and suppose that $A(t) = A$, so that there is no technological progress of the usual kind. However, assume that the relationship between investment and capital accumulation is modified to

$$K(t+1) = (1-\delta) K(t) + q(t) I(t),$$

where $[q(t)]_{t=0}^{\infty}$ is an exogenously given time-varying process. Intuitively, when $q(t)$ is high, the same investment expenditure translates into a greater increase in the capital stock. Therefore, we can think of $q(t)$ as the inverse of the relative prices of machinery to output. When $q(t)$ is high, machinery is relatively cheaper. Gordon (1990) documented that the relative prices of durable machinery has been declining relative to output throughout the postwar era. This is quite plausible, especially given our recent experience with the decline in the relative price of computer hardware and software. Thus we may want to suppose that $\dot{q}(t) > 0$. This exercise asks you to work through a model with this feature based on Greenwood, Hercowitz and Krusell (1997).

(1) Suppose that $\dot{q}(t)/q(t) = \gamma_K > 0$. Show that for a general production function, $F(K, L)$, there exists no steady-state equilibrium.

(2) Now suppose that the production function is Cobb-Douglas, $F(K, L) = K^{\alpha} L^{1-\alpha}$, and characterize the unique steady-state equilibrium.

(3) Show that this steady-state equilibrium does not satisfy the Kaldor fact of constant K/Y. Is this a problem? [Hint: how is "K" measured in practice? How is it measured in this model?].

CHAPTER 3

The Solow Model and the Data

In this chapter, we will see how the Solow model or its simple extensions can be used to interpret both economic growth over time and cross-country output differences. The focus is on *proximate causes* of economic growth, that is, the factors such as investment or capital accumulation highlighted by the basic Solow model, as well as technology and human capital differences. What lies underneath these proximate causes is the topic of the next chapter.

There are multiple ways of using the basic Solow model to look at the data. I start with the growth accounting framework, which is most commonly applied for decomposing the sources of growth over time. After briefly discussing the theory of growth accounting and some of its uses, I will discuss applications of the Solow model to cross-country output differences. In this context, I will also introduce the augmented Solow model with human capital, and show how various different regression-based approaches can be motivated from this framework. Finally, I will illustrate how the growth accounting framework can be modified to a "development accounting framework" to form another bridge between the Solow model and the data. A constant theme that emerges from many of these approaches concerns the importance of productivity differences over time and across countries. The chapter ends with a brief discussion of various other approaches to estimating cross-country productivity differences.

3.1. Growth Accounting

As discussed in the previous chapter, at the center of the Solow model is the aggregate production function, (2.1), which I rewrite here in its general form:

$$Y(t) = F[K(t), L(t), A(t)]. \tag{3.1}$$

Another major contribution of Bob Solow to the study of economic growth was the observation that this production function, combined with competitive factor markets, also gives us a framework for accounting for the sources of economic growth. In particular, Solow (1957) developed what has become one of the most common tools in macroeconomics, the *growth accounting framework.*

For our purposes, it is sufficient to expose the simplest version of this framework. Consider a continuous-time economy and differentiate the production function (3.1) with respect to time. Dropping time dependence and denoting the partial derivatives of F with respect to

its arguments by F_A, F_K and F_L, this yields

$$\frac{\dot{Y}}{Y} = \frac{F_A A}{Y}\frac{\dot{A}}{A} + \frac{F_K K}{Y}\frac{\dot{K}}{K} + \frac{F_L L}{Y}\frac{\dot{L}}{L}. \tag{3.2}$$

Now denote the growth rates of output, capital stock and labor by $g \equiv \dot{Y}/Y$, $g_K \equiv \dot{K}/K$ and $g_L \equiv \dot{L}/L$, and also define

$$x \equiv \frac{F_A A}{Y}\frac{\dot{A}}{A}$$

as the contribution of technology to growth. Next, recall from the previous chapter that with competitive factor markets, we have $w = F_L$ and $R = F_K$ (eq.'s (2.6) and (2.7)) and define the factor shares as $\alpha_K \equiv RK/Y$ and $\alpha_L \equiv wL/Y$. Putting all these together, (3.2) can be written as

$$x = g - \alpha_K g_K - \alpha_L g_L. \tag{3.3}$$

This is the *fundamental growth accounting* equation. At some level it is no more than an identity. However, it also allows us to estimate the contribution of technological progress to economic growth using data on factor shares, output growth, labor force growth and capital stock growth. The contribution from technological progress, x, is typically referred to as *Total Factor Productivity* (TFP) or sometimes as Multi Factor Productivity.

In particular, denoting an estimate by "^", we have the estimate of TFP growth at time t as:

$$\hat{x}(t) = g(t) - \alpha_K(t)\, g_K(t) - \alpha_L(t)\, g_L(t). \tag{3.4}$$

I only put the "^" on x, but one may want to take into account that all of the terms on the right-hand side are also "estimates" obtained with a range of assumptions from national accounts and other data sources.

In addition, if the production function takes the standard labor-augmenting form $Y(t) = \tilde{F}[K(t), A(t) L(t)]$, then

$$\frac{\dot{A}}{A} = \frac{1}{\alpha_L}\left[g - \alpha_K g_K - \alpha_L g_L\right],$$

but this equation is not particularly useful, since $\dot{A}/A$ is not something we are inherently interested in. The economically interesting object is in fact $\hat{x}$ in (3.4), since it directly measures the effect of technological progress on output growth.

In continuous time, eq. (3.4) is exact because it is defined in terms of instantaneous changes (derivatives). In practice, instead of instantaneous changes, we look at changes over discrete time intervals, for example over a year (or sometimes with better data, over a quarter or a month). With discrete time intervals, there is a potential problem in using (3.4); over the time horizon in question, factor shares can change; should we use beginning-of-period or end-of-period values of α_K and α_L? It can be shown that the use of either beginning-of-period or end-of-period values might lead to seriously biased estimates of the contribution of TFP to output growth, $\hat{x}$. This is particularly likely when the distance between the two time periods

is large (see Exercise 3.1). The best way of avoiding such biases is to use as high-frequency data as possible.

For now, taking the available data as given, let us look at how one could use the growth accounting framework with data over discrete intervals. The most common way of dealing with the problems pointed out above is to use factor shares calculated as the average of the beginning of period and end of period values. Therefore in discrete time, for a change between times t and $t+1$, the analog of eq. (3.4) becomes

$$\hat{x}_{t,t+1} = g_{t,t+1} - \bar{\alpha}_{K,t,t+1} g_{K,t,t+1} - \bar{\alpha}_{L,t,t+1} g_{L,t,t+1}, \tag{3.5}$$

where $g_{t,t+1}$ is the growth rate of output between t and $t+1$, and other growth rates are defined analogously. Moreover,

$$\bar{\alpha}_{K,t,t+1} \equiv \frac{\alpha_K(t) + \alpha_K(t+1)}{2} \text{ and } \bar{\alpha}_{L,t,t+1} \equiv \frac{\alpha_L(t) + \alpha_L(t+1)}{2}$$

are average factor shares between t and $t+1$. Equation (3.5) would be a fairly good approximation to (3.4) when the difference between t and $t+1$ is small and the capital-labor ratio does not change much during this time interval.

Solow's (1957) article not only developed this growth accounting framework but also applied it to US data for a preliminary assessment of the "sources of growth" during the early 20th century. The question Bob Solow asked was this: how much of the growth of the US economy can be attributed to increased labor and capital inputs, and how much of it is due to the residual, "technological progress"? Solow's conclusion was quite striking: a large part of the growth was due to technological progress.

This has been a landmark finding, emphasizing the importance of technological progress as the driver of economic growth not only in theory, but also in practice. It focused the attention of economists on sources of technology differences over time, and across nations, industries and firms.

From early days, however, it was recognized that calculating the contribution of technological progress to economic growth in this manner has a number of pitfalls. Moses Abramovitz (1956), famously, dubbed the $\hat{x}$ term "the measure of our ignorance"—after all, it was the residual that we could not explain and we decided to call it "technology".

In its extreme form, this criticism is unfair, since $\hat{x}$ does correspond to technology according to eq. (3.4); thus the growth accounting framework is an example of using theory to inform measurement. Yet at another level, the criticism has validity. If we mismeasure the growth rates of labor and capital inputs, g_L and g_K, we will arrive at inflated estimates of $\hat{x}$. And in fact there are good reasons for suspecting that Solow's estimates and even the higher-quality estimates that came later may be mismeasuring the growth of inputs. The most obvious reason for this is that what matters is not labor hours, but effective labor hours, so it is important—though difficult—to make adjustments for changes in the *human capital* of workers. I discuss issues related to human capital in Section 3.3 below and then in greater detail in Chapter 10. Similarly, measurement of capital inputs is not straightforward. In the theoretical model, capital corresponds to the final good used as input to produce more

goods. But in practice, capital is machinery, and in measuring the amount of capital used in production one has to make assumptions about how relative prices of machinery change over time. The typical assumption, adopted for a long time in national accounts and also naturally in applications of the growth accounting framework, was to use capital expenditures. However, if the same machines are much cheaper today than they had been in the past (as has been the case for computers, for example), then this methodology would severely underestimate g_K (recall Exercise 2.23 in the previous chapter). Therefore, when applying eq. (3.4), underestimates of g_L and g_K will naturally inflate in the estimates of the role of technology as a source of economic growth.

What the best way of making adjustments to labor and capital inputs in order to arrive to the best estimate of technology is still a hotly debated area. Dale Jorgensen, for example, has shown that the "residual" technology can be reduced very substantially (perhaps almost to 0) by making adjustments for changes in the quality of labor and capital (see, for example, Jorgensen, Gollop and Fraumeni, 1987, or Jorgensen, 2005). These issues will also become relevant when we attempt to decompose the sources of cross-country output differences. Before doing this, let us look at applications of the Solow model to data using regression analysis.

3.2. Solow Model and Regression Analyses

Another popular approach of taking the Solow model to data is to use *growth regressions*, which involve estimating regression models with country growth rates on the left-hand side. These growth regressions have been used extensively following the work by Barro (1991). To see how these regressions are motivated and what their shortcomings are, let us return to the basic Solow model with constant population growth and labor-augmenting technological change in continuous time. Recall that, in this model, the equilibrium is described by the following equations:

$$y(t) = A(t) f(k(t)), \tag{3.6}$$

and

$$\frac{\dot{k}(t)}{k(t)} = \frac{sf(k(t))}{k(t)} - \delta - g - n, \tag{3.7}$$

where $A(t)$ is the labor-augmenting (Harrod-neutral) technology term, $k(t) \equiv K(t)/(A(t)L(t))$ is the effective capital labor ratio and $f(\cdot)$ is the per capita production function. Equation (3.7) follows from the constant technological progress and constant population growth assumptions, $\dot{A}(t)/A(t) = g$ and $\dot{L}(t)/L(t) = n$. Now differentiating (3.6) with respect to time and dividing both sides by $y(t)$,

$$\frac{\dot{y}(t)}{y(t)} = g + \varepsilon_f(k(t)) \frac{\dot{k}(t)}{k(t)}, \tag{3.8}$$

where

$$\varepsilon_f(k(t)) \equiv \frac{f'(k(t)) k(t)}{f(k(t))} \in (0,1)$$

is the elasticity of the $f(\cdot)$ function. The fact that it is between 0 and 1 follows from Assumption 1. For example, with the Cobb-Douglas technology from Example 2.1 in the previous chapter, $\varepsilon_f(k(t)) = \alpha$, that is, it is a constant independent of $k(t)$ (see Example 3.1 below). However, generally, this elasticity is a function of $k(t)$.

Now let us consider a first-order Taylor expansion of (3.7) with respect to $\log k(t)$ around the steady-state value k^* (and recall that $\partial y/\partial \log x = (\partial y/\partial x) \cdot x$). This expansion implies that for $k(t)$ in the neighborhood of k^*, we have

$$\begin{aligned}\frac{\dot{k}(t)}{k(t)} &\simeq \left(\frac{sf(k^*)}{k^*} - \delta - g - n\right) + \left(\frac{f'(k^*)k^*}{f(k^*)} - 1\right) s\frac{f(k^*)}{k^*}(\log k(t) - \log k^*). \\ &\simeq (\varepsilon_f(k^*) - 1)(\delta + g + n)(\log k(t) - \log k^*).\end{aligned}$$

The use of the symbol "$\simeq$" here is to emphasize that this is an approximation, ignoring second-order terms. In particular, the first line follows simply by differentiating $\dot{k}(t)/k(t)$ with respect to $\log k(t)$ and evaluating the derivatives at k^* (and ignoring second-order terms). The second line uses the fact that the first term in the first line is equal to zero by definition of the steady-state value k^* (recall that from eq. (2.48) in the previous chapter, $sf(k^*)/k^* = \delta + g + n$), the definition of the elasticity of the f function, $\varepsilon_f(k(t))$, and again the fact that $sf(k^*)/k^* = \delta + g + n$. Now substituting this into (3.8),

$$\frac{\dot{y}(t)}{y(t)} \simeq g - \varepsilon_f(k^*)(1 - \varepsilon_f(k^*))(\delta + g + n)(\log k(t) - \log k^*).$$

Let us define $y^*(t) \equiv A(t) f(k^*)$ as the "steady-state level of output per capita," that is, the level of per capita output that would apply if the effective capital-labor ratio were at its steady-state value and technology were at its time t level. A first-order Taylor expansion of $\log y(t)$ with respect to $\log k(t)$ around $\log k^*(t)$ then gives

$$\log y(t) - \log y^*(t) \simeq \varepsilon_f(k^*)(\log k(t) - \log k^*).$$

Combining this with the previous equation yields the following "convergence equation":

$$\frac{\dot{y}(t)}{y(t)} \simeq g - (1 - \varepsilon_f(k^*))(\delta + g + n)(\log y(t) - \log y^*(t)). \tag{3.9}$$

Equation (3.9) makes it clear that, in the Solow model, there are two sources of growth in output per capita: the first is g, the rate of technological progress, and the second is "convergence". This latter is a consequence of the impact of the gap between the current level of output per capita and the steady-state level of output per capita on the rate of capital accumulation (recall that $0 < \varepsilon_f(k^*) < 1$). Intuitively, the further is a country below its steady-state capital-labor ratio, the more capital it will accumulate and the faster it will grow. Put differently, the lower is $y(t)$ relative to $y^*(t)$, and thus the lower is $k(t)$ relative to k^*, the greater is the average product of capital $f(k^*)/k^*$, and this leads to faster growth in the effective capital-labor ratio. This pattern is already visible in Figure 2.7 in the previous chapter.

Another noteworthy feature is that the speed of convergence in eq. (3.9), measured by the term $(1 - \varepsilon_f(k^*))(\delta + g + n)$ multiplying the gap between $\log y(t)$ and $\log y^*(t)$, depends on

$\delta + g + n$, and on the elasticity of the production function, $\varepsilon_f(k^*)$. Both of these capture intuitive effects. As discussed in the previous chapter, the term $\delta + g + n$ determines the rate at which the effective capital-labor ratio needs to be replenished. The higher is this rate of replenishment, the larger is the amount of investment in the economy (recall Figure 2.7 in the previous chapter) and thus there is room for faster adjustment. On the other hand, when $\varepsilon_f(k^*)$ is high, we are close to a linear—AK—production function, and as demonstrated in the previous chapter, in this case convergence should be slow. In the extreme case where $\varepsilon_f(k^*)$ is equal to 1, the economy takes the AK form and there is no convergence.

Example 3.1. **(Cobb-Douglas Production Function and Convergence)** Consider briefly the Cobb-Douglas production function from Example 2.1 in the previous chapter, where $Y(t) = A(t) K(t)^{\alpha} L(t)^{1-\alpha}$. This implies that $y(t) = A(t) k(t)^{\alpha}$. Consequently, as noted above, $\varepsilon_f(k(t)) = \alpha$. Therefore, (3.9) becomes

$$\frac{\dot{y}(t)}{y(t)} \simeq g - (1-\alpha)(\delta + g + n)(\log y(t) - \log y^*(t)).$$

This equation also enables us to "calibrate" the speed of convergence in practice—meaning to obtain a back-of-the-envelope estimate of the speed of convergence by using plausible values of parameters. Let us focus on advanced economies. In that case, plausible values for these parameters might be $g \simeq 0.02$ for approximately 2% per year output per capita growth, $n \simeq 0.01$ for approximately 1% population growth and $\delta \simeq 0.05$ for about 5% per year depreciation. Recall also from the previous chapter that the share of capital in national income is about 1/3, so with the Cobb-Douglas production function we should have $\alpha \simeq 1/3$. Consequently, we may expect the convergence coefficient in front of $\log y(t) - \log y^*(t)$ to be around 0.054 ($\simeq 0.67 \times 0.08$). This is a very rapid rate of convergence and would imply that the income gap between two similar countries that have the same technology, the same depreciation rate and the same rate of population growth should narrow rather quickly. For example, it can be computed that with these numbers, the gap of income between two similar countries should be halved in little more than 10 years (see Exercise 3.4). This is clearly at odds with the patterns we saw in Chapter 1.

Next using a discrete-time approximation, eq. (3.9) yields the regression equation:

$$g_{i,t,t-1} = b^0 + b^1 \log y_{i,t-1} + \varepsilon_{i,t}, \tag{3.10}$$

where $g_{i,t,t-1}$ is the growth rate of country i between dates $t-1$ and t, $\log y_{i,t-1}$ is the "initial" (time $t-1$) log output per capita of this country, and $\varepsilon_{i,t}$ is a stochastic term capturing all omitted influences. Regressions on this form have been estimated by, among others, Baumol (1986), Barro (1991) and Barro and Sala-i-Martin (1992). If such an equation is estimated in the sample of core OECD countries, b^1 is indeed estimated to be negative; countries like Greece, Spain and Portugal that were relatively poor at the end of World War II have grown faster than the rest as shown in Figure 1.14 in Chapter 1.

Yet, Figure 1.13 in Chapter 1 shows, when we look at the whole world, there is no evidence for a negative b^1. Instead, this figure makes it clear that, if anything, b^1 would be positive. In other words, there is no evidence of world-wide convergence.

Barro and Sala-i-Martin refer to this type of convergence as "unconditional convergence," meaning the convergence of countries regardless of differences in characteristics and policies. However, this notion of unconditional convergence may be too demanding. It requires that there should be a tendency for the income gap between any two countries to decline, regardless of the technological opportunities, investment behavior, policies and institutions of these countries. If they do differ with respect to these factors, the Solow model would *not* predict that they should converge in income level. Instead, each should converge to their own level of steady-state income per capita or balanced growth path. Thus in a world where countries differ according to their characteristics, a more appropriate regression equation may take the form:

$$g_{i,t,t-1} = b_i^0 + b^1 \log y_{i,t-1} + \varepsilon_{i,t}, \tag{3.11}$$

where the key difference is that now the constant term, b_i^0, is country specific. (In principle, the slope term, measuring the speed of convergence, b^1, should also be country specific, but in empirical work, this is generally taken to be a constant, and I assume the same here to simplify the exposition). One may then model b_i^0 as a function of country characteristics.

If the true equation is (3.11), in the sense that the Solow model applies but certain determinants of economic growth differ across countries, eq. (3.10) would not be a good fit to the data. Put differently, there is no guarantee that the estimates of b^1 resulting from this equation will be negative. In particular, it is natural to expect that $Cov\left(b_i^0, \log y_{i,t-1}\right) < 0$ (where Cov refers to the population covariance), since economies with certain growth-reducing characteristics will have low levels of output. This implies a negative bias in the estimate of b^1 in eq. (3.10), when the more appropriate equation is (3.11).

With this motivation, Barro (1991) and Barro and Sala-i-Martin (1992, 2004) favor the notion of "conditional convergence," which means that the convergence effects emphasized by the Solow model should lead to negative estimates of b^1 once b_i^0 is allowed to vary across countries. To implement this idea of conditional convergence empirically, they estimate models where b_i^0 is assumed to be a function of, among other things, the male schooling rate, the female schooling rate, the fertility rate, the investment rate, the government-consumption ratio, the inflation rate, changes in terms of trades, openness and institutional variables such as rule of law and democracy. In regression form, this can be written as

$$g_{i,t,t-1} = \mathbf{X}_{i,t}^T \boldsymbol{\beta} + b^1 \log y_{i,t-1} + \varepsilon_{i,t}, \tag{3.12}$$

where $\mathbf{X}_{i,t}$ is a (column) vector including the variables mentioned above (as well as a constant), with a vector of coefficients $\boldsymbol{\beta}$ (recall that $\mathbf{X}^T$ denotes the transpose of $\mathbf{X}$). In other words, this specification imposes that b_i^0 in eq. (3.11) can be approximated by $\mathbf{X}_{i,t}^T \boldsymbol{\beta}$. Consistent with the emphasis on conditional convergence, regressions of eq. (3.12) tend to show a

negative estimate of b^1, but the magnitude of this estimate is much lower than that suggested by the computations in Example 3.1.

Regressions similar to (3.12) have not only been used to support "conditional convergence," but also to estimate the "determinants of economic growth". In particular, it may appear natural to presume that the estimates of the coefficient vector $\boldsymbol{\beta}$ will contain information about the *causal effects* of various variables on economic growth. For example, the fact that the schooling variables enter with positive coefficients in the estimates of regression (3.12) is interpreted as evidence that "schooling causes growth". The simplicity of the regression equations of the form (3.12) and the fact that they create an attractive bridge between theory and data have made them very popular over the past two decades.

Nevertheless, there are several problematic features with regressions of this form. These include:

1. Most, if not all, of the variables in $\mathbf{X}_{i,t}$ as well as $\log y_{i,t-1}$, are *econometrically endogenous* in the sense that they are jointly determined with the rate of economic growth between dates $t-1$ and t. For example, the same factors that make a country relatively poor in 1950, thus reducing $\log y_{i,t-1}$, should also affect its growth rate after 1950. Or the same factors that make a country invest little in physical and human capital could have a direct effect on its growth rate (through other channels such as its technology or the efficiency with which the factors of production are being utilized). This creates an obvious source of bias (and lack of econometric consistency) in the regression estimates of the coefficients. This bias makes it unlikely that the effects captured in the coefficient vector $\boldsymbol{\beta}$ correspond to causal effects of these characteristics on the growth potential of economies. One may argue that the convergence coefficient b^1 is of interest, even if it does not have a "causal interpretation". This argument is not entirely compelling, either. A basic result in econometrics is that if $\mathbf{X}_{i,t}$ is econometrically endogenous, so that the parameter vector $\boldsymbol{\beta}$ is estimated inconsistently, the estimate of the parameter b^1 will also be inconsistent unless $\mathbf{X}_{i,t}$ is independent from $\log y_{i,t-1}$.[1] This makes the estimates of the convergence coefficient, b^1, difficult to interpret.

2. Even if the $\mathbf{X}_{i,t}$'s were econometrically exogenous, a negative coefficient estimate for b^1 could be caused by other econometric problems, such as measurement error or other transitory shocks to $y_{i,t}$. For example, because national accounts data are measured with error, data on $\log y_{i,t}$ will naturally contain measurement error. Suppose, for example, that we only observe estimates of output per capita $\tilde{y}_{i,t} = y_{i,t} \exp(u_{i,t})$, where $y_{i,t}$ is the true output per capita and $u_{i,t}$ is a random and serially uncorrelated error term. When the variable $\log \tilde{y}_{i,t}$ is used in the regression, the error term $u_{i,t-1}$ appears both on the left-hand side and the right-hand side of (3.12). In particular, note that

$$\log \tilde{y}_{i,t} - \log \tilde{y}_{i,t-1} = \log y_{i,t} - \log y_{i,t-1} + u_{i,t} - u_{i,t-1}.$$

Since the measured growth is $\tilde{g}_{i,t,t-1} \approx \log \tilde{y}_{i,t} - \log \tilde{y}_{i,t-1} = \log y_{i,t} - \log y_{i,t-1} + u_{i,t} - u_{i,t-1}$, the measurement error $u_{i,t-1}$ will be part of both the error term $\varepsilon_{i,t}$ and the right-hand side

[1] An example of the potential endogeneity of these variables will be given in Section 3.4 below.

variable $\log \tilde{y}_{i,t-1} = \log y_{i,t-1} + u_{i,t-1}$ in the regression equation

$$\tilde{g}_{i,t,t-1} = \mathbf{X}_{i,t}^T \boldsymbol{\beta} + b^1 \log \tilde{y}_{i,t-1} + \varepsilon_{i,t}.$$

This will naturally lead a negative bias in the estimation of b^1. Therefore, we can end up with a negative estimate of b^1, *even when there is no conditional convergence.*

3. The economic interpretation of regression equations like (3.12) is not always straightforward either. Many of the regressions used in the literature include the investment rate as part of the vector $\mathbf{X}_{i,t}$ (and all of them include the schooling rate). However, in the Solow model, differences in investment rates are *the* channel via which convergence takes place—in the sense that economies below their steady state grow faster by having higher investment rates. Therefore, according to the Solow model, conditional on the investment rate there should be no further effect of the gap between the current level of output and the steady-state level of output. The same concern applies to the interpretation of the effects of the variables in $\mathbf{X}_{i,t}$ (such as institutions or openness), which are typically included as potential determinants of economic growth. However, many of these variables affect growth primarily by affecting the investment rate (or the schooling rate). Therefore, once we condition on the investment rate and the schooling rate, the coefficients on these variables no longer measure their impact on economic growth. Consequently, estimates of (3.12) with investment-like variables on the right-hand side are difficult to link to theory.

4. Regressions of the form (3.12) are often thought to be appealing because they model the "process of economic growth" and may give us information about the potential determinants of economic growth (to the extent that these are included in the vector $\mathbf{X}_{i,t}$). Nevertheless, there is a sense in which growth regressions are not much different than "levels regressions"—similar to those discussed in Section 3.4 below and further in Chapter 4. In particular, again noting that $g_{i,t,t-1} \approx \log y_{i,t} - \log y_{i,t-1}$, eq. (3.12) can be rewritten as

$$\log y_{i,t} \approx \mathbf{X}_{i,t}^T \boldsymbol{\beta} + \left(1 + b^1\right) \log y_{i,t-1} + \varepsilon_{i,t}.$$

Therefore, essentially the level of output is being regressed on $\mathbf{X}_{i,t}$, so that regression of the form (3.12) will typically uncover the correlations between the variables in the vector $\mathbf{X}_{i,t}$ and output per capita. Thus growth regressions are often informative about which economic or social processes go hand-in-hand with high levels of output. A specific example is life expectancy. When life expectancy is included in the vector $\mathbf{X}_{i,t}$ in a growth regression, it is often highly significant. But this simply reflects the high correlation between income per capita and life expectancy (recall Figure 1.6 in Chapter 1). It does not imply that life expectancy has a causal effect on economic growth, and more likely, it reflects the joint determination of life expectancy and income per capita and the impact of prosperity on life expectancy.

5. Finally, the motivating equation for the growth regression, (3.9), is derived for a closed Solow economy. When we look at cross-country income differences or growth experiences, the use of this equation imposes the assumption that "each country is an island". In other words, the world is being interpreted as a collection of *non-interacting* closed economies. In

practice, countries trade goods, exchange ideas and borrow and lend in international financial markets. This implies that the behavior of different countries will not be given by eq. (3.9), but by a system of equations characterizing the joint world equilibrium. Even though a world equilibrium is typically a better way of representing differences in income per capita and their evolution, in the first part of the book I follow the usual approach of treating "each country as an island," which is simpler and thus a useful starting point to develop the main ideas. I return to models of world equilibrium in Chapters 18 and 19.

This discussion suggests that growth regressions need to be used with caution and ought to be supplemented with other methods of mapping the basic Solow model to data. What other methods will be useful in empirical analyses of economic growth? A number of answers emerge from our discussion above. To start with, for many of the central questions of economic growth, regressions that control for fixed country characteristics might be equally useful as, or more useful than, cross-sectional growth regressions. For example, a more natural regression framework for investigating the economic (or statistical) relationship between the variables in the vector $\mathbf{X}_{i,t}$ and economic growth might be

$$\log y_{i,t} = \alpha \log y_{i,t-1} + \mathbf{X}_{i,t}^T \boldsymbol{\beta} + \delta_i + \mu_t + \varepsilon_{i,t}, \tag{3.13}$$

where δ_i's denote a full set of country fixed effects and μ_t's denote a full set of year effects. This regression framework differs from the growth regressions in a number of respects. First, the regression equation is specified in levels rather than with the growth rate on the left-hand side. But this is mainly a rearrangement of (3.13). More importantly, by including the country fixed effects, this regression equation takes out fixed country characteristics that might be simultaneously affecting economic growth (or the level of income per capita) and the right-hand side variables of interest. For instance, in terms of the example discussed in (4) above, if life expectancy and income per capita are correlated because of some omitted factors, the country fixed effects, the δ_i's, will remove the influence of these factors. Therefore, panel data regressions as in (3.13) may be more informative about the relationship between a range of factors and income per capita. Nevertheless, it is important to emphasize that including country fixed effects is not a panacea against all omitted variable biases and econometric endogeneity problems. Simultaneity bias often results from time-varying influences, which cannot be removed by including fixed effects. Moreover, to the extent that some of the variables in the vector $\mathbf{X}_{i,t}$ are slowly-varying themselves, the inclusion of country fixed effects will make it difficult to uncover the statistical relationship between these variables and income per capita, and may increase potential biases due to measurement error in the right-hand side variables.

This discussion highlights that econometric models similar to (3.13) are often useful and a good complement to (or substitute for) growth regressions. But they are not a good substitute for specifying the structural economic relationships fully and estimating the causal relationships of interest. The next chapter shows how some progress can be made in this regard by looking at the historical determinants of long-run economic growth and using

specific historical episodes to generate potential sources of exogenous variation in the variables of interest that can allow us to develop an instrumental-variables strategy.

In the remainder of this chapter, I discuss how the structure of the Solow model can be further exploited to look at the data. Before doing this, I present an augmented version of the Solow model incorporating human capital, which is useful in these empirical exercises.

3.3. The Solow Model with Human Capital

Human capital is a term we use to represent the stock of skills, education, competencies and other productivity-enhancing characteristics embedded in labor. Put differently, human capital represents the efficiency units of labor embedded in raw labor hours.

The term "human capital" originates from the observation that individuals will invest in their skills, competencies and earning capacities in the same way as firms invest in their physical capital—to increase their productivity. The seminal work by Ted Schultz, Jacob Mincer and Gary Becker brought the notion of human capital to the forefront of economics. For now, all we need to know is that labor hours supplied by different individuals do not contain the same efficiency units; a highly trained carpenter can produce a chair in a few hours, while an amateur would spend many more hours to perform the same task. Economists capture this notion by thinking that the trained carpenter has more efficiency units of labor embedded in the labor hours he supplies, or alternatively he has more human capital. The theory of human capital is very rich and some of the important notions will be discussed in Chapter 10. For now, our objective is more modest, to investigate how including human capital makes the Solow model a better fit to the data. The inclusion of human capital will enable us to embed all three of the main *proximate sources* of income differences; physical capital, human capital and technology.

For the purposes of this section, let us focus on the continuous-time economy and suppose that the aggregate production function of the economy is given by a variant of eq. (2.1):

$$Y = F\left(K, H, AL\right), \tag{3.14}$$

where H denotes "human capital". Notice that this production function is somewhat unusual, since it separates human capital H from labor L as potential factors of production. I start with this form because it is used commonly in the growth literature. The more micro-founded models in Chapter 10 will assume that human capital is embedded in workers. How human capital is measured in the data will be discussed below.

Let us also modify Assumptions 1 and 2 as follows

Assumption 1′: The production function $F : \mathbb{R}^3_+ \to \mathbb{R}_+$ in (3.14) is twice differentiable in K, H and L, and satisfies

$$\frac{\partial F\left(K, H, AL\right)}{\partial K} > 0, \quad \frac{\partial F\left(K, H, AL\right)}{\partial H} > 0, \quad \frac{\partial F\left(K, H, AL\right)}{\partial L} > 0$$

$$\frac{\partial^2 F\left(K, H, AL\right)}{\partial K^2} < 0, \quad \frac{\partial^2 F\left(K, H, AL\right)}{\partial H^2} < 0, \quad \frac{\partial^2 F\left(K, H, AL\right)}{\partial L^2} < 0.$$

Moreover, F exhibits constant returns to scale in its three arguments.

Assumption 2′: F satisfies the Inada conditions

$$\begin{aligned}
\lim_{K\to 0}\frac{\partial F(K,H,AL)}{\partial K} &= \infty \text{ and } \lim_{K\to\infty}\frac{\partial F(K,H,AL)}{\partial K}=0 \text{ for all } H>0 \text{ and } AL>0,\\
\lim_{H\to 0}\frac{\partial F(K,H,AL)}{\partial H} &= \infty \text{ and } \lim_{H\to\infty}\frac{\partial F(K,H,AL)}{\partial H}=0 \text{ for all } K>0 \text{ and } AL>0,\\
\lim_{L\to 0}\frac{\partial F(K,H,AL)}{\partial L} &= \infty \text{ and } \lim_{L\to\infty}\frac{\partial F(K,H,AL)}{\partial L}=0 \text{ for all } K,\ H,A>0.
\end{aligned}$$

Let us assume that investments in human capital take a similar form to investments in physical capital; households save a fraction s_k of their income to invest in physical capital and a fraction s_h to invest in human capital. Human capital also depreciates in the same way as physical capital, and we denote the depreciation rates of physical and human capital by δ_k and δ_h, respectively.

There is again constant population growth and a constant rate of labor-augmenting technological progress, that is,

$$\frac{\dot{L}(t)}{L(t)}=n \text{ and } \frac{\dot{A}(t)}{A(t)}=g.$$

Now defining effective human and physical capital ratios as

$$k(t)\equiv\frac{K(t)}{A(t)L(t)} \text{ and } h(t)\equiv\frac{H(t)}{A(t)L(t)},$$

and using the constant returns to scale feature in Assumption 1′, output per effective unit of labor can be written as

$$\begin{aligned}
\hat{y}(t) &\equiv \frac{Y(t)}{A(t)L(t)}\\
&= F\left(\frac{K(t)}{A(t)L(t)},\frac{H(t)}{A(t)L(t)},1\right)\\
&\equiv f(k(t),h(t)).
\end{aligned}$$

With the same steps as in Chapter 2, the law of motion of $k(t)$ and $h(t)$ can then be obtained as:

$$\begin{aligned}
\dot{k}(t) &= s_k f(k(t),h(t))-(\delta_k+g+n)k(t),\\
\dot{h}(t) &= s_h f(k(t),h(t))-(\delta_h+g+n)h(t).
\end{aligned}$$

A steady-state equilibrium is now defined not only in terms of effective capital-labor ratio, but effective human and physical capital ratios, (k^*,h^*), which satisfy the following two equations:

$$s_k f(k^*,h^*)-(\delta_k+g+n)k^*=0, \tag{3.15}$$

and

$$s_h f(k^*,h^*)-(\delta_h+g+n)h^*=0. \tag{3.16}$$

As in the basic Solow model, the focus is on steady-state equilibria with $k^*>0$ and $h^*>0$ (if $f(0,0)=0$, then there exists a trivial steady state with $k=h=0$, which I ignore for the same reasons as in the previous chapter).

Let us first prove that this steady-state equilibrium is unique. To see this heuristically, consider Figure 3.1, which is drawn in the (k, h) space. The two curves represent the two equations (3.15) and (3.16). Both lines are upward sloping. For example, in (3.15) a higher level of h^* implies greater $f(k^*, h^*)$ from Assumption $1'$, thus the level of k^* and that will satisfy the equation is higher. The same reasoning applies to (3.16). However, the proof of the next proposition shows that (3.16) is always shallower in the (k, h) space, so the two curves can only intersect once.

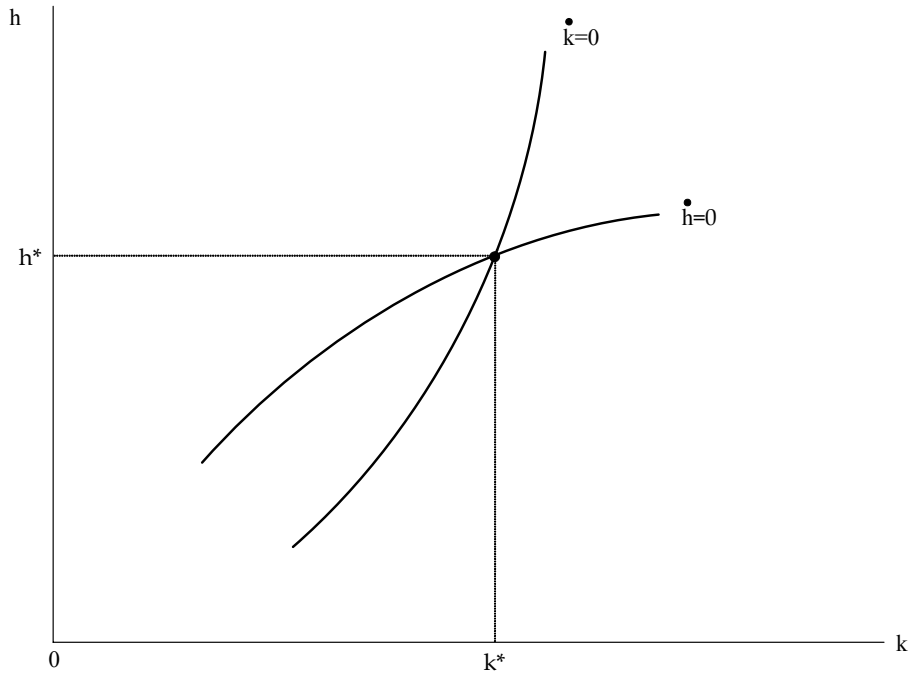

FIGURE 3.1. Steady-state equilibrium in the Solow model with human capital.

PROPOSITION 3.1. *Suppose Assumptions 1′ and 2′ are satisfied. Then, in the augmented Solow model with human capital, there exists a unique steady-state equilibrium* (k^*, h^*).

PROOF. First consider the slope of the curve (3.15), corresponding to the $\dot{k} = 0$ locus, in the (k, h) space. Using the Implicit Function Theorem, we have

$$\left.\frac{dh}{dk}\right|_{\dot{k}=0} = \frac{(\delta_k + g + n) - s_k f_k(k^*, h^*)}{s_k f_h(k^*, h^*)}, \tag{3.17}$$

where $f_k \equiv \partial f / \partial k$. Rewriting (3.15), we have $s_k f(k^*, h^*)/k^* - (\delta_k + g + n) = 0$. Now recall that since f is strictly concave in k in view of Assumption $1'$ and $f(0, h^*) \geq 0$,

$$\begin{aligned} f(k^*, h^*) &> f_k(k^*, h^*) k^* + f(0, h^*) \\ &> f_k(k^*, h^*) k^*. \end{aligned}$$

Therefore, $(\delta_k + g + n) - s_k f_k (k^*, h^*) > 0$, and (3.17) is strictly positive.

Similarly, defining $f_h \equiv \partial f / \partial h$ and applying the Implicit Function Theorem to the $\dot{h} = 0$ locus, (3.16),

$$\left.\frac{dh}{dk}\right|_{\dot{h}=0} = \frac{s_h f_k (k^*, h^*)}{(\delta_h + g + n) - s_h f_h (k^*, h^*)}. \tag{3.18}$$

With the same argument as that used for (3.17), this expression is also strictly positive.

Next, we prove that (3.17) is steeper than (3.18) whenever (3.15) and (3.16) hold, so that can it most be one intersection. First, observe that

$$\begin{aligned}
\left.\frac{dh}{dk}\right|_{\dot{h}=0} &< \left.\frac{dh}{dk}\right|_{\dot{k}=0} \\
&\Updownarrow \\
\frac{s_h f_k (k^*, h^*)}{(\delta_h + g + n) - s_h f_h (k^*, h^*)} &< \frac{(\delta_k + g + n) - s_k f_k (k^*, h^*)}{s_k f_h (k^*, h^*)} \\
&\Updownarrow \\
s_k s_h f_k f_h &< s_k s_h f_k f_h + (\delta_h + g + n)(\delta_k + g + n) \\
&\quad - (\delta_h + g + n) s_k f_k - (\delta_k + g + n) s_h f_h.
\end{aligned}$$

Now using (3.15) and (3.16) and substituting for $(\delta_k + g + n) = s_k f (k^*, h^*) / k^*$ and $(\delta_h + g + n) = s_h f (k^*, h^*) / h^*$, this is equivalent to

$$f (k^*, h^*) > f_k (k^*, h^*) k^* + f_h (k^*, h^*) h^*,$$

which is satisfied by the fact that $f (k^*, h^*)$ is a strictly concave function.

Finally, to establish existence note that Assumption 2′ implies that $\lim_{h \to 0} f (k, h) / h = \infty$, $\lim_{k \to 0} f (k, h) / k = \infty$, $\lim_{h \to \infty} f (k, h) / h = 0$ and $\lim_{k \to \infty} f (k, h) / k = 0$, so that the curves look as in Figure 3.1, that is, (3.15) is below (3.16) as $k \to 0$ and $h \to \infty$, but it is above (3.16) as $k \to \infty$ and $h \to 0$. This implies that the two curves must intersect at least once. □

This proposition shows that a unique steady state exists when the Solow model is augmented with human capital. The comparative statics are similar to the basic Solow model (see Exercise 3.7). Most importantly, both greater s_k and greater s_h will translate into higher normalized output per capita, $\hat{y}^*$.

Now turning to cross-country behavior, consider a set of countries experiencing the same rate of labor-augmenting technological progress, g. Then, those with greater propensity to invest in physical and human capital will bthise relatively richer. This is the type of prediction that can be investigated empirically to see whether the augmented Solow model (with similar technological possibilities across countries) provides us with a useful way of looking at cross-country income differences.

Before doing this, however, let us verify that the unique steady state is globally stable. The next proposition shows that this is the case.

PROPOSITION 3.2. *Suppose Assumptions 1′ and 2′ are satisfied. Then, the unique steady-state equilibrium of the augmented Solow model with human capital,* (k^*, h^*)*, is globally stable in the sense that starting with any* $k(0) > 0$ *and* $h(0)$*, we have* $(k(t), h(t)) \to (k^*, h^*)$.

PROOF. See Exercise 3.6. □

Figure 3.2 gives the intuition for this result, by showing the law of motion of k and h depending on whether the economy is above or below the two curves representing the loci for $\dot{k} = 0$ and $\dot{h} = 0$, respectively, (3.15) and (3.16). When we are to the right of the (3.15) curve, there is too much physical capital relative to the amount of labor and human capital, and consequently, $\dot{k} < 0$. When we are to its left, we are in the converse situation and $\dot{k} > 0$. Similarly, when we are above the (3.16) curve, there is too little human capital relative to the amount of labor and physical capital, and thus $\dot{h} > 0$. When we are below it, $\dot{h} < 0$. Given these arrows, the global stability of the dynamics follows.

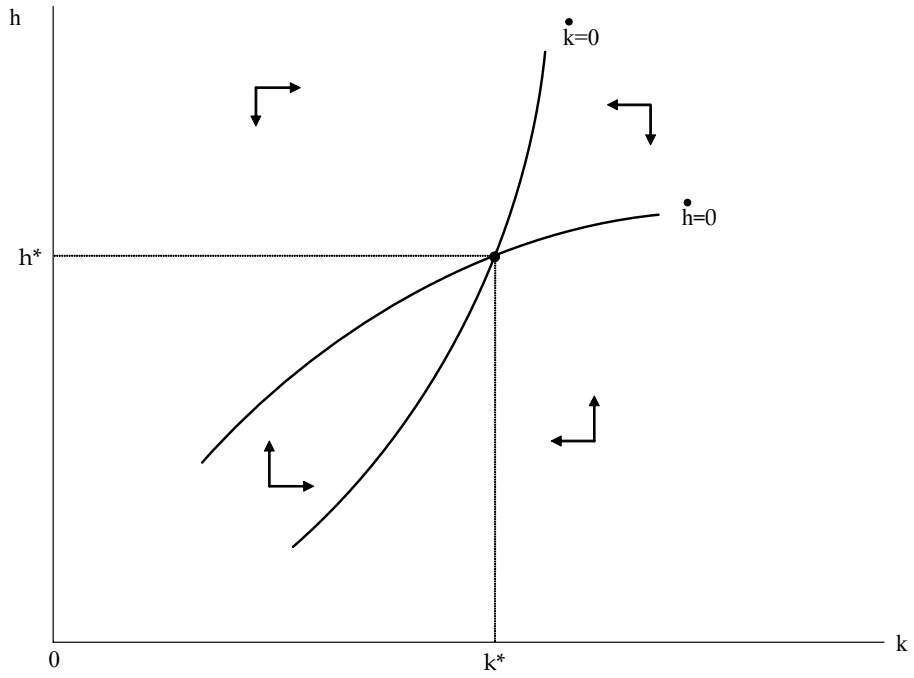

FIGURE 3.2. Dynamics of physical capital-labor and human capital-labor ratios in the Solow model with human capital.

EXAMPLE 3.2. **(Augmented Solow Model with Cobb-Douglas Production Functions)** Let us now work through a special case of the above model with Cobb-Douglas production function. In particular, suppose that the aggregate production function is

$$Y(t) = K(t)^{\alpha} H(t)^{\beta} (A(t) L(t))^{1-\alpha-\beta}, \tag{3.19}$$

where $0 < \alpha < 1$, $0 < \beta < 1$ and $\alpha + \beta < 1$. Output per effective unit of labor can then be written as

$$\hat{y}(t) = k^{\alpha}(t) h^{\beta}(t),$$

with the same definition of $\hat{y}(t)$, $k(t)$ and $h(t)$ as above. Using this functional form, (3.15) and (3.16) give the unique steady-state equilibrium as

$$k^* = \left(\left(\frac{s_k}{n+g+\delta_k} \right)^{1-\beta} \left(\frac{s_h}{n+g+\delta_h} \right)^{\beta} \right)^{\frac{1}{1-\alpha-\beta}} \tag{3.20}$$

$$h^* = \left(\left(\frac{s_k}{n+g+\delta_k} \right)^{\alpha} \left(\frac{s_h}{n+g+\delta_h} \right)^{1-\alpha} \right)^{\frac{1}{1-\alpha-\beta}},$$

which shows that a higher saving rate in physical capital not only increases k^*, but also h^*. The same applies for a higher saving rate in human capital. This reflects the facts that the higher saving rate in physical capital, by increasing, k^*, raises overall output and thus the amount invested in schooling (since s_h is constant). Given (3.20), output per effective unit of labor in steady state is obtained as

$$\hat{y}^* = \left(\frac{s_k}{n+g+\delta_k} \right)^{\frac{\alpha}{1-\alpha-\beta}} \left(\frac{s_h}{n+g+\delta_h} \right)^{\frac{\beta}{1-\alpha-\beta}}. \tag{3.21}$$

This expression shows that the relative contributions of the saving rates for physical and human capital on (normalized) output per capita depends on the shares of physical and human capital—the larger is α, the more important is s_k and the larger is β, the more important is s_h.

3.4. Solow Model and Cross-Country Income Differences: Regression Analyses

3.4.1. A World of Augmented Solow Economies. An important paper by Mankiw, Romer and Weil (1992) used regression analysis to take the augmented Solow model, with human capital, to data. In line with our main emphasis here, let us focus on the cross-country part of Mankiw, Romer and Weil's analysis. To do this, we will use the Cobb-Douglas model in Example 3.2 and envisage a world consisting of $j = 1, ..., N$ countries.

Mankiw, Romer and Weil, like many other authors, start with the assumption mentioned above, that "each country is an island"; in other words, they assume that countries do not interact (perhaps except for sharing some common technology growth, see below). This assumption enables us to analyze the behavior of each economy as a self-standing Solow model. Even though "each country is an island" is an unattractive assumption, it is a useful starting point both because of its simplicity and because this is where much of the literature started from (and in fact, it is still where much of the literature stands).

Following Example 3.2, let us assume that country $j = 1, ..., N$ has the aggregate production function:

$$Y_j(t) = K_j(t)^{\alpha} H_j(t)^{\beta} (A_j(t) L_j(t))^{1-\alpha-\beta}.$$

This production function nests the basic Solow model without human capital when $\alpha = 0$. First, assume that countries differ in terms of istheir saving rates, $s_{k,j}$ and $s_{h,j}$, population growth rates, n_j, and technology growth rates $\dot{A}_j(t)/A_j(t) = g_j$. As usual, define $k_j \equiv K_j/A_jL_j$ and $h_j \equiv H_j/A_jL_j$.

Since our main interest here is cross-country income differences, rather than studying the dynamics of a particular country over time, let us focus on a world in which each country is in steady state (thus ignoring convergence dynamics, which was the focus in the previous section). To the extent that countries are not too far from their steady state, there will be little loss of insight from this assumption, though naturally this approach will not be satisfactory when we think of countries experiencing very large growth spurts or growth collapses, as in some of the examples discussed in Chapter 1.

Given the steady-state assumption, equivalents of eq.'s (3.20) apply here and imply that the steady state physical and human capital to effective labor ratios of country j, $\left(k_j^*, h_j^*\right)$, are given by:

$$\begin{aligned} k_j^* &= \left(\left(\frac{s_{k,j}}{n_j+g_j+\delta_k}\right)^{1-\beta}\left(\frac{s_{h,j}}{n_j+g_j+\delta_h}\right)^{\beta}\right)^{\frac{1}{1-\alpha-\beta}} \\ h_j^* &= \left(\left(\frac{s_{k,j}}{n_j+g_j+\delta_k}\right)^{\alpha}\left(\frac{s_{h,j}}{n_j+g_j+\delta_h}\right)^{1-\alpha}\right)^{\frac{1}{1-\alpha-\beta}}. \end{aligned}$$

Consequently, using (3.21), the "steady-state" income per capita of country j can be written as

$$\begin{aligned} y_j^*(t) &\equiv \frac{Y(t)}{L(t)} \\ &= A_j(t)\left(\frac{s_{k,j}}{n_j+g_j+\delta_k}\right)^{\frac{\alpha}{1-\alpha-\beta}}\left(\frac{s_{h,j}}{n_j+g_j+\delta_h}\right)^{\frac{\beta}{1-\alpha-\beta}}. \end{aligned} \tag{3.22}$$

Here $y_j^*(t)$ stands for output per capita of country j along the balanced growth path. An immediate implication of this equation is that if g_j's are not equal across countries, income per capita will diverge, since the terms in front, the $A_j(t)$'s, will be growing at different rates for different countries. As discussed in Chapter 1, there is some evidence for this type of divergent behavior, but the world (per capita) income distribution can also be approximated by a relatively stable distribution. Recall that this is an area of current research and there is an active debate on whether the world economy in the postwar era should be modeled as having an expanding or a stable distribution of income per capita. The former would be consistent with a specification in which the g_j's differ across countries, while the latter would require all countries to have the same rate of technological progress, g (recall the discussion in Chapter 1).

Since technological progress is taken as exogenous in the Solow model, it is, in many ways, more appropriate for the Solow model to assume a common rate of technical progress. Motivated by this, Mankiw, Romer and Weil make the following assumption:

Common technology advances assumption: $A_j(t) = \bar{A}_j \exp(gt)$.

That is, countries differ according to their technology *level*, in particular, according to their initial level of technology, $\bar{A}_j$, but they share the same common technology growth rate, g.

Now using this assumption together with (3.22) and taking logs, we obtain the following convenient log-linear equation for the balanced growth path of income for country $j = 1, ..., N$:

$$\ln y_j^*(t) = \ln \bar{A}_j + gt + \frac{\alpha}{1-\alpha-\beta} \ln\left(\frac{s_{k,j}}{n_j + g + \delta_k}\right) + \frac{\beta}{1-\alpha-\beta} \ln\left(\frac{s_{h,j}}{n_j + g + \delta_h}\right). \tag{3.23}$$

This is a simple and attractive equation, and can be estimated easily with cross-country data. Estimates for $s_{k,j}$, $s_{h,j}$ and n_j can be computed from the available data, and combined with values for the constants δ_k, δ_h and g, these can be used to construct measures of the two key right-hand side variables. Given these measures, eq. (3.23) can be estimated by ordinary least squares (by regressing income per capita on these measures) to uncover the values of α and β.

Mankiw, Romer and Weil take $\delta_k = \delta_h = \delta$ and $\delta + g = 0.05$ as approximate depreciation rates for physical and human capital and as the growth rate for the world economy. These numbers are somewhat arbitrary, but their exact values are not important for the estimation. The literature typically approximates $s_{k,j}$ with average investment rates (investments/GDP). Investment rates, average population growth rates n_j, and log output per capita are from the Summers-Heston dataset discussed in Chapter 1. In addition, they use estimates of the fraction of the school-age population that is enrolled in secondary school as a measure of the investment rate in human capital, $s_{h,j}$. I return to a discussion of this variable below.

However, even with all of these assumptions, eq. (3.23) can still not be estimated consistently. This is because the $\ln \bar{A}_j$ term is unobserved (at least to the econometrician) and thus will be captured by the error term. Most reasonable models of economic growth would suggest that technological differences, the $\ln \bar{A}_j$'s, should be correlated with investment rates in physical and human capital. Thus an estimation of (3.23) would lead to the most standard form of omitted variable bias and inconsistent estimates. Consistency would only follow under a stronger assumption than the common technology advances assumption introduced above. Therefore, implicitly, Mankiw, Romer and Weil make another *crucial* assumption:

Orthogonal technology assumption: $\bar{A}_j = \varepsilon_j A$, with ε_j orthogonal to all other variables.

Under the orthogonal technology assumption, $\ln \bar{A}_j$, which is part of the error term, is orthogonal to the key right-hand side variables and eq. (3.23) can be estimated consistently.

3.4.2. Mankiw, Romer and Weil Estimation Results. Mankiw, Romer and Weil first estimate eq. (3.23) without the human capital term (that is, imposing $\alpha = 0$) for the cross-sectional sample of non-oil producing countries. In particular, their estimating equation in this case is:

$$\ln y_j^* = \text{constant} + \frac{\alpha}{1-\alpha} \ln(s_{k,j}) - \frac{\alpha}{1-\alpha} \ln(n_j + g + \delta_k) + \varepsilon_j.$$

This equation is obtained from (3.23) by setting $\beta = 0$ and specializing it to a single cross section. In addition, the terms $\ln(s_{k,j})$ and $\ln(n_j + g + \delta_k)$ are separated to test the restriction that their coefficients should be equal in absolute value and of opposite signs. Finally, this equation also includes ε_j as an error term, capturing all omitted factors and influences on income per capita.

Their results from this estimation exercise are replicated in columns 1 of Table 3.1 using the original Mankiw, Romer and Weil data (standard errors in parentheses). Their estimates suggest a coefficient of around 1.4 for $\alpha/(1-\alpha)$, which implies that α must be around 2/3. Since α is also the share of capital in national income, it should be around 1/3. Thus, the regression estimates without human capital appear to lead to overestimates of α. Columns 2 and 3 report the same results with updated data. The fit on the model is slightly less good than was the case with the Mankiw, Romer and Weil data, but the general pattern is similar. The implied values of α are also a little smaller than the original estimates, but still substantially higher than the 1/3 number one would expect on the basis of the underlying model.

Table 3.1
Estimates of the Basic Solow Model

	MRW	Updated data	
	1985	1985	2000
$\ln(s_k)$	1.42	1.01	1.22
	(.14)	(.11)	(.13)
$\ln(n+g+\delta)$	-1.97	-1.12	-1.31
	(.56)	(.55)	(.36)
Adj R^2	.59	.49	.49
Implied α	.59	.50	.55
No. of observations	98	98	107

The most natural reason for the high implied values of the parameter α in Table 3.1 is that ε_j is correlated with $\ln(s_{k,j})$, either because the orthogonal technology assumption is not a good approximation to reality or because there are also human capital differences correlated with $\ln(s_{k,j})$—so that there is an omitted variable bias.

Mankiw, Romer and Weil favor the second interpretation and estimate the augmented model, in particular the equation

$$\begin{aligned}(3.24) \qquad \ln y_j^* &= \text{constant} + \frac{\alpha}{1-\alpha-\beta}\ln(s_{k,j}) - \frac{\alpha}{1-\alpha-\beta}\ln(n_j+g+\delta_k) \\ &\quad + \frac{\beta}{1-\alpha-\beta}\ln(s_{h,j}) - \frac{\beta}{1-\alpha-\beta}\ln(n_j+g+\delta_h) + \varepsilon_j.\end{aligned}$$

This requires a proxy for $\ln(s_{h,j})$. Mankiw, Romer and Weil use the fraction of the working age population that is in school. With this proxy and again under the orthogonal technology assumption, the original Mankiw, Romer and Weil estimates are given in column 1 of Table 3.2. Now the estimation is more successful. Not only is the Adjusted R^2 quite high (about 78%), the implied value for α is around 1/3. On the basis of this estimation result, Mankiw, Romer and Weil and others have interpreted the fit of the augmented Solow model to the data as a success: with common technology, human and physical capital investments appear to explain 78% of the cross-country income per capita differences and the implied parameter values are reasonable. Columns 2 and 3 of the table show the results with the updated data. The implied values of α are similar, though the Adjusted R^2 is somewhat lower.

Table 3.2
Estimates of the Augmented Solow Model

	MRW	Updated data	
	1985	1985	2000
$\ln(s_k)$	.69	.65	.96
	(.13)	(.11)	(.13)
$\ln(n+g+\delta)$	-1.73	-1.02	-1.06
	(.41)	(.45)	(.33)
$\ln(s_h)$	.66	.47	.70
	(.07)	(.07)	(.13)
Adj R^2	.78	.65	.60
Implied α	.30	.31	.36
Implied β	.28	.22	.26
No. of observations	98	98	107

To the extent that these regression results are reliable, they give a big boost to the augmented Solow model. In particular, the estimate of Adjusted R^2 suggests that over (or close to) three quarters of income per capita differences across countries can be explained by differences in their physical and human capital investment behavior. The immediate implication is that technology (TFP) differences have a somewhat limited role, confined to at most accounting for about a quarter of the cross-country income per capita differences. If this conclusion were appropriate, it would imply that, as far as the proximate causes of prosperity are concerned, we could confine our attention to physical and human capital, and also assume that countries have access to more or less the same world technology. The implications for the modeling of economic growth are of course quite major.

The next subsection, however, will question the conclusion that technology differences are small and that physical and human capital differences are the major proximate cause of income per capita differences.

3.4.3. Challenges to the Regression Analyses of Growth Models. There are two major (and related) problems with this approach.

The first relates to the assumption that technology differences across countries are orthogonal to all other variables. While the constant technology advances assumption may be defended, the orthogonality assumption is too strong, almost untenable. When $\bar{A}_j$ varies across countries, it should also be correlated with measures of s_j^h and s_j^k; countries that are more productive also invest more in physical and human capital. This is for two reasons. The first is a version of the *omitted variable bias* problem; technology differences are also outcomes of investment decisions. Thus societies with high levels of $\bar{A}_j$ will be those that

have invested more in technology for various reasons; it is then natural to expect the same reasons to induce greater investment in physical and human capital as well. Second, even ignoring the omitted variable bias problem, there is a *reverse causality* problem; complementarity between technology and physical or human capital imply that countries with high $\bar{A}_j$ will find it more beneficial to increase their stock of human and physical capital.

In terms of the regression eq. (3.24), this implies that the key right-hand side variables are correlated with the error term, ε_j. Consequently, ordinary least squares regressions of eq. (3.24) will lead to upwardly biased estimates of α and β. In addition, the estimate of the R^2, which is a measure of how much of the cross-country variability in income per capita can be explained by physical and human capital, will also be biased upwards.

The second problem relates to the magnitudes of the estimates of α and β in eq. (3.24). The regression framework above is attractive in part because we can gauge whether the estimate of α was plausible. We should do the same for the estimate of β. However, such an exercise reveals that the coefficient on the investment rate in human capital, s_j^h, appears too large relative to microeconometric evidence.

Recall first that Mankiw, Romer and Weil use the fraction of the working age population enrolled in school. This variable ranges from 0.4% to over 12% in the sample of countries used for this regression. Their estimates therefore imply that, holding all other variables constant, a country with approximately 12% school enrollment should have income per capita about 9 times that of a country with $s_j^h = 0.4$. More explicitly, the predicted log difference in incomes between these two countries is

$$\frac{\alpha}{1-\alpha-\beta}\left(\ln 12 - \ln\left(0.4\right)\right) = 0.66 \times \left(\ln 12 - \ln\left(0.4\right)\right) \approx 2.24.$$

This implies that, holding all other factors constant, a country with school enrollment of over 12% should be about $\exp\left(2.24\right) - 1 \approx 8.5$ times richer than a country with a level of schooling investment of around 0.4.

In practice, the difference in average years of schooling between any two countries in the Mankiw-Romer-Weil sample is less than 12. Chapter 10 will show that there are good economic reasons to expect additional years of schooling to increase earnings proportionally, for example as in Mincer regressions of the form:

$$\ln w_i = \mathbf{X}_i^T \boldsymbol{\gamma} + \phi S_i, \tag{3.25}$$

where w_i denotes the wage earnings of individual i, $\mathbf{X}_i$ is a set of demographic controls, and S_i is years of schooling. The estimate of the coefficient ϕ is the rate of returns to education, measuring the proportional increase in earnings resulting from one more year of schooling. The microeconometrics literature suggests that eq. (3.25) provides a good approximation to the data and estimates ϕ to be between 0.06 and 0.10, implying that a worker with one more year of schooling earns about 6 to 10 percent more than a comparable worker with one less year of schooling. If labor markets are competitive, or at the very least, if wages are, on average, proportional to productivity, this also implies that one more year of schooling increases worker productivity by about 6 to 10 percent.

Can we deduce from this information how much richer a country with 12 more years of average schooling should be? The answer is yes, but with two caveats. First, we need to assume that the micro-level relationship as captured by (3.25) applies identically to all countries. In other words, the implicit assumption in wage regressions in general, and in eq. (3.25) in particular, is that the human capital (and the earnings capacity) of each individual is a function of his or her years of schooling. For example, ignoring other potential determinants, the wage earnings of individual i is a function of his or her schooling and can be written as $w_i = \tilde{\phi}(S_i)$. The first key assumption is that this $\tilde{\phi}$ function is identical across countries and can be approximated by an exponential function of the form $\tilde{\phi}(S_i) \approx \exp(\phi S_i)$ so that we obtain eq. (3.25). The reasons why this may be a reasonable assumption will be further discussed in Chapter 10.

Second, we need to assume that there are no *human capital externalities*—meaning that the human capital of a worker does not directly increase the productivity of other workers. There are reasons for why human capital externalities may exist and some economists believe that they are important. This evidence discussed in Chapter 10, however, suggests that human capital externalities—except those working through innovation—are unlikely to be very large. Thus it is reasonable to start without them. The key result which will enable us to go from the microeconometric wage regressions to cross-country differences is that, with constant returns to scale, perfectly competitive markets and no human capital externalities, differences in worker productivity directly translate into differences in income per capita. To see this, suppose that each firm f in country j has access to the production function

$$y_{fj} = K_f^{\alpha} (A_j H_f)^{1-\alpha},$$

where A_j is the productivity of all the firms in the country, K_f is the capital stock and H_f denotes the efficiency units of human capital employed by firm f (thus here the production function takes the more usual form where human capital is embedded in workers rather than the form in (3.14)). Here the Cobb-Douglas production function is chosen for simplicity and does not affect the argument. Suppose also that firms in this country face a cost of capital equal to R_j. With perfectly competitive factor markets, profit maximization implies that the cost of capital must equal its marginal product,

$$R_j = \alpha \left(\frac{K_f}{A_j H_f} \right)^{-(1-\alpha)}. \tag{3.26}$$

This implies that all firms ought to function at the same physical to human capital ratio, and consequently, all workers, regardless of their level of schooling, ought to work at the same physical to human capital ratio. Another direct implication of competitive labor markets is that in country j, wages per unit of human capital will be equal to

$$w_j = (1-\alpha)\, \alpha^{\alpha/(1-\alpha)} A_j R_j^{-\alpha/(1-\alpha)}.$$

Consequently, a worker with human capital h_i will receive a wage income of $w_j h_i$. Once again, this is a more general result; with aggregate constant returns to scale production technology, wage earnings are linear in the effective human capital of the worker, so that a

worker with twice as much effective human capital as another should earn twice as much as this other worker (see Exercise 3.9). Next, substituting for capital from (3.26), we have total income in country j as $Y_j = (1-\alpha)\,\alpha^{\alpha/(1-\alpha)} R_j^{-\alpha/(1-\alpha)} A_j H_j$, where H_j is the total efficiency units of labor in country j. This equation implies that ceteris paribus (in particular, holding constant capital intensity corresponding to R_j and technology, A_j), a doubling of human capital will translate into a doubling of total income. Notice that both A_j and R_j are being kept constant here. While it may be reasonable to keep technology, A_j, constant, one may wonder whether R_j is likely change systematically in response to changes in H_j. Even though this is a possibility, such changes are likely to be small in practice. First, international capital flows will reduce the sensitivity of the rate of return to capital to domestic labor supply (see Chapter 19). Second, recall that Theorem 2.6 established the constancy of the capital-labor ratio as a requirement for a balanced growth path, and R_j will indeed be constant when the capital-labor ratio is constant (see Exercise 3.10). This implies that under constant returns and perfectly competitive factor markets, a doubling of human capital (a doubling of the efficiency units of labor) has the same effects on the earnings of an individual as the effect of a doubling of aggregate human capital has on total output.

This analysis implies that the estimated Mincerian rates of return to schooling can be used to calculate differences in the stock of human capital across countries. So in the absence of *human capital externalities*, a country with 12 more years of average schooling should have a stock of human capital somewhere between $\exp(0.10 \times 12) \simeq 3.3$ and $\exp(0.06 \times 12) \simeq 2.05$ times the stock of human capital of a county with fewer years of schooling. This implies that, holding other factors constant, this country should be about 2-3 times as rich as the country with zero years of average schooling, which is much less than the 8.5 fold difference implied by the Mankiw-Romer-Weil analysis.

This discussion suggests that the estimate for β that is implied by the Mankiw-Romer-Weil regressions is too high relative to the estimates that would be implied by the microeconometric evidence and thus is likely to be upwardly biased. The cause of this overestimation is, in turn, most likely related to the possible correlation between the error term ε_j and the key right-hand side regressors in eq. (3.24). Consequently, regression analyses based on equation (3.24) appear unlikely to provide us with an accurate picture of the extent of cross-country productivity differences and of the proximate causes of income differences.

3.5. Calibrating Productivity Differences

What other approach can we use to gauge the importance of physical and human capital and technology differences? An alternative approach is to "calibrate" the (total factor) productivity differences across countries rather than estimating them using a regression framework. These total factor productivity differences are then interpreted as a measure of the contribution of "technology" to cross-country income differences.

The calibration approach was proposed and used by Klenow and Rodriguez (1997) and by Hall and Jones (1999). Here I follow Hall and Jones's approach, which is slightly simpler.

The advantage of the calibration approach is that the omitted variable bias underlying the estimates of Mankiw, Romer and Weil will be less important (since micro-level evidence will be used to anchor the contribution of human capital to economic growth). The disadvantage is that certain assumptions on functional forms have to be taken much more seriously and we explicitly have to assume no human capital externalities.

3.5.1. Basics. Suppose that each country has access to the Cobb-Douglas aggregate production function:

$$Y_j = K_j^{\alpha} \left(A_j H_j\right)^{1-\alpha}, \tag{3.27}$$

where H_j is the stock of human capital of country j, capturing the amount of efficiency units of labor available to this country. K_j is its stock of physical capital and A_j is labor-augmenting technology. Since our focus is on cross-country comparisons, time arguments are omitted.

Suppose that each worker in country j has S_j years of schooling. Then, using the Mincer eq. (3.25) from the previous section, ignoring the other covariates and taking exponents, H_j can be estimated as

$$H_j = \exp\left(\phi S_j\right) L_j,$$

where L_j is employment in country j and ϕ is the rate on returns to schooling estimated from eq. (3.25). This approach may not lead to very good estimates of the stock of human capital of a country, however. First, it does not take into account differences in other "human capital" factors, such as training or experience (which will be discussed in greater detail in Chapter 10). Second, countries may differ not only in the years of schooling of their labor forces, but in the quality of schooling and the amount of post-schooling human capital. Third, the rate of return to schooling may vary systematically across countries (for example, it may be lower in countries with a greater abundance of human capital). It is possible to deal with each of these problems to some extent by constructing better estimates of the stocks of human capital.

Following Hall and Jones, let us make only a partial correction for the last factor. Let us assume that the rate of return to schooling does not vary by country, but is potentially different for different years of schooling. For example, one year of primary schooling may be more valuable than one year of graduate school (for example, because learning how to read might increase productivity more than a solid understanding of growth theory). In particular, let the rate of return to acquiring the Sth year of schooling be $\phi\left(S\right)$. The above equation would be the special case where $\phi\left(S\right) = \phi$ for all S. With this assumption and with estimates of the returns to schooling for different years (e.g., primary schooling, secondaary schooling, and so on), a somewhat better estimate of the stock of human capital can be constructed as

$$H_j = \sum_{S} \exp\left\{\phi\left(S\right) S\right\} L_j\left(S\right)$$

where $L_j\left(S\right)$ now refers to the total employment of workers with S years of schooling in country j.

A series for $K_j(t)$ can be constructed from the Summers-Heston dataset using investment data and the *perpetual inventory* method. In particular, recall that, with exponential depreciation, the stock of physical capital evolves according to

$$K_j(t+1) = (1-\delta) K_j(t) + I_j(t),$$

where $I_j(t)$ is the level of investment in country j at time t. The perpetual inventory method involves using information on the depreciation rate, δ, and investments, $I_j(t)$'s, to estimate $K_j(t)$. Let us assume, following Hall and Jones that $\delta = 0.06$. With a complete series for $I_j(t)$, this equation can be used to calculate the stock of physical capital at any point in time. However, the Summers-Heston dataset does not contain investment information before the 1960s. This equation can still be used by assuming that each country's investment was growing at the same rate before the sample in order to compute the initial capital stock. Using this assumption, Hall and Jones calculate the physical capital stock for each country in the year 1985. I do the same here for various years. Finally, with the same arguments as before, I choose a value of 1/3 for α.

Given series for H_j and K_j and a value for α, we can construct "predicted" incomes at a point in time using the following equation

$$\hat{Y}_j = K_j^{1/3} (A_{US} H_j)^{2/3}$$

for each country j, where A_{US} is the labor-augmenting technology level of the United States, computed so that this equation fits the United States perfectly—$Y_{US} = K_{US}^{1/3} (A_{US} H_{US})^{2/3}$. Throughout, time indices are dropped.

Once a series for $\hat{Y}_j$ has been constructed, it can be compared to the actual output series. The gap between the two series represents the contribution of technology. Alternatively, we could explicitly back out country-specific technology terms (relative to the United States) as

$$\frac{A_j}{A_{US}} = \left(\frac{Y_j}{Y_{US}}\right)^{3/2} \left(\frac{K_{US}}{K_j}\right)^{1/2} \left(\frac{H_{US}}{H_j}\right).$$

Figures 3.3 and 3.4 show the results of these exercises for 1980, 1990 and 2000.

The following features are noteworthy:

(1) Differences in physical and human capital still matter a lot; the predicted and actual incomes are highly correlated. Thus the regression analysis was not entirely misleading in emphasizing the importance of physical and human capital.

(2) However, differently from the regression analysis, this exercise shows that there are significant *technology (productivity) differences*. There are often large gaps between predicted and actual incomes, showing the importance of technology differences across countries. This can be most easily seen in Figure 3.3, where practically all observations are above the 45°, which implies that the neoclassical model is over predicting the income level of countries that are poorer than the United States.

(3) The same pattern is visible in Figure 3.4, which plots the estimates of the technology differences, A_j/A_{US}, against log GDP per capita in different years. These differences

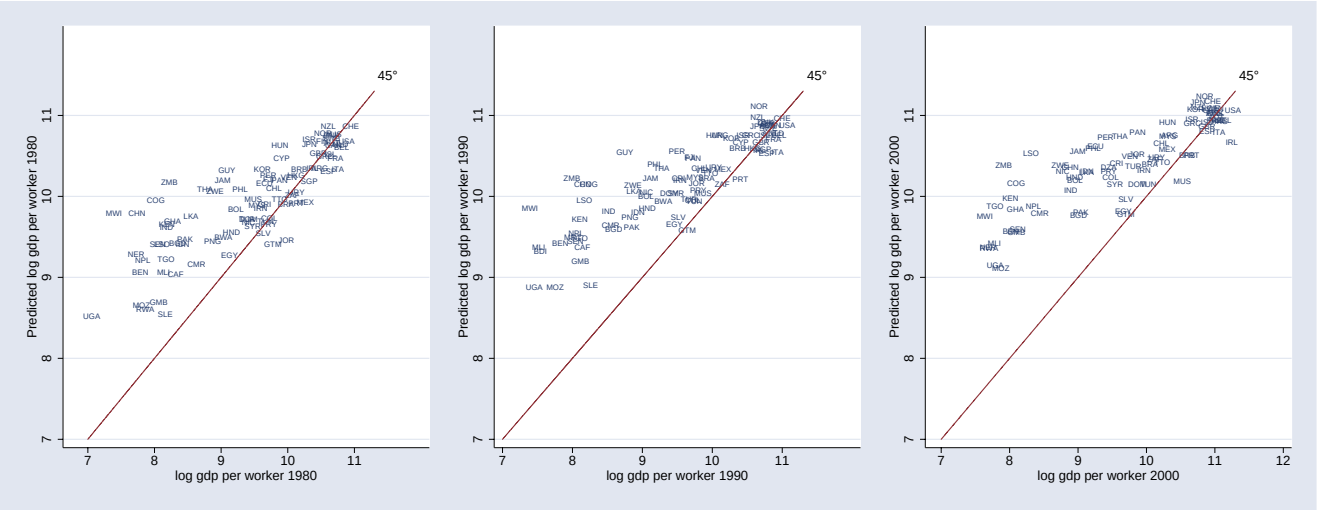

Figure 3.3. Calibrated technology levels relative to the US technology (from the Solow growth model with human capital) versus log GDP per worker, 1980, 1990 and 2000.

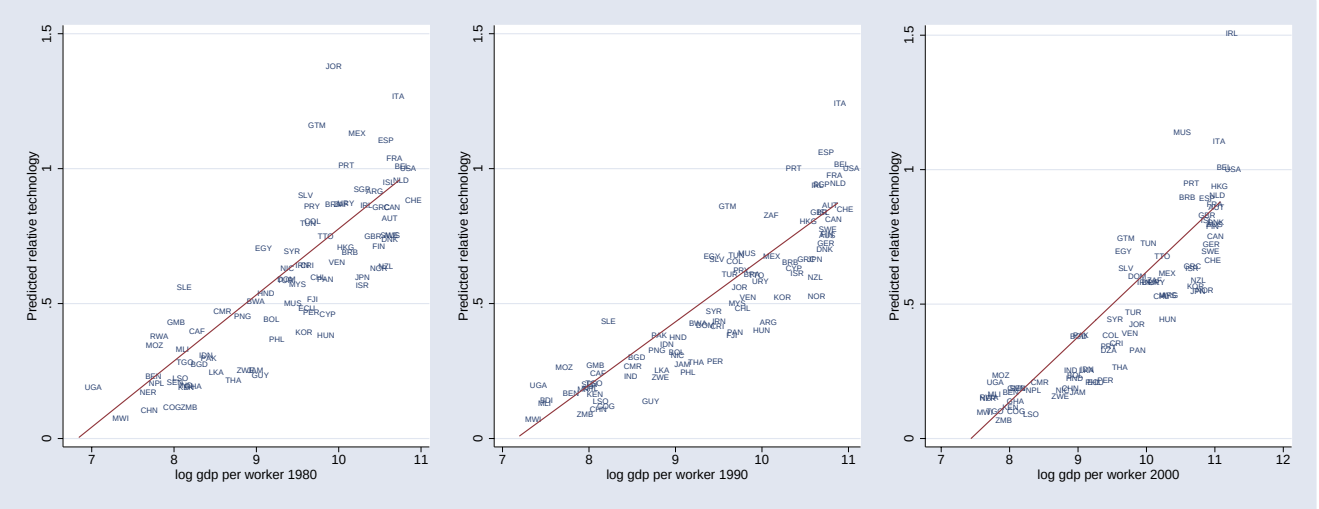

Figure 3.4. Calibrated technology levels relative to the US technology (from the Solow growth model with human capital) versus log GDP per worker, 1980, 1990 and 2000.

are often substantial. More important, these differences are also strongly correlated with income per capita; richer countries appear to have "better technologies".

(4) Also interesting is the pattern that the empirical fit of the neoclassical growth model seems to deteriorate over time. In the Figure 3.3, the observations are further above the 45° in 2000 than in 1980, and in Figure 3.4, the relative technology differences become larger over time. Why the fit of the simple neoclassical growth model is better in 1980 than in 2000 is an interesting and largely unanswered question.

3.5.2. Challenges. In the same way as the regression analysis was based on a number of stringent assumptions (in particular, the assumption that technology differences across countries were orthogonal to other factors), the calibration approach also relies on certain important assumptions. The above exposition highlighted several of those. In addition to the standard assumption that factor markets are competitive, the calibration exercise had to assume no human capital externalities, impose a Cobb-Douglas production function, and also make a range of approximations to measure cross-country differences in the stocks of physical and human capital.

Let us focus on the functional form assumptions. Could we get away without the Cobb-Douglas production function? The answer is yes, but not perfectly. The reason for this is that the exercise here is very similar to growth accounting, which does not need to make strong functional form assumptions (and this similarity to growth accounting is the reason why this exercise is sometimes referred to as "levels accounting"). In particular, recall eq. (3.5), which showed how TFP estimates can be obtained from a general constant returns to scale production function (under competitive labor markets) by using average factor shares. Now instead imagine that the production function that applies to all countries in the world is given by

$$F\left(K_j, H_j, A_j\right),$$

and countries differ according to their physical and human capital as well as technology—but not according to F. Suppose also that we have data on K_j and H_j as well as capital and labor share for each country. Then, a natural adaptation of eq. (3.5) can be used across countries rather than over time. In particular, let us a rank countries in descending order according to their physical capital to human capital ratios, K_j/H_j (use Exercise 3.1 to think about why this is the right way to rank countries rather than doing so randomly). Then,

$$\hat{x}_{j,j+1} = g_{j,j+1} - \bar{\alpha}_{K,j,j+1} g_{K,j,j+1} - \bar{\alpha}_{Lj,j+1} g_{H,j,j+1}, \tag{3.28}$$

where $g_{j,j+1}$ is the proportional difference in output between countries j and $j+1$, $g_{K,j,j+1}$ is the proportional difference in capital stock between these countries and $g_{H,j,j+1}$ is the proportional difference in human capital stocks. In addition, $\bar{\alpha}_{K,j,j+1}$ and $\bar{\alpha}_{Lj,j+1}$ are the average capital and labor shares between the two countries. The estimate $\hat{x}_{j,j+1}$ is then the proportional TFP difference between the two countries.

Using this method, and taking one of the countries, for example the United States, as the base country, we can calculate relative technology differences across countries. This levels accounting exercise faces two major challenges, however. One is data-related and the other one theoretical. First, data on capital and labor shares across countries are not available for most countries. This makes the use of eq. (3.28) far from straightforward. Consequently, almost all calibration or levels accounting exercises that estimate technology (productivity) differences use the Cobb-Douglas approach of the previous subsection (a constant value of α_K equal to 1/3).

Second, even if data on capital and labor shares were available, the differences in factor proportions, e.g., differences in K_j/H_j, across countries are large. An equation like (3.28) is a good approximation for small (infinitesimal) changes. As illustrated in Exercise 3.1, when differences in factor proportions are significant between the two observations, the use of this type of equation can lead to significant biases.

To sum up, the approach of calibrating productivity differences across countries is a useful alternative to the regression analysis, but has to rely on a range of stringent assumptions on the form of the production function and can also lead to biased estimates of technology differences when factors are mismeasured.

3.6. Estimating Productivity Differences

In the previous section, productivity/technology differences are obtained as "residuals" from a calibration exercise, so we have to trust the functional form assumptions used in this strategy. But if we are willing to trust the functional forms, we can also estimate these differences econometrically rather than rely on calibration. The great advantage of econometrics relative to calibration is that not only do we obtain estimates of the objects of interest, but we also have standard errors, which show us how much we can trust these estimates. In this section, I briefly discuss two different approaches to estimating productivity differences.

3.6.1. A Naïve Approach. The first possibility is to take a production function of the form (3.27) as given and try to estimate this using cross country data. In particular, taking logs:

$$\log Y_j = \alpha \log K_j + (1-\alpha) \log H_j + \alpha \log A_j. \tag{3.29}$$

Given series for Y_j, K_j and H_j, this equation can be estimated with ordinary least squares with the restriction that the coefficients on $\log K_j$ and $\log H_j$ sum to one, and the residuals can be interpreted as estimates of technology differences. Unfortunately, this approach is not particularly attractive, since the potential correlation between $\log A_j$ and $\log K_j$ or $\log H_j$ implies that the estimates of α need not be unbiased even when constant returns to scale is imposed. Moreover, if we do not impose the assumption that these coefficients sum to one and test this restriction, it will be rejected. Thus, this regression approach runs into the same difficulties as the Mankiw, Romer and Weil approach discussed previously.

What this discussion highlights is that, even if we are willing to presume that we know the functional form of the aggregate production function, it is difficult to directly estimate productivity differences. So how can we improve over this naïve approach? The answer involves making more use of economic theory. Estimating an equation of the form (3.29) does not make use of the fact that we are looking at the equilibrium of an economic system. A more sophisticated approach would use more of the restrictions imposed by equilibrium behavior (and bring additional relevant data). I next illustrate this using a specific attempt

based on international trade. The reader who is not familiar with trade theory may want to skip this subsection.

3.6.2. Learning from International Trade*. Models of growth and international trade are studied in Chapter 19. Even without a detailed discussion of international trade theory, we can use data from international trade flows and some simple principles of international trade theory to obtain another way of estimating productivity differences across countries.

Let us follow an important paper by Trefler (1993), which uses an augmented version of the standard Heckscher-Ohlin approach to international trade. The Heckscher-Ohlin approach assumes that countries differ according to their factor proportions (e.g., some countries have much more physical capital relative to their labor supply than others). In a closed economy, this will lead to differences in relative factor costs and differences in the relative prices of products using these factors in different intensities. International trade results as a way of taking advantage of these relative price differences. The most extreme form of the theory assumes no costs of shipping goods and no policy impediments to trade, so that international trade can happen costlessly between countries.

Trefler starts from the standard Heckscher-Ohlin model of international trade, but allows for factor-specific productivity differences, so that capital in country j has productivity A_j^k, thus a stock of capital K_j in this country is equivalent to an effective supply of capital $A_j^k K_j$. Similarly for labor (human capital), country j has productivity A_j^h. In addition, Trefler assumes that all countries have the same homothetic preferences and there are sufficient factor intensity differences across goods to ensure international trade between countries to arbitrage relative price and relative factor costs differences (or in the jargon of international trade, countries are said to be in the *cone of diversification*). This latter assumption is important: when all countries have the same productivities both in physical and human capital, it leads to the celebrated *factor price equalization* result—all factor prices would be equal in all countries, because the world economy is sufficiently integrated. When there are productivity differences across countries, this assumption instead leads to *conditional factor price equalization*, meaning that factor prices are equalized once their different "effective" productivities are taken into consideration.

Under these assumptions, a standard equation in international trade links the *net factor exports* of each country to the abundance of that factor in the country relative to the world as a whole. The term "net factor exports" needs some explanation. It does not refer to actual trade in factors (such as migration of people or capital flows). Instead trading goods is a way of trading the factors that are embodied in that particular good. For example, a country that exports cars made with capital and imports corn made with labor is implicitly exporting capital and importing labor. More specifically, the net export of capital by country j, $X_j^K$is calculated by looking at the total exports of country j and computing how much capital is necessary to produce these and then subtracting the amount of capital necessary to produce its total imports. For our purposes here, how this is calculated is not important (it suffices

to say that as with all things empirical, the devil is in the detail and these calculations are far from straightforward and require a range of assumptions). Then, the absence of trading frictions across countries and identical homothetic preferences imply that

$$X_j^K = A_j^k K_j - \gamma_j^s \sum_{i=1}^{N} A_i^k K_i \tag{3.30}$$

$$X_j^H = A_j^h H_j - \gamma_j^s \sum_{i=1}^{N} A_i^h H_i$$

where γ_j^s is the share of country j in world consumption (the value of this country's consumption divided by world consumption) and N is the total number of countries in the world. These equations simply restate the conclusion in the previous paragraph that a country will be a net exporter of capital if its effective supply of capital, $A_j^k K_j$, exceeds a fraction, here γ_j^s, of the world's effective supply of capital, $\sum_{i=1}^{N} A_i^k K_i$.

Consumption shares are easy to calculate. Then, given estimates for X_j^K and X_j^H, the above system of $2 \times N$ equations can be solved for the same number of unknowns, the A_i^k and A_i^h's for N countries. This gives estimates for factor-specific productivity differences across countries, generated from an entirely different source of variation than those exploited before. In fact, this exercise provides us with not a single productivity parameter, but a separate labor-augmenting (or human-capital-augmenting) and a capital-augmenting productivity term for each country.

How do we know that these numbers provide a good approximation to cross-country factor productivity differences? This is in some sense the same problem as we had in judging whether the calibrated productivity (technology) differences in the previous section were reliable. Fortunately, international trade theory gives us one more set of equations to check whether these numbers are reliable. As noted above, under the assumption that the world economy is sufficiently integrated, there is conditional factor price equalization. This implies that for any two countries j and j':

$$\frac{R_j}{A_j^k} = \frac{R_{j'}}{A_{j'}^k}, \tag{3.31}$$

$$\frac{w_j}{A_j^h} = \frac{w_{j'}}{A_{j'}^h}, \tag{3.32}$$

where R_j is the rental rate of capital in country j and w_j is the observed wage rate (which includes the compensation to human capital) in country j. Equation (3.32), for example, states that if workers in a particular country have, on average, half the efficiency units as those in the United States, their earnings should be roughly half of American workers.

With data on factor prices, we can therefore construct an alternative series for A_j^k and A_j^h's. It turns out that the series for A_j^k and A_j^h's implied by (3.30), (3.31) and (3.32) are very similar, so there appears to be some validity to this approach. This validation gives us some confidence that there is some information in the numbers that Trefler obtains.

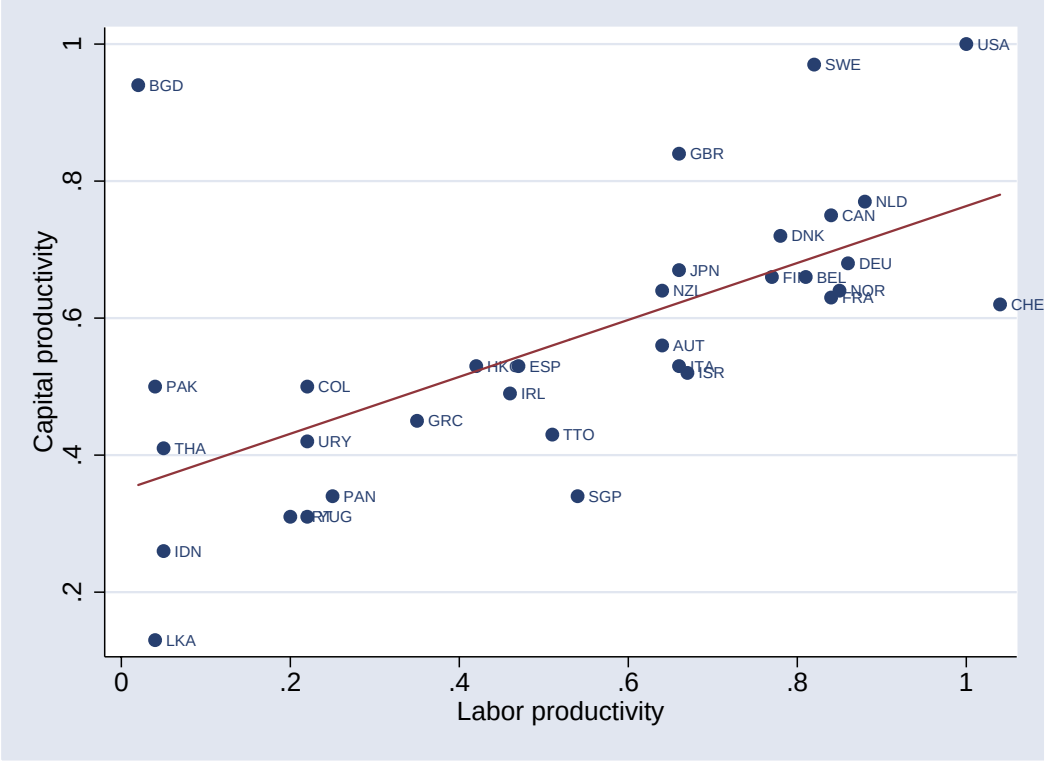

Figure 3.5. Comparison of labor-productivity and capital-productivity differences across countries.

Figure 3.5 shows Trefler's original estimates. The numbers in this figure imply that there are very large differences in labor productivity, and some substantial, but much smaller differences in capital productivity. For example, labor in Pakistan is 1/25th as productive as labor in the United States. In contrast, capital productivity differences are much more limited than labor productivity differences; capital in Pakistan is only half as productive as capital in the United States. This finding is not only intriguing in itself, but is also quite consistent with models of directed technological change in Chapter 15 that may explain why technological change is labor-augmenting in the long run.

It is also informative to compare the productivity difference estimates from Trefler's approach to those from the previous section. Figures 3.6 and 3.7 undertake this comparison. The first plots the labor-productivity difference estimates from the Trefler approach against the calibrated overall productivity differences from the Cobb-Douglas specification in the previous section. The similarity between the two series is remarkable. This gives us some confidence that both approaches are capturing some features of reality and that in fact there are significant productivity (technology) differences across countries. Interestingly, however, Figure 3.7 shows that the relationship between the calibrated productivity differences and the capital-productivity differences is considerably weaker than for labor productivity.

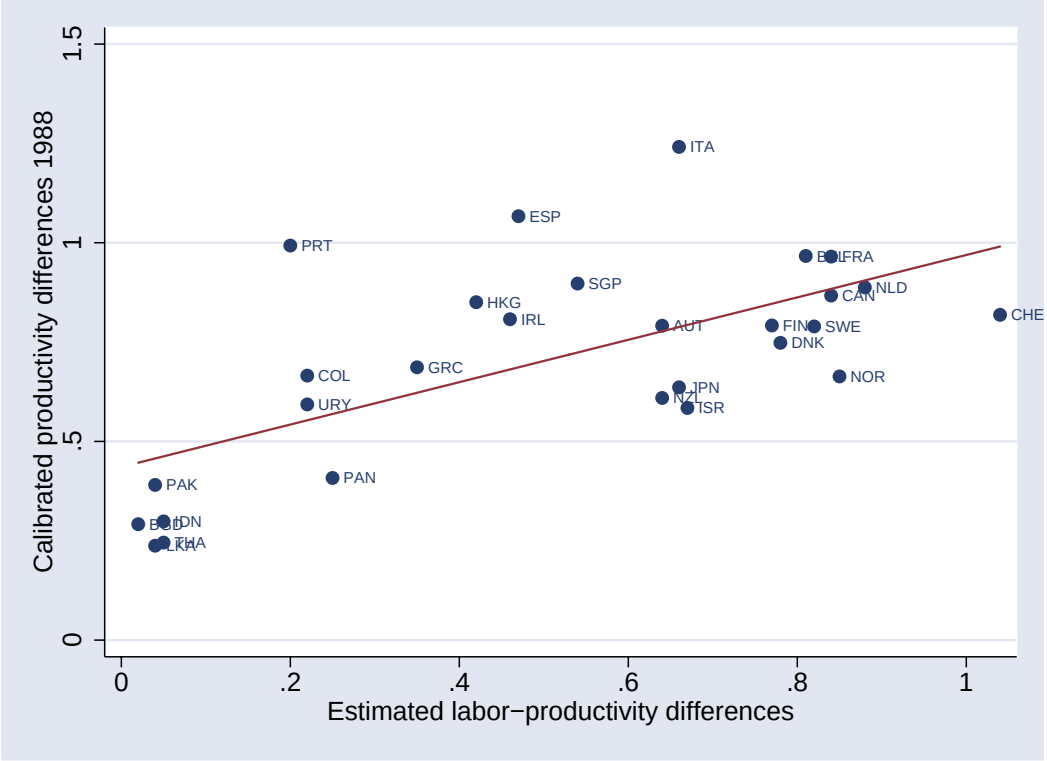

Figure 3.6. Comparison of the labor productivity estimates from the Trefler approach with the calibrated productivity differences from the Hall-Jones approach.

Despite its apparent success, it is important to emphasize that Trefler's approach relies on very stringent assumptions. To recap, the three major assumptions are:

(1) No international trading costs;
(2) Identical homothetic preferences;
(3) Sufficiently integrated world economy, leading to conditional factor price equalization.

All three of these assumptions are rejected in the data in one form or another. There are clearly international trading costs, including freight costs, tariff costs and other trading restrictions. There is very well-documented home bias in consumption violating the identical homothetic preferences assumption. Finally, most trade economists believe that conditional factor price equalization is not a good description of factor price differences across countries. In view of all of these, the results from the Trefler exercise have to be interpreted with caution. Nevertheless, this approach is important in showing how different sources of data and additional theory can be used to estimate cross-country technology differences and in providing a cross-validation for the calibration and estimation results discussed previously.

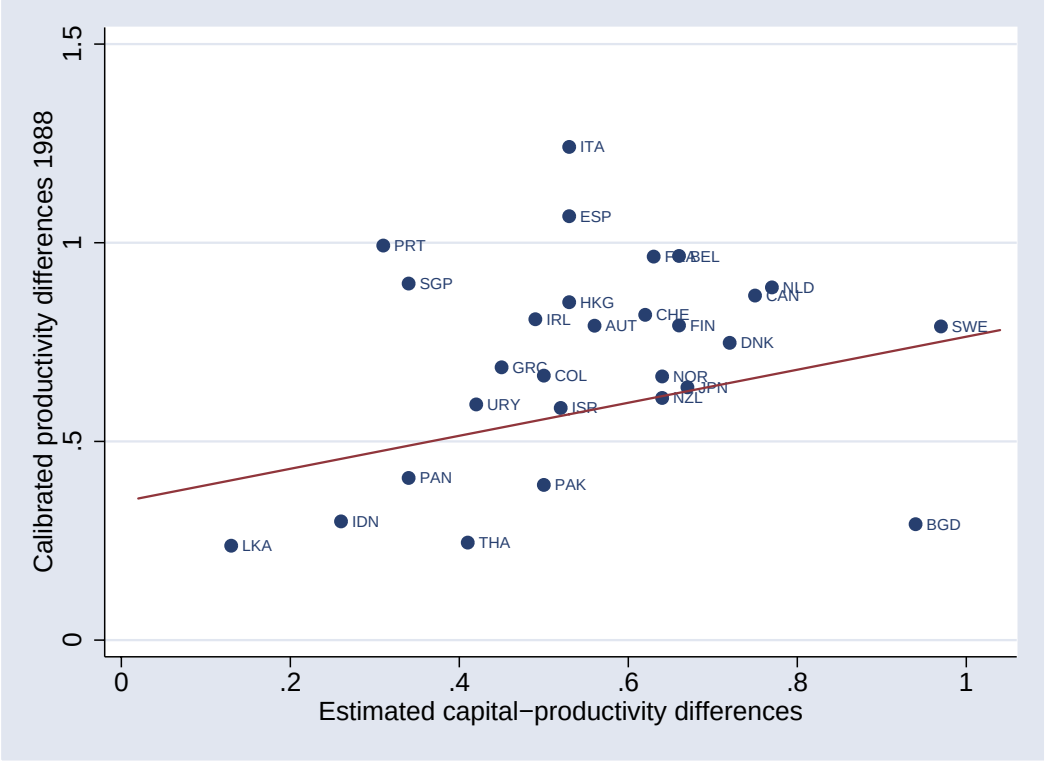

FIGURE 3.7. Comparison of the capital productivity estimates from the Trefler approach with the calibrated productivity differences from the Hall-Jones approach.

3.7. Taking Stock

What have we learned? The major point of this chapter has not been the development of new theory. Our major objective in this chapter has been to see whether we could use the Solow model to have a more informed interpretation of cross-country differences and also use data in order to understand the strengths and shortcomings of the Solow growth model.

At the end of this brief journey, the message is somewhat mixed. On the positive side, despite its simplicity, the Solow model has enough substance that we can take it to data in various different forms, including TFP accounting, regression analysis and calibration. Moreover, each of these different methods gives us some idea about the sources of economic growth over time and of income differences across countries.

On the negative side, however, no single approach is entirely convincing. Each relies on a range of stringent auxiliary assumptions. Consequently, no firm conclusions can be drawn. The simplest applications of the Solow accounting framework suggest that technology is the main source of economic growth over time. However, this conclusion is disputed by those who point out that sufficient adjustments to the quality of physical and human capital substantially reduce or perhaps even totally eliminate residual TFP growth. The same debate is seen in the context of cross-country income differences; while some believe that accounting

for differences in physical and human capital across countries leaves little need for technology differences, others show that, with reasonable models, most of the cross-country differences are due to technology.

While complete agreement is not possible, it is safe to say that the consensus in the literature today favors the interpretation that cross-country differences in income per capita cannot be understood solely on the basis of differences in physical and human capital; in other words there are technology differences across countries and these technology differences are likely to be at the heart of cross-country income and growth differences.

Hence one important potential lesson from this data detour is that technological progress is not only important in generating economic growth in the basic Solow model, but also likely to be a major factor in cross-country differences in prosperity. A detailed study of technological progress and technology adoption decisions of households and firms is therefore necessary. This motivates the detailed analysis of technological progress and technology adoption later in the book. It is also useful to emphasize once again that differences in TFP are not necessarily due to technology in the narrow sense. If two countries have access to the same technology but make use of these techniques in different ways with different degrees of efficiency or if they are subject to different degrees of market or organizational failures, these differences will show up as TFP differences. One indication that TFP differences arising from market or organizational failures are important comes from episodes of severe crises. When countries have large drops in their income, due to civil wars, political instability, financial crises or other reasons, this is almost always associated with a corresponding decline in TFP (and little change in the capital stock and a much smaller change in labor inputs). Naturally, these drops in TFP are not caused by "technological regress," but result from the breakdown of the market or increases in some other sources of inefficiency. Therefore, technology differences should always be construed rather broadly and we should pay special attention to cross-country differences in the efficiency of production. By implication, to understand TFP differences across countries, we must study not only differences in the techniques that they use but the way they organize markets and firms and how to incentivize different agents in the economy. This again shapes our agenda for the rest of the book, especially paving the way for our investigation of endogenous technological change in Part 4 and of differences in technology and productive efficiency across countries in Parts 6 and 7.

There is one more sense in which what we have learned in this chapter is limited. What the Solow model makes us focus on, physical capital, human capital and technology, are proximate causes of economic growth in cross-country differences. It is important to know which of these proximate causes are important and how they affect economic performance both to have a better understanding of the mechanics of economic growth and also to know which classes of models to focus on. But at some level (and exaggerating somewhat) to say that a country is poor because it has insufficient physical capital, human capital and inefficient technology is like saying that a person is poor because he does not have money. There are, in turn, other reasons making some countries more abundant in physical capital,

human capital and technology, in the same way as there are factors that make a person have more money than another. I have referred to these as the *fundamental causes* of differences in prosperity, contrasting them with the proximate causes. A satisfactory understanding of economic growth and of differences in prosperity across countries requires both an analysis of proximate causes and of fundamental causes of economic growth. The former is essential for the study of the mechanics of economic growth and to develop the appropriate formal models incorporating these insights. The latter is important for understanding why some societies make choices that lead them to low physical capital, low human capital and inefficient technology and thus to relative poverty. This is the issue I turn to in the next chapter.

3.8. References and Literature

The growth accounting framework is introduced and applied in Solow (1957). Jorgensen, Gollop and Fraumeni (1987) give a comprehensive development of this framework, emphasizing how competitive markets are necessary and essentially sufficient for this approach to work. They also highlight the measurement difficulties and emphasize how underestimates of the quality improvements in physical and human capital will lead to overestimates of the contribution of technology to economic growth. Jorgensen (2005) contains a more recent survey.

Regression analysis based on the Solow model has a long history. More recent contributions include Baumol (1986), Barro (1991) and Barro and Sala-i-Martin (1992). Barro (1991) has done more than anybody else to popularize growth regressions, which have become a very commonly-used technique over the past two decades. See Durlauf (1996), Durlauf, Johnson and Temple (2005) and Quah (1993) for various critiques of growth regressions, especially focusing on issues of convergence. Wooldridge (2002) contains an excellent discussion of issues of omitted variable bias and how different approaches can be used (see, for example, Chapters 4, 5, 8, 9 and 10). The difficulties involved in estimating models with fixed effects and lagged dependent variables are discussed in Chapter 11. You should read more about the economic limitations of growth regressions and the econometric problems facing such regressions before embarking upon your own empirical analyses.

The augmented Solow model with human capital is a generalization of the model presented in Mankiw, Romer and Weil (1992). As noted in the text, treating human capital as a separate factor of production this somewhat unusual and difficult to micro-found. Different ways of introducing human capital in the basic growth model are discussed in Chapter 10 below.

Mankiw, Romer and Weil (1992) also provide the first regression estimates of the Solow and the augmented Solow models. A detailed critique of the Mankiw, Romer and Weil is provided in Klenow and Rodriguez (1997). Hall and Jones (1999) and Klenow and Rodriguez (1997) provide the first calibrated estimates of productivity (technology) differences across countries. Caselli (2005) gives an excellent overview of this literature, with a detailed discussion of how one might correct for differences in the quality of physical and human capital

across countries. He reaches the conclusion that such corrections will not change the basic conclusions of Klenow and Rodriguez and Hall and Jones, that cross-country technology differences are important.

The last subsection draws on Trefler (1993). Trefler does not emphasize the productivity estimates implied by this approach, focusing more on this method as a way of testing the Heckscher-Ohlin model. Nevertheless, these productivity estimates are an important input for growth economists. Trefler's approach has been criticized for various reasons, which are secondary for our focus here. The interested reader should look at Gabaix (2000) and Davis and Weinstein (2001).

3.9. Exercises

Exercise 3.1. Suppose that output is given by the neoclassical production function $Y(t) = F[K(t), L(t), A(t)]$ satisfying Assumptions 1 and 2, and that we observe output, capital and labor at two dates t and $t+T$. Suppose that we estimate TFP growth between these two dates using the equation

$$\hat{x}(t,t+T) = g(t,t+T) - \alpha_K(t)\, g_K(t,t+T) - \alpha_L(t)\, g_L(t,t+T),$$

where $g(t,t+T)$ denotes output growth between dates t and $t+T$, and other growth terms are defined similarly, while $\alpha_K(t)$ and $\alpha_L(t)$ denote the factor shares at the beginning date. Let $x(t,t+T)$ be the true TFP growth between these two dates. Show that there exists functions F such that $\hat{x}(t,t+T)/x(t,t+T)$ can be arbitrarily large or small. Next show the same result when the TFP estimate is constructed using the end date factor shares:

$$\hat{x}(t,t+T) = g(t,t+T) - \alpha_K(t+T)\, g_K(t,t+T) - \alpha_L(t+T)\, g_L(t,t+T).$$

Explain the importance of differences in factor proportions (capital-labor ratio) between the beginning and end dates in these results.

Exercise 3.2. Consider the economy with labor market imperfections as in the second part of Exercise 2.15 from the previous chapter, where workers were paid a fraction $\beta > 0$ of output. Show that in this economy the fundamental growth accounting equation leads to biased estimates of TFP.

Exercise 3.3. For the Cobb-Douglas production function from Example 3.1 $Y(t) = A(t) K(t)^{a} L(t)^{1-\alpha}$, derive an exact analog of (3.9) and show how the rate of convergence, that is, the coefficient in front of $(\log y(t) - \log y^*(t))$, changes as a function of $\log y(t)$.

Exercise 3.4. Consider once again the production function in Example 3.1. Suppose that two countries, 1 and 2, have exactly the same technology and the same parameters α, n, δ and g, thus the same $y^*(t)$. Suppose that they start with $y_1(0) = 2y_2(0)$ at time $t = 0$. Using the parameter values in Example 3.1 calculate how long it would take for the income gap between the two countries to decline to 10%.

Exercise 3.5. Consider a collection of Solow economies, each with different levels of δ, s and n. Show that an equivalent of the conditional convergence regression eq. (3.11) can be derived from an analog of (3.9) in this case.

Exercise 3.6. Prove Proposition 3.2.

Exercise 3.7. In the augmented Solow model (cfr Proposition 3.2) determine the impact of increases in s_k, s_h and n on h^* and k^*.

Exercise 3.8. Consider a world economy consisting of countries represented by the augmented Solow growth model, with the production functions given by (3.14). Derive the equivalent of the fundamental growth accounting equation in this case and explain how one might use available data in order to estimate TFP growth using this equation.

Exercise 3.9. Consider the basic Solow model with no population growth and no technological progress, and a production function of the form $F(K, H)$, where H denotes the efficiency units of labor (human capital), given by $H = \sum_{i \in \mathcal{N}} h_i$, where $\mathcal{N}$ is the set of all individuals in the population and h_i is the human capital of individual i. Assume that H is fixed. Suppose there are no human capital externalities and factor markets are competitive.

(1) Calculate the steady-state equilibrium of this economy.
(2) Prove that if 10% higher h at the individual level is associated with a% higher earnings, then a 10% increase in the country's stock of human capital H will lead to a% increase in steady-state output. Compare this to the immediate impact of an unanticipated 10% increase in H (that is, consider the impact of a 10% increase in H with the stock of capital unchanged).

Exercise 3.10. Consider a constant returns to scale production function for country j, $Y_j = F(K_j, A_j H_j)$, where K_j is physical capital, H_j denotes the efficiency units of labor and A_j is labor-augmenting technology. Prove that if $K_j/Y_j = K_{j'}/Y_{j'}$ in two different countries j and j', then the rental rates of capital in the two countries, R_j and $R_{j'}$ will also be equal.

Exercise 3.11. Imagine you have a cross-section of countries, $i = 1, ..., N$, and for each country, at a single point in time, you observe labor L_i, capital K_i, total output Y_i, and the share of capital in national income, α_i^K. Assume that all countries have access to a production technology of the following form

$$F(K, L, A)$$

where A is technology. Assume that F exhibits constant returns to scale in K and L, and all markets are competitive.

(1) Explain how you would estimate relative differences in technology/productivity across countries due to the term A without making any further assumptions. Write down the equations that are involved in estimating the contribution of A to cross-country income differences explicitly.
(2) Suppose that the exercise in part 1 leads to large differences in productivity due to the A term. How would you interpret this? Does it imply that countries have access to different production possibility sets?
(3) Now suppose that the true production function is $F(K, H, A)$, where H denotes efficiency units of labor. What other types of data would you need in order to estimate the contribution of technology/productivity across countries to output differences.

(4) Show that if H is calculated as in Section 3.5, but there are significant quality-of-schooling differences and no differences in A, this strategy will lead to significant differences in the estimates of A.

CHAPTER 4

Fundamental Determinants of Differences in Economic Performance

4.1. Proximate Versus Fundamental Causes

> "...the factors we have listed (innovation, economies of scale, education, capital accumulation etc.) are not causes of growth; *they are growth.*" (North and Thomas, 1973, p. 2, italics in original).

The previous chapter illustrated how the Solow growth model can be used to understand cross-country income differences and the process of economic growth. In the context of the Solow growth model, the process of economic growth is driven by technological progress. Cross-country income differences, on the other hand, are due to a combination of technology differences, differences in physical capital per worker and in human capital per worker. While this approach provides us with a good starting point and delineates potential sources of economic growth and cross-country income differences, these sources are only *proximate causes* of economic growth and economic success. Let us focus on cross-country income differences, for example. As soon as we attempt to explain these differences with technology, physical capital and human capital differences, an obvious next question presents itself: if technology, physical capital and human capital are so important in understanding differences in the wealth of nations and if they can account for five-fold, ten-fold, twenty-fold or even thirty-fold differences in income per capita across countries, then why is it that societies do not improve their technologies, invest more in physical capital and accumulate more human capital?

It appears therefore that any explanation that simply relies on technology, physical capital and human capital differences across countries is, at some level, incomplete. There must be some other reasons underneath those, reasons which we will refer to as *fundamental causes* of economic growth. It is these reasons that are preventing many countries from investing enough in technology, physical capital and human capital.

An investigation of fundamental causes of economic growth is important for at least two reasons. First, any theory that focuses on the intervening variables (proximate causes) alone, without understanding what the underlying driving forces are, would be incomplete. Thus growth theory will not fulfill its full promise until it comes to grips with these fundamental causes. Second, if part of our study of economic growth is motivated by improving the growth performance of certain nations and the living standards of their citizens, understanding fundamental causes is central to this objective, since attempting to increase growth just

focusing on proximate causes would be tantamount to dealing with symptoms of diseases without understanding what the diseases themselves are. While such attacks on symptoms can sometimes be useful, they are no substitute for a fuller understanding of the causes of the disease, which may allow a more satisfactory treatment. In the same way, we may hope that an understanding of the fundamental causes of economic growth could one day offer more satisfactory solutions to the major questions of social science concerning why some countries are poor and some are rich and how we can ensure that more nations grow faster.

What could these fundamental causes be? Can we make progress in understanding them? And, perhaps most relevant for this book, is growth theory useful in such an endeavor?

In this chapter, I attempted develop some answers to these questions. Let us start with the last two questions. The argument in this book is that a good understanding of the mechanics of economic growth, thus the detailed models of the growth process, are essential for a successful investigation of the fundamental causes of economic growth. This is for at least two reasons; first, we can only pose useful questions about the fundamental causes of economic growth by understanding what the major proximate causes are and how they impact economic outcomes. Second, only models that provide a good approximation to reality and are successful in qualitatively and quantitatively matching the major features of the growth process can inform us about whether the potential fundamental causes that are proposed could indeed play a significant role in generating the huge income per capita differences across countries. Our analysis of the mechanics of economic growth will often enable us to discard or refine certain proposed fundamental causes. As to the question of whether we can make progress, the vast economic growth literature is evidence that progress is being made and more progress is certainly achievable. In some sense, it is part of the objective of this book to convince you that the answer to this question is yes.

Returning to the first question, there are innumerable fundamental causes of economic growth that various economists, historians and social scientists have proposed over the ages. Clearly, listing them and cataloging them will be neither informative nor useful. Instead, I will classify the major candidate fundamental causes of economic growth into four categories of hypotheses. While such a classification undoubtedly fails to do justice to some of the nuances of the previous literature, it is satisfactory for our purposes of bringing out the main factors affecting cross-country income differences and economic growth. These are:

(1) The luck hypothesis.
(2) The geography hypothesis.
(3) The culture hypothesis.
(4) The institutions hypothesis.

By *luck*, I refer to the set of fundamental causes that explain divergent paths of economic performance among otherwise-identical countries, either because some small uncertainty or heterogeneity between them have led to different choices with far-ranging consequences, or because of different selection among multiple equilibria. Multiple equilibria correspondent to different equilibrium configurations arising for the same underlying economic environment.

When our models exhibit multiple equilibria, we are often unable to make specific predictions as to which of these equilibria will be selected by different countries and it is possible for two otherwise-identical countries to end up in different equilibria with quite different implications for economic growth and living standards. Luck and multiple equilibria can manifest themselves through any of the proximate causes discussed so far (and through some additional mechanisms that will be discussed later in the book). For example, multiple equilibria can exist in technology adoption or in models that focus on human capital and physical capital investments. Therefore, explanations based on luck or multiple equilibria are often theoretically well grounded. Whether they are empirically plausible is another matter.

By *geography*, I refer to all factors that are imposed on individuals as part of the physical, geographic and ecological environment in which they live. Geography can affect economic growth through a variety of proximate causes. Geographic factors that can influence the growth process include soil quality, which can affect agricultural productivity; natural resources, which directly contribute to the wealth of a nation and may facilitate industrialization by providing certain key resources, such as coal and iron ore during critical times; climate, which may affect productivity and attitudes directly; topography, which can affect the costs of transportation and communication; and disease environment, which can affect individual health, productivity and incentives to accumulate physical and human capital. For example, in terms of the aggregate production function of the Solow model, poor soil quality, lack of natural resources or an inhospitable climate may correspond to a low level of A, that is, to a type of "inefficient technology". Many philosophers and social scientists have suggested that climate also affects preferences in a fundamental way, so perhaps individuals living in certain climates have a preference for earlier rather than later consumption, thus reducing their saving rates both in physical and human capital. Finally, differences in the disease burden across areas may affect the productivity of individuals and their willingness to accumulate human capital. Thus geography-based explanations can easily be incorporated into both the simple Solow model as well as into many of the more sophisticated models discussed later in the book.

By *culture*, I refer to beliefs, values and preferences that influence individual economic behavior. Differences in religious beliefs across societies are among the clearest examples of cultural differences that may affect economic behavior. Differences in preferences, for example, regarding how important wealth is relative to other status-generating activities and how patient individuals should be, might be as important as, or even more important than, luck, geography and institutions in affecting economic performance. Broadly speaking, culture can affect economic outcomes through two major channels. First, it can affect the willingness of individuals to tradeoff different activities or consumption today versus consumption tomorrow. Via this channel, culture will influence societies' occupational choices, market structure, saving rates and their willingness to accumulate physical and human capital. Second, culture may also affect the degree of cooperation among individuals, and cooperation and trust are often important foundations for productive activities in societies.

By *institutions*, I refer to rules, regulations, laws and policies that affect economic incentives and thus the incentives to invest in technology, physical capital and human capital. It is a truism of economic analysis that individuals will only take actions that are rewarded. Institutions, which shape these rewards, must therefore be important in affecting all three of the proximate causes of economic growth. What distinguishes institutions from geography, luck and culture is that they are *social choices*. While laws and regulations are not directly chosen by individuals and some institutional arrangements may be historically persistent, in the end the laws, policies and regulations under which a society lives are the choices of the members of that society. If the members of the society collectively decide to change them, they can change them. This implies that if institutions are a major fundamental cause of economic growth and cross-country differences in economic performance, they can be potentially reformed so as to achieve better outcomes. Such reforms may not be easy, they may encounter a lot of opposition, and often we may not exactly know which reforms will work. But they are still within the realm of the possible, and further research might help us understand how such reforms will affect economic incentives and how they can be implemented.

There is a clear connection between institutions and culture. Both affect individual behavior and both are important determinants of incentives. Nevertheless, a crucial difference between the theories put into these two categories justifies their separation. Institutions are directly under the control of the members of the society, in the sense that by changing the distribution of resources, constitutions, laws and policies, individuals can collectively influence the institutions under which they live. In contrast, culture refers to a set of beliefs that have evolved over time and outside the direct control of individuals.[1] Even though institutions might be hard to change in practice, culture is much harder to influence, and any advice to a society that it should change its culture is almost vacuous.

It is also important to emphasize that institutions themselves, even if they are a fundamental cause of economic growth and income differences across countries, are *endogenous*. They are equilibrium choices made either by the society at large or by some powerful groups in society. One can then argue that at some level luck, geography or culture should be more important, because they can be "more exogenous" in the sense that they are *not* equilibrium choices in the same way as institutions are and institutions will vary across societies largely because of geographic, cultural or random factors. While at some philosophical level this is correct, it is not a particularly useful observation. It neither obviates the need to understand the *direct* effects of luck, geography, culture and institutions (these direct effects have been the focus of much of the debate in this area) nor does it imply that understanding the specific role of institutions and economic development is secondary in any sense. After all, if we can understand what the effects of institutions are and which specific types of institutions matter, institutional reform can lead to major changes in economic behavior (even if part of the original variation in institutions was due to geography, luck or culture).

[1] A major and important exception to this is the effect of education on the beliefs and values of individuals.

In the rest of this chapter, I will explain what the reasoning for these different hypotheses are and provide a brief overview of the empirical evidence pertaining to various fundamental causes of economic growth. The theoretical underpinnings and implications of the institutions view will be further developed in Part 8 of the book. At this point, the reader should be warned that the author of this book is not an objective outside observer in this debate, but a strong proponent of the institutions hypothesis. Therefore, not surprisingly, this chapter will conclude that the institutional differences are at the root of the important proximate causes that we have already listed. Nevertheless, the same evidence can be interpreted in different ways and the reader should feel free to draw his or her own conclusions.

Before delving into a discussion of the fundamental causes, one other topic deserves a brief discussion. This is where I start in the next section.

4.2. Economies of Scale, Population, Technology and World Growth

As emphasized in Chapter 1, cross-country income differences result from the differential growth experiences of countries over the past two centuries. This makes it important for us to understand the process of economic growth. Equally remarkable is the fact that world economic growth is, by and large, a phenomenon of the past 200 years or so. Thus another major question concerns why economic growth started so recently and why there was little economic growth before. The growth literature has provided a variety of interesting answers to this question. Many of them focus on the role of economies of scale and population. The argument goes as follows: in the presence of economies of scale (or increasing returns to scale), population needs to have reached a certain critical level so that technological progress can gather speed. Alternatively, some natural (steady) progress of technology that may have been going on in the background needs to reach a critical threshold for the process of growth to begin. These stories are quite plausible. World population has indeed increased tremendously over the past one million years and the world's inhabitants today have access to a pool of knowledge and technology unimaginable to our ancestors. Could these long-run developments of the world economy also account for cross-country differences? Is the increase in world population a good explanation for the take off of the world economy?

Let us focus on population to give a preliminary answer to these questions. The simplest way of thinking of the relationship between population and technological change is the Simon-Kremer model (after the demographer Julian Simon and the economist Michael Kremer). This model is implicitly one of the entire world economy, since there are no cross-country differences and proponents of this model do not try to explain differences across countries by their populations. Imagine that there is a small probability that each individual will discover a new idea that will contribute to the knowledge pool of the society. Crucially, these random discoveries are independent across individuals, so that a larger pool of individuals implies the discovery of more new ideas, increasing aggregate productivity. Let output be determined simply by technology (this can be generalized so that technology and capital determine output

as in the Solow model, but this does not affect the point I would like to make here):

$$Y(t) = A(t) L(t)^{\alpha} Z^{1-\alpha},$$

where $\alpha \in (0,1)$, $Y(t)$ is world output, $A(t)$ is the world stock of technology, $L(t)$ is world population, and Z is some other fixed factor of production, for example, land. I normalize $Z = 1$ without loss of any generality. Time is continuous and ideas are discovered at the rate λ so that the knowledge pool of the society evolves according to the differential equation

$$\dot{A}(t) = \lambda L(t), \tag{4.1}$$

with $A(0) > 0$ is taken as given. Population, in turn, is a function of output, for example because of the Malthusian channels discussed in Chapter 21. For example, suppose that population increases linearly in output,

$$L(t) = \phi Y(t). \tag{4.2}$$

Combining these three eq.'s (see Exercise 4.1):

$$\dot{A}(t) = \lambda \phi^{\frac{1}{1-\alpha}} A(t). \tag{4.3}$$

The solution to this differential equation involves

$$A(t) = \exp\left(\lambda \phi^{1/(1-\alpha)} t\right) A(0). \tag{4.4}$$

This shows how a model of economies of scale (increasing returns) in population can generate a steady increase in technology. It is also straightforward to verify that

$$Y(t) = \phi^{\frac{\alpha}{1-\alpha}} A(t), \tag{4.5}$$

so that aggregate income also grows at the constant level $\lambda \phi^{1/(1-\alpha)}$. Such a model would generate steady growth but no acceleration. Simon and Kremer, instead, assume that there are stronger externalities to population than in (4.1). They impose the following equation governing the accumulation of ideas:

$$\frac{\dot{A}(t)}{A(t)} = \lambda L(t). \tag{4.6}$$

This implies that the law of motion of technology is given by (see Exercise 4.2):

$$A(t) = \frac{1}{A(0)^{-1} - \lambda \phi^{1/(1-\alpha)} t}. \tag{4.7}$$

In contrast to (4.4), this equation implies an accelerating output level. Starting from a low-level of $A(0)$ (or $L(0)$), this model would generate a long period of low output, followed by an acceleration or a *takeoff*, reminiscent to the modern economic growth experience discussed in Chapter 1. Therefore, a model with significant economies of scale is capable of generating the pattern of take off we see in the data.

While such a story, which has been proposed by many economists, may have some appeal for accounting for world growth, it is important to emphasize that it has little to say about cross-country income differences or why modern economic growth started in some countries (Western Europe) and not others (Asia, South America, Africa). In fact, if we take Western

Europe and Asia as the relevant economic units, European population has consistently been less than that of Asia over the past 2000 years, thus it is unlikely that simple economies of scale to population are responsible for the economic takeoff in Western Europe while Asia stagnated.

This discussion therefore suggests that models based on economies of scale of one sort or another do not provide us with fundamental causes of cross-country income differences. At best, they are theories of world growth (the world taken as a whole). Moreover, once we recognize that the modern economic growth process was uneven, meaning that it took place in some parts of the world and not others, the appeal of such theories diminishes further. If economies of scale were responsible for modern economic growth, it should also be able to explain when and where this process of economic growth started. Existing models based on economies of scale do not. In this sense, they are unlikely to provide the fundamental causes of modern economic growth. Does this mean that these types of economies of scale and increasing returns to population are unimportant? Certainly not. They may well be part of the proximate causes of the growth process (for example, the part lying in the blackbox of technology). But this discussion suggests that these models need to be augmented by other fundamental causes in order to explain why, when and where the takeoff occurred. This further motivates the investigation of the fundamental causes.

4.3. The Four Fundamental Causes

4.3.1. Luck and Multiple Equilibria. Chapter 21 presents a number of models in which multiple equilibria or multiple steady states can arise because of coordination failures in the product market or because of imperfections in credit markets. These models suggest that an economy, with given parameter values, can exhibit very different types of equilibrium behavior, some with higher levels of income or perhaps sustained growth, while others correspond to poverty and stagnation. To give a flavor of these models, consider the following simple game of investment played by a large number of agents in the society:

everybody else →	high investment	low investment
individual ↓		
high investment	y^H, y^H	$y^L - \varepsilon, y^L$
low investment	$y^L, y^L - \varepsilon$	y^L, y^L

We focus on symmetric equilibria. The first column indicates that all agents (except the individual in question) have chosen high investment, while the second corresponds to low investment I'll agents. The first row, on the other hand, corresponds to high investment by the individual in question, and the second row is for low investment. In each cell, the first number refers to the income of the individual in question, while the second number is the payoff to each of the other agents in the economy. Suppose that $y^H > y^L$ and $\varepsilon > 0$. This payoff matrix implies that high investment is more profitable when others are also undertaking

high investment. For example, this may be because of technological complementarities or aggregate demand spillovers (see Chapter 21).

It is then clear that there are two (pure-strategy) symmetric equilibria in this game. In one, the individual expects all other agents to choose high investment and he does so himself as well. In the other, the individual expects all others to choose low investment and it is the best response for him to choose low investment. Since the same calculus applies to each agent, this argument establishes the existence of the two symmetric equilibria. This simple game captures, in a very reduced-form way, the essence of the "Big Push" models discussed in Chapter 21.

Two points are worth noting. First, depending on the extent of complementarities and other economic interactions, y^H can be quite large relative to y^L, so there may be significant income differences in the allocations implied by the two different equilibria. Thus if we believe that such a game is a good approximation to reality and different countries can end up in different equilibria, it could help in explaining very large differences in income per capita. Second, the two equilibria in this game are also "Pareto-ranked"—all individuals are better-off in the equilibrium in which everybody chooses high investment.

In addition to models of multiple equilibria, stochastic models, in which the realization of certain random variables determine when a particular economy transitions from low-productivity to high-productivity technologies and starts the process of takeoff, might also be relevant in this context (see Section 17.6 in Chapter 17).

Both models of multiple equilibria and those in which stochastic variables determine the long-run growth properties of the economy are attractive as descriptions of certain aspects of the development process. They are also informative about the mechanics of economic development in an interesting class of models. But do they inform us about the fundamental causes of economic growth? Can we say that the United States is rich today while Nigeria is poor because the former has been lucky in its equilibrium selection while the latter has been unlucky? Can we pinpoint their divergent development paths to some small stochastic events 200, 300 or 400 years ago? The answer seems to be no.

US economic growth is the cumulative result of a variety of processes, ranging from innovations and free entrepreneurial activity to significant investments in human capital and rapid capital accumulation. It is difficult to reduce these to a simple lucky break or to the selection of the right equilibrium. Even 400 years ago, the historical conditions were very different in the United States and in Nigeria, and this led to different opportunities, different institutional paths and different incentives. It is the combination of the historical experiences of countries and different economic incentives that underlies their different processes of economic growth.

Equally important, models based on luck or multiple equilibria can explain why there might be a 20-year or perhaps a 50-year divergence between two otherwise-identical economies. But how are we to explain a 500-year divergence? It certainly does not seem plausible to imagine that Nigeria, today, can suddenly switch equilibria and quickly achieve

the level of income per capita in the United States.[2] Most models of multiple equilibria are unsatisfactory in another fundamental sense. As in the simple example discussed above, most models of multiple equilibria involve the presence of Pareto-ranked equilibria. This implies that one equilibrium gives higher utility or welfare to *all* agents than another. While such Pareto-ranked equilibria are a feature of our parsimonious models, which do not specify many relevant dimensions of heterogeneity that are important in practice, it is not clear whether they are useful in thinking about why some countries are rich and some are poor. If indeed it were possible for Nigerians to change their behavior and for all individuals in the nation to become better-off (say by switching from low to high investment in terms of the game above), it is very difficult to believe that for 200 years they have not been able to coordinate on such a better action. Most readers will be aware that Nigerian history is shaped by religious and ethnic conflict and by a civil war that ravaged the nation, and is still adversely affected by the extreme corruption of politicians, bureaucrats and soldiers that have enriched themselves at the expense of the population at large. That an easy Pareto improving change exists against this historical and social background seems improbable to say the least.

To be fair, not all models of "multiple equilibria" allow easy transitions from a Pareto inferior equilibrium to a superior equilibrium. In the literature, a useful distinction is between models of multiple equilibria, where different equilibria can be reached if individuals change their beliefs and behaviors simultaneously, versus models of *multiple steady states with history dependence*, where once a particular path of equilibrium is embarked upon, it becomes much harder (perhaps impossible) to transition to the other steady state equilibrium (see Chapter 21). Such models are much more attractive for understanding persistent differences in economic performance across countries. Nevertheless, unless some other significant source of conflict of interest or distortions are incorporated, it seems unlikely that the difference between the United States and Nigeria can be explained by using models where the two countries have identical parameters, but have made different choices and stuck with them. The mechanics of how a particular steady-state equilibrium can be maintained would be the most important element of such a theory, and other fundamental causes of economic growth, including institutions, policies or perhaps culture, must play a role in explaining this type of persistence. Put differently, in today's world of free information, technology and capital flows, if Nigeria had the same parameters, the same opportunities and the same "institutions" as the United States, there should exist some arrangement such that these new technologies can be imported and everybody could be made better-off.

Another challenge to models of multiple steady states concerns the ubiquity of growth miracles such as South Korea and Singapore, which we discussed in Chapter 1. If cross-country income differences are due to multiple steady states, from which escape is totally or nearly impossible, then how can we explain countries that embark upon a very rapid growth

[2]Naturally, one can argue that reforms or major changes in the growth trajectory are always outcomes of a switch from one equilibrium to another. But such an explanation would not have much empirical content, unless it is based on a well-formulated model of *equilibrium selection* and can make predictions about when we might expect such switches.

process? The example of China may be even more telling here. While China stagnated under communism until Mao's death, the changes in economic institutions and policies that took place thereafter have led to very rapid economic growth. If China was in a low-growth steady state before Mao's death, then we need to explain how it escaped from the steady state after 1978, and why it did not do so before. Inevitably this takes us to the role of other fundamental causes, such as institutions, policies and culture.

A different, and perhaps more promising, argument about the importance of luck can be made by emphasizing the role of leaders. Perhaps it was Mao who held back China, and his death and the identity, beliefs and policies of his successors were at the root of its subsequent growth. Perhaps the identity of the leader of a country can thus be viewed as a stochastic event, shaping economic performance. This point of view probably has a lot of merit. Recent empirical work by Jones and Olken (2005) shows that leaders seem to matter for the economic performance of nations. Thus luck could play a major role in cross-country income and growth differences by determining whether growth-enhancing or growth-retarding leaders are selected. Nevertheless, such an explanation is closer to the institutional approaches than the pure luck category. First, leaders often influence the economic performance of their societies by the policies they set and the institutions they develop. Thus, the selection and behavior of leaders and the policies that they pursue should be a part of the institutional explanations. Second, Jones and Olken's research points to an important interaction between the effect of leaders and a society's institutions. Leaders seem to matter for economic growth only in countries where institutions are non-democratic or weak (in the sense of not placing constraints on politicians or elites). In democracies and in societies where other institutions appear to place checks on the behavior of politicians and leaders, the identity of leaders seems to play almost no role in economic performance.

Given these considerations, I tentatively conclude that models emphasizing luck and multiple equilibria are useful for our study of the mechanics of economic development, but they are unlikely to provide us with the fundamental causes of why world economic growth started 200 years ago and why some countries are rich while others are poor today.

4.3.2. Geography. While the approaches in the last subsection emphasize the importance of luck and multiple equilibria among otherwise-identical societies, an alternative is to emphasize the deep heterogeneity across societies. The geography hypothesis is, first and foremost, about the fact that not all areas of the world are created equal. "Nature", that is, the physical, ecological and geographical environment of nations, plays a major role in their economic experiences. As pointed out above, geographic factors can play this role by determining both the preferences and the opportunity set of individual economic agents in different societies. There are at least three main versions of the geography hypothesis, each emphasizing a different mechanism for how geography affects prosperity.

The first and earliest version of the geography hypothesis goes back to Montesquieu ([1748], 1989). Montesquieu, who was a brilliant French philosopher and an avid supporter of Republican forms of government, was also convinced that climate was among the main

determinants of the fate of nations. He believed that climate, in particular heat, shaped human attitudes and effort, and via this channel, affected both economic and social outcomes. He wrote in his classic book *The Spirit of the Laws*:

> "The heat of the climate can be so excessive that the body there will be absolutely without strength. So, prostration will pass even to the spirit; no curiosity, no noble enterprise, no generous sentiment; inclinations will all be passive there; laziness there will be happiness,"
>
> "People are ... more vigorous in cold climates. The inhabitants of warm countries are, like old men, timorous; the people in cold countries are, like young men, brave..."

Today some of the pronouncements in these passages appear somewhat naïve and perhaps bordering on "political incorrectness". They still have many proponents, however. Even though Montesquieu's eloquence makes him stand out among those who formulated this perspective, he was neither the first nor the last to emphasize such geographic fundamental causes of economic growth. Among economists a more revered figure is one of the founders of our discipline, Alfred Marshall. Almost a century and a half after Montesquieu, Marshall wrote:

> "...vigor depends partly on race qualities: but these, so far as they can be explained at all, seem to be chiefly due to climate." (1890, p. 195).

While the first version of the geography hypothesis appears naïve and raw to many of us, its second version, which emphasizes the impact of geography on the technology available to a society, especially in agriculture, is more palatable and has many more supporters. This view is developed by an early Nobel Prize winner in economics, Gunnar Myrdal, who wrote

> "...serious study of the problems of underdevelopment ... should take into account the climate and its impacts on soil, vegetation, animals, humans and physical assets—in short, on living conditions in economic development." (1968, volume 3, p. 2121).

More recently, Jared Diamond, in his widely popular *Guns, Germs and Steel,* espouses this view and argues that geographical differences between the Americas and Europe (or more appropriately, Eurasia) have determined the timing and nature of settled agriculture and via this channel, shaped whether societies have been able to develop complex organizations and advanced civilian and military technologies (1997, e.g., p. 358). The economist Jeffrey Sachs has been a recent and forceful proponent of the importance of geography in agricultural productivity, stating that

> "By the start of the era of modern economic growth, if not much earlier, temperate-zone technologies were more productive than tropical-zone technologies ..." (2001, p. 2).

There are reasons for questioning this second, and more widely-held second view, of geographic determinism as well. Most of the technological differences emphasized by these

authors refer to agriculture. But as Chapter 1 emphasized, the origins of differential economic growth across countries goes back to the age of industrialization. Modern economic growth came with industry, and it is the countries that have failed to industrialize that are poor today. Low agricultural productivity, if anything, should create a comparative advantage in industry, and thus encourage those countries with the "unfavorable geography" to start investing in industry before others. One might argue that reaching a certain level of agricultural productivity is a prerequisite for industrialization. While this is plausible (at least possible), many of the societies that later failed to industrialize had already achieved a certain level of agricultural productivity, and in fact were often ahead of those who later industrialized very rapidly (see below). Thus a simple link between unfavorable agricultural conditions and the failure to take off seems to be absent.[3]

The third variant of the geography hypothesis, which has become particularly popular over the past decade, links poverty in many areas of the world to their "disease burden," emphasizing that: "The burden of infectious disease is ... higher in the tropics than in the temperate zones" (Sachs, 2000, p. 32). Bloom and Sachs (1998) and Gallup and Sachs (2001, p. 91) claim that the prevalence of malaria alone reduces the annual growth rate of sub-Saharan African economies by as much as 2.6 percent a year. Such a magnitude implies that had malaria been eradicated in 1950, income per capita in sub-Saharan Africa would have been double of what it is today. If we add to this the effect of other diseases, we would obtain even larger effects. The World Health Organization also subscribes to this view and in its recent report writes:

> "...in today's world, poor health has particularly pernicious effects on economic development in sub-Saharan Africa, South Asia, and pockets of high disease and intense poverty elsewhere..." (p. 24) and
>
> "...extending the coverage of crucial health services... to the world's poor could save millions of lives each year, reduce poverty, spur economic development and promote global security." (p. i).

This third version of the geography hypothesis may be much more plausible than the first two, especially since the microeconomics literature shows that unhealthy individuals are less productive and perhaps less able to learn and thus accumulate human capital. I will discuss both the general geography hypothesis and this specific version of it in greater detail below. But even at this point, an important caveat needs to be mentioned. The fact that the burden of disease is heavier in poor nations today is as much a consequence as a cause of poverty. European nations in the 18th and even 19th centuries were plagued by many diseases. The process of economic development enabled them to eradicate these diseases and create healthier environments for living. The fact that many poor countries have unhealthy environments is, at least in part, a consequence of their failure to develop economically.

[3]Ex post, one can in fact tell the opposite story: perhaps poor nations of today had agriculturally *superior* land, and this created a comparative advantage against industry and they failed to benefit from the increasing returns to scale in manufacturing.

4.3.3. Institutions. An alternative fundamental cause of differences in economic growth and income per capita is institutions. One problem with the institutions hypothesis is that it is somewhat difficult to define what "institutions" are. In daily usage, the word institutions refers to may different things, and the academic literature is sometimes not clear about its definition.

The economic historian Douglass North was awarded the Nobel Prize in economics largely because of his work emphasizing the importance of institutions in the historical development process. North (1990, p. 3) offers the following definition:

> "Institutions are the rules of the game in a society or, more formally, are the humanly devised constraints that shape human interaction."

He goes on to emphasize the key implications of institutions:

> "In consequence [institutions] structure incentives in human exchange, whether political, social, or economic."

This definition encapsulates the three important elements that make up institutions. First, they are "humanly devised"; that is, in contrast to geography, which is outside human control, institutions refer to man-made factors. Institutions are about the effect of the societies' own choices on their own economic fates. Second, institutions are about placing constraints on individual behavior. These do not need to be unassailable constraints. Any law can be broken, any regulation can be ignored. Nevertheless, policies, regulations and laws that punish certain types of behavior while rewarding others will naturally have an effect on behavior. And this brings the third important element in the definition. The constraints placed on individuals by institutions will shape human interaction and affect incentives. In some deep sense, institutions, much more than the other candidate fundamental causes, are about the importance of incentives.

The reader may have already noted that the above definition makes institutions a rather broad concept. In fact, this is precisely the sense in which I will use the concept of institutions throughout this book; institutions will refer to a *broad cluster* of arrangements that influence various economic interactions among individuals. These include economic, political and social relations among households, individuals and firms. The importance of political institutions, which determine the process of collective decision-making in society, cannot be overstated and will be the topic of analysis in Part 8 of this book.

A more natural starting point for the study of the fundamental causes of income differences across countries is with *economic institutions*, which comprise such things as the structure of property rights, the presence and (well or ill) functioning of markets, and the contractual opportunities available to individuals and firms. Economic institutions are important because they influence the structure of economic incentives in society. Without property rights, individuals will not have the incentive to invest in physical or human capital or adopt more efficient technologies. Economic institutions are also important because they ensure the allocation of resources to their most efficient uses, and they determine who obtains profits,

revenues and residual rights of control. When markets are missing or ignored (as was the case in many former socialist societies, for example), gains from trade go unexploited and resources are misallocated. Economic theory therefore suggests that societies with economic institutions that facilitate and encourage factor accumulation, innovation and the efficient allocation of resources should prosper relative to societies that do not have such institutions.

The hypothesis that differences in economic institutions are a fundamental cause of different patterns of economic growth is intimately linked to the models we will develop in this book. In all of our models, especially in those that endogenize physical capital, human capital and technology accumulation, individuals will respond to (profit) incentives. Economic institutions, determined by the way in which individuals organize their societies, shape these incentives. Some ways of organizing societies encourage people to innovate, to take risks, to save for the future, to find better ways of doing things, to learn and educate themselves, to solve problems of collective action and to provide public goods. Others do not. Our theoretical models will then pinpoint exactly what specific policy and institutional variables are important in retarding or encouraging economic growth.

Part 8 of the book will develop theoretical approaches to the analysis of what "good economic institutions" that encourage physical and human capital accumulation and the development and adoption of better technologies may be (though "good economic institutions" may change from environment to environment and from time to time). It should already be intuitive to the reader that economic institutions that tax productivity-enhancing activities will not encourage economic growth. Economic institutions that ban innovation will not lead to technological improvements. Therefore, enforcement of some basic *property rights* and some amount of free enterprise are indispensable. But other aspects of economic institutions matter as well. Human capital, for example, is important both for increasing productivity and for technology adoption. However, for a broad cross-section of society to be able to accumulate human capital, some degree of equality of opportunity is necessary. Economic institutions that only protect the rights of a rich elite or the already-privileged will not achieve such equality of opportunity and will often create other distortions, potentially retarding economic growth. Chapter 14 will emphasize that the process of Schumpeterian creative destruction, where new firms improve over and destroy incumbents, is an essential element of economic growth. Schumpeterian creative destruction requires a level playing field, so that incumbents are unable to block technological progress. Economic growth based on creative destruction therefore also requires economic institutions that guarantee some degree of equality of opportunity in the society.

Another question may have already occurred to the reader: why should any society have economic and political institutions that retard economic growth? Would it not be better for all parties to maximize the size of the national pie (level of GDP, consumption or economic growth)? There are two possible answers to this question. The first takes us back to multiple equilibria. It may be that the members of the society cannot coordinate on the "right," that is, growth-enhancing, institutions. This answer is not satisfactory for the same reasons

as other broad explanations based on multiple equilibria are unsatisfactory; if there exists an equilibrium institutional improvement that will make *all* members of a society richer and better-off, it seems unlikely that the society will be unable to coordinate on this improvement for extended periods of time.

The second answer, instead, recognizes that there are inherent *conflicts of interest* within the society. There are no reforms, no changes, no advances that would make everybody better-off; as in the Schumpeterian creative destruction stories, every reform, every change and every advance creates winners and losers. Part 8 will show that institutional explanations are intimately linked with the conflicts of interests in society. Put simply, the distribution of resources cannot be separated from the aggregate economic performance of the economy—or perhaps in a more familiar form, efficiency and distribution cannot be decoupled. This implies that institutions that fail to maximize the growth potential of an economy may nevertheless create benefits for some segments of the society, who will then form a constituency in favor of these institutions. Thus to understand the sources of institutional variations we have to study the winners and losers of different institutional reforms and why winners are unable to buy off or compensate losers, and why they are not powerful enough to overwhelm the losers, even when the institutional change in question may increase the size of the national pie. Such a study will not only help us understand why some societies choose or end up with institutions that do not encourage economic growth, but will also enable us to make predictions about institutional change. After all, the fact that institutions can and do change is a major difference between the institutions hypothesis and the geography and culture hypotheses. Questions of equilibrium institutions and endogenous institutional change are central for the institutions hypothesis, but must be postponed until Part 8. Here, note that the endogeneity of institutions has another important implication; endogeneity of institutions makes empirical work on assessing the role of institutions more challenging, because it implies that the standard "simultaneity" biases in econometrics will be present when we look at the effect of institutions on economic outcomes.[4]

In this chapter, I focus on the empirical evidence in favor and against the various different hypotheses. by argue that this evidence, by and large, suggests that institutional differences that societies choose and end up with are a primary determinant of their economic fortunes. The further discussion below and a summary of recent empirical work will try to bolster this case. Nevertheless, it is important to emphasize that this does not mean that only institutions matter and luck, geography and culture are not important. The four fundamental causes are potentially complementary. The evidence suggests that institutions are the most important one among these four causes, but does not deny the potential role of other factors, such as cultural influences.

[4]Note also that geography is indeed "exogenous" in the sense that, with some notable exceptions such as climate change and global warming, it is not much influenced by economic decisions. But this does not make it "econometrically exogenous". Geographic characteristics may still be (and in fact likely are) correlated with other factors that influence economic growth.

4.3.4. Culture. The final fundamental explanation for economic growth emphasizes the idea that different societies (or perhaps different races or ethnic groups) have different cultures, because of different shared experiences or different religions. Culture is viewed, by some social scientists, as a key determinant of the values, preferences and beliefs of individuals and societies and, the argument goes, these differences play a key role in shaping economic performance.

At some level, culture can be thought of as influencing equilibrium outcomes for a given set of institutions. Recall that in the presence of multiple equilibria, there is a central question of equilibrium selection. For example, in the simple game discussed above, perhaps culture may affect whether the society coordinates on the high-investment or the low-investment equilibrium. "Good" cultures can be thought of as ways of coordinating on better (Pareto superior) equilibria. Naturally, the arguments discussed above, that an entire society being stuck in an equilibrium in which *all* individuals are worse-off than in an alternative equilibrium is implausible, would militate against the importance of this particular role of culture. Alternatively, different cultures generate different sets of beliefs about how people behave and this can alter the set of equilibria for a given specification of institutions (for example, some beliefs will allow punishment strategies to be used whereas others will not).

The most famous link between culture and economic development is that proposed by Max Weber (1930), who argued that the origins of industrialization in Western Europe could be traced to a cultural factor—the Protestant reformation and particularly the rise of Calvinism. Interestingly, Weber provided a clear summary of his views as a comment on Montesquieu's arguments:

> "Montesquieu says of the English that they 'had progressed the farthest of all peoples of the world in three important things: in piety, in commerce, and in freedom'. Is it not possible that their commercial superiority and their adaptation to free political institutions are connected in some way with that record of piety which Montesquieu ascribes to them?"

Weber argued that English piety, in particular, Protestantism, was an important driver of capitalists development. Protestantism led to a set of beliefs that emphasized hard work, thrift and saving. It also interpreted economic success as consistent with, even as signalling, being chosen by God. Weber contrasted these characteristics of Protestantism with those of other religions, such as Catholicism and other religions, which he argued did not promote capitalism. More recently, similar ideas have been applied to emphasize different implications of other religions. Many historians and scholars have argued that the rise of capitalism, the process of economic growth and industrialization are intimately linked to cultural and religious beliefs. Similar ideas are also proposed as explanations for why Latin American countries, with their Iberian heritage, are poor and unsuccessful, while their North American neighbors are more prosperous because of their Anglo-Saxon culture.

A related argument, originating in anthropology, argues that societies may become "dysfunctional" because their cultural values and their system of beliefs do not encourage cooperation. An original and insightful version of this argument is developed in Banfield's (1958) analysis of poverty in Southern Italy. His ideas were later popularized by Putnam (1993), who suggested the notion of "social capital," as a stand-in for cultural attitudes that lead to cooperation and other "good outcomes". Many versions of these ideas are presented in one form or another in the economics literature as well.

Two challenges confront theories of economic growth based on culture. The first is the difficulty of measuring culture. While both Putnam himself and some economists have made some progress in measuring certain cultural characteristics with self-reported beliefs and attitudes in social surveys, simply stating that the North of Italy is rich because it has good social capital while the South is poor because it has poor social capital runs the risk of circularity. The second difficulty confronting cultural explanations is for accounting for growth miracles, such as those of South Korea and Singapore. As mentioned above, if some Asian cultural values are responsible for the successful growth experiences of these countries, it becomes difficult to explain why these Asian values did not lead to growth before. Why do these values not spur economic growth in North Korea? If Asian values are important for Chinese growth today, why did they not lead to a better economic performance under Mao's dictatorship? Both of these challenges are, in principle, surmountable. One may be able to develop models of culture, with better mapping to data, and also with an associated theory of how culture may change rapidly under certain circumstances. While possible in principle, such theories have not been developed yet. Moreover, the evidence presented in the next section suggests that cultural effects are not the major force behind the large differences in economic growth experienced by many countries over the past few centuries.

4.4. The Effect of Institutions on Economic Growth

I now argue that there is convincing empirical support for the hypothesis that differences in economic institutions, rather than luck, geography or culture, *cause* differences in incomes per-capita. Let us start by looking at the simplest correlation between a measure of economic institutions and income per capita.

Figure 4.1 shows the cross-country correlation between the log of GDP per-capita in 1995 and a broad measure of property rights, "protection against expropriation risk", averaged over the period 1985 to 1995. The data on this measure of economic institutions come from Political Risk Services, a private company which assesses the expropriation risk that foreign investments face in different countries. These data are not perfect. They reflect the subjective assessments of some analysts about how secure property rights are. Nevertheless, they are useful for our purposes. First, they emphasize the security of property rights, which is an essential aspect of economic institutions, especially in regards to their effect on economic incentives. Second, these measures are purchased by businessmen contemplating investment in these countries, thus they reflect the "market assessment" of security of property rights.

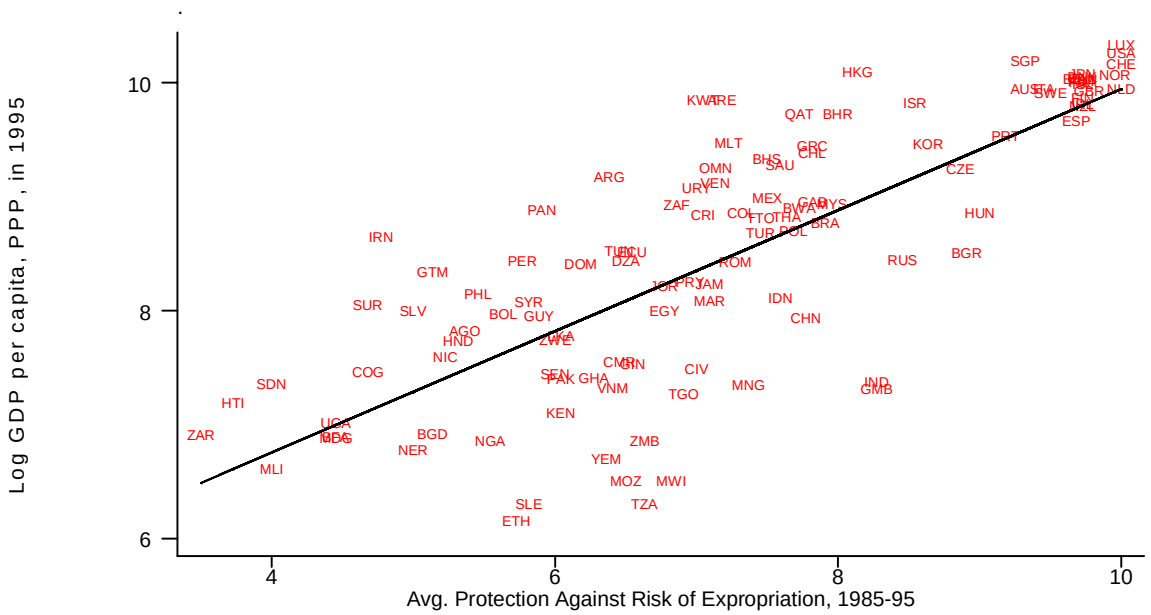

Figure 4.1. Relationship between economic institutions, as measured by average expropriation risk 1985-1995, and GDP per capita.

Figure 4.1 shows that countries with more secure property rights—thus better economic institutions—have higher average incomes. One should not interpret the correlation in this figure as depicting a causal relationship—that is, as establishing that secure property rights cause prosperity. First, the correlation might reflect reverse causation; it may be that only countries that are sufficiently wealthy can afford to enforce property rights. Second and more importantly, there might be a problem of omitted variable bias. It could be something else, for example, geography or culture, that explains both why countries are poor and why they have insecure property rights. Thus if omitted factors determine institutions and incomes, we would spuriously infer the existence of a causal relationship between economic institutions and incomes when in fact no such relationship exists. This is the standard identification problem in econometrics resulting from simultaneity or omitted variable biases. Finally, security of property rights—or other proxy measures of economic institutions—are themselves equilibrium outcomes, presumably resulting from the underlying political institutions and political conflict. While this last point is important, a satisfactory discussion requires formal models of political economy of institutions, which will have to wait until Part 8.

To further illustrate these potential identification problems, suppose that climate or geography matter for economic performance. In fact, a simple scatterplot shows a positive association between latitude (the absolute value of distance from the equator) and income per capita consistent with the views of Montesquieu and other proponents of the geography hypothesis. Interestingly, Montesquieu not only claimed that warm climate makes people

lazy and thus unproductive, but also unfit to be governed by democracy. He argued that despotism would be the political system in warm climates. Therefore, a potential explanation for the patterns in Figure 4.1 is that there is an omitted factor, geography, which explains both economic institutions and economic performance. Ignoring this potential third factor would lead to mistaken conclusions.

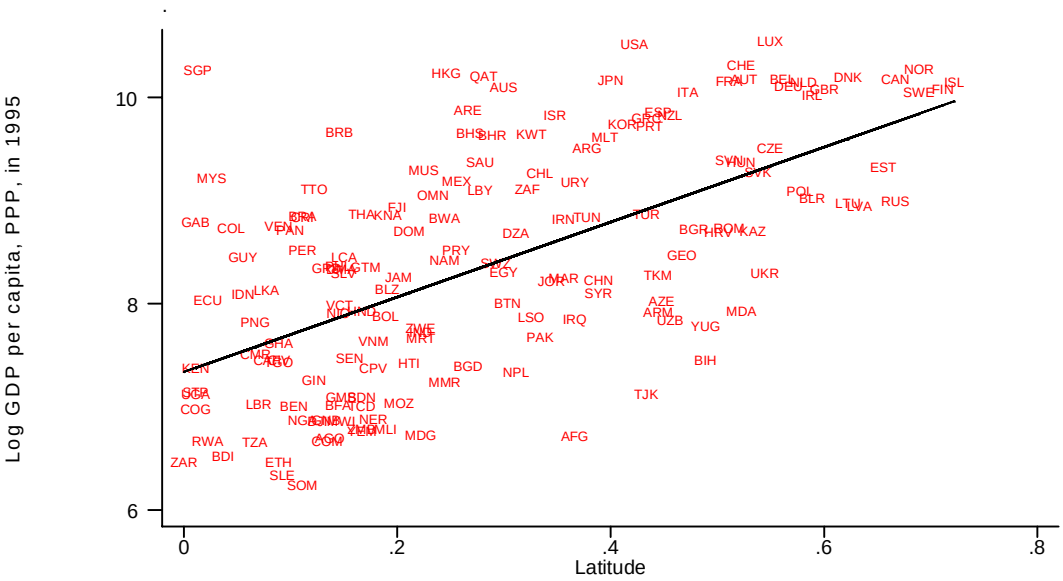

FIGURE 4.2. Relationship between latitude (distance of capital from the equator) and income per capita in 1995.

Even if Montesquieu's story appears both unrealistic and condescending to our modern sensibilities, the general point should be taken seriously: the correlations depicted in Figure 4.1, and for that matter that shown in Figure 4.2, do not necessarily reflect causal relationships. As noted in the context of the effect of religion or social capital on economic performance, these types of scatterplots, correlations, or their multidimensional version in ordinary least squares regressions, *cannot* establish causality. A residual doubt about the effect of omitted variables will almost always remain, even after careful regression analyses.

How can we overcome the challenge of establishing a causal relationship between (economic) institutions and economic outcomes? The answer to this question is to specify econometric approaches based on convincing identifying restrictions. This can be done by estimating structural econometric models or by using more reduced-form approaches, based on instrumental-variables strategies. At the moment we do not know enough about the evolution of economic institutions and their impact on economic outcomes to be able to specify and estimate fully-structural econometric models. Thus as a first step, we can look at more reduced-form evidence that might still be informative about the causal relationship between

institutions and economic growth. One way of doing so is to learn from history, in particular from the "natural experiments", which are unusual historical events where, while other fundamental causes of economic growth are held constant, institutions change because of potentially-exogenous reasons. I now discuss lessons from two such natural experiments.

4.4.1. The Korean Experiment. Until the end of World War II, Korea was under Japanese occupation. Korean independence came shortly after the war. The major fear of the United States during this time period was the takeover of the entire Korean peninsula either by the Soviet Union or by communist forces under the control of the former guerrilla fighter, Kim Il Sung. US authorities therefore supported the influential nationalist leader Syngman Rhee, who was in favor of separation rather than a united communist Korea. Elections in the South were held in May 1948, amidst a widespread boycott by Koreans opposed to separation. The newly elected representatives proceeded to draft a new constitution and established the Republic of Korea to the south of the 38th parallel. The North became the Democratic People's Republic of Korea, under the control of Kim Il Sung.

These two independent countries organized themselves in radically different ways and adopted completely different sets of (economic and political) institutions. The North followed the model of Soviet socialism and the Chinese Revolution in abolishing private property in land and capital. Economic decisions were not mediated by the market, but by the communist state. The South instead maintained a system of private property and capitalist economic institutions.

Before this "natural experiment" in institutional change, North and South Korea shared the same history and cultural roots. In fact, Korea exhibited an unparalleled degree of ethnic, linguistic, cultural, geographic and economic homogeneity. There are few geographic distinctions between the North and South, and both share the same disease environment. Moreover, before the separation the North and the South were at the same level of development. If anything, there was slightly more industrialization in the North. Maddison (2001) estimates that at the time of separation, North and South Korea had approximately the same income per capita.

We can therefore think of the splitting on the Koreas 60 years ago as a natural experiment that we can use to identify the causal influence of institutions on prosperity. Korea was split into two, with the two halves organized in radically different ways, and with geography, culture and many other potential determinants of economic prosperity held fixed. Thus any differences in economic performance can plausibly be attributed to differences in institutions.

In the 60 years following the split, the two Koreas have experienced dramatically diverging paths of economic development since separation. By the late 1960's South Korea was transformed into one of the Asian "miracle" economies, experiencing one of the most rapid surges of economic prosperity in history while North Korea stagnated. By 2000 the level of income in South Korea was $16,100 while in North Korea it was only $1,000. There is only one plausible explanation for the radically different economic experiences of the two Koreas after 1950: their very different institutions led to divergent economic outcomes. In this context,

it is noteworthy that the two Koreas not only shared the same geography, but also the same culture, so that neither geographic nor cultural differences could have much to do with the divergent paths of the two Koreas. Of course one can say that South Korea was lucky while the North was unlucky (even though this was not due to any kind of multiple equilibria, but a result of the imposition of different institutions). Nevertheless, the perspective of "luck" is unlikely to be particularly useful in this context, since what is remarkable is the persistence of the dysfunctional North Korean institutions. Despite convincing evidence that the North Korean system has been generating poverty and famine, the leaders of the Communist Party in North Korea have opted to use all the means available to them to maintain their regime.

However convincing on its own terms, the evidence from this natural experiment is not sufficient for the purposes of establishing the importance of economic institutions as the primary factor shaping cross-country differences in economic prosperity. First, this is only one case, and in the better-controlled experiments in the natural sciences, a relatively large sample is essential. Second, here we have an example of an extreme case, the difference between a market-oriented economy and an extreme communist one. Few social scientists today would deny that a lengthy period of totalitarian centrally-planned rule has significant economic costs. And yet, many might argue that differences in economic institutions among capitalist economies or among democracies are not the major factor leading to differences in their economic trajectories. To establish the major role of economic institutions in the prosperity and poverty of nations we need to look at a larger scale "natural experiment" in institutional divergence.

4.4.2. The Colonial Experiment: The Reversal of Fortune. The colonization of much of the world by Europeans provides such a large scale natural experiment. Beginning in the early fifteenth century and especially after 1492, Europeans conquered many other nations. The colonization experience transformed the institutions in many diverse lands conquered or controlled by Europeans. Most importantly, Europeans imposed very different sets of institutions in different parts of their global empire, as exemplified most sharply by the contrast of the institutional structure that developed in the Northeastern United States, based on small-holder private property and democracy, versus the institutions in the Caribbean plantation economies, based on repression and slavery. As a result, while geography was held constant, Europeans initiated very significant changes in the economic institutions of different societies.

The impact of European colonialism on economic institutions is perhaps most dramatically conveyed by a single fact—historical evidence shows that there has been a remarkable Reversal of Fortune in economic prosperity within former European colonies. Societies like the Mughals in India, and the Aztecs and the Incas in the Americas were among the richest civilizations in 1500, yet the nation-states that now exist in their boundaries are among the poorer nations of today. In contrast, countries occupying the territories of the less-developed civilizations of North America, New Zealand and Australia are now much richer than those in the lands of the Mughals, Aztecs and Incas.

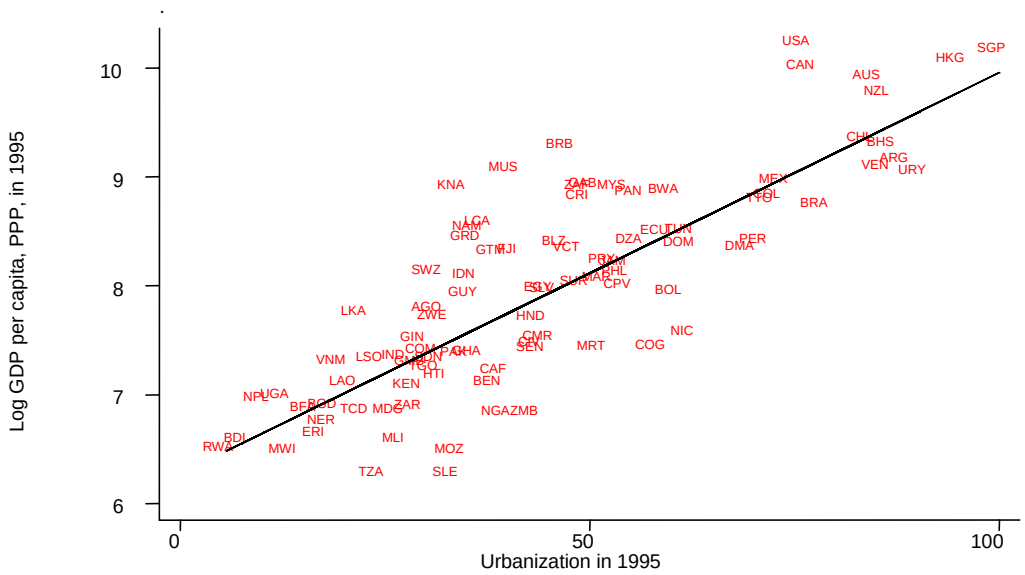

Figure 4.3. Urbanization and Income today.

The Reversal of Fortune is not confined to such comparisons. To document the reversal more broadly, we need a proxy for prosperity 500 years ago. Fortunately, urbanization rates and population density can serve the role of such proxies. Only societies with a certain level of productivity in agriculture and a relatively developed system of transport and commerce can sustain large urban centers and a dense population. Figure 4.3 shows the relationship between income per capita and urbanization (fraction of the population living in urban centers with greater than 5,000 inhabitants) today, and demonstrates that even today, long after industrialization, there is a significant relationship between urbanization and prosperity.

Naturally, high rates of urbanization do not mean that the majority of the population lived in prosperity. In fact, before the twentieth century urban areas were often centers of poverty and ill health. Nevertheless, urbanization is a good proxy for average prosperity and closely corresponds to the GDP per capita measures we are using to look at prosperity today. Another variable that is useful for measuring pre-industrial prosperity is the density of the population, which is closely related to urbanization.

Figures 4.4 and 4.5 show the relationship between income per capita today and urbanization rates and (log) population density in 1500 for the sample of (former) European colonies. Let us focus on 1500 since it is before European colonization had an effect on any of these societies. A strong negative relationship, indicating a reversal in the rankings in terms of economic prosperity between 1500 and today, is clear in both figures. In fact, the figures show that in 1500 the temperate areas were generally less prosperous than the tropical areas, but this pattern too was reversed by the twentieth century.

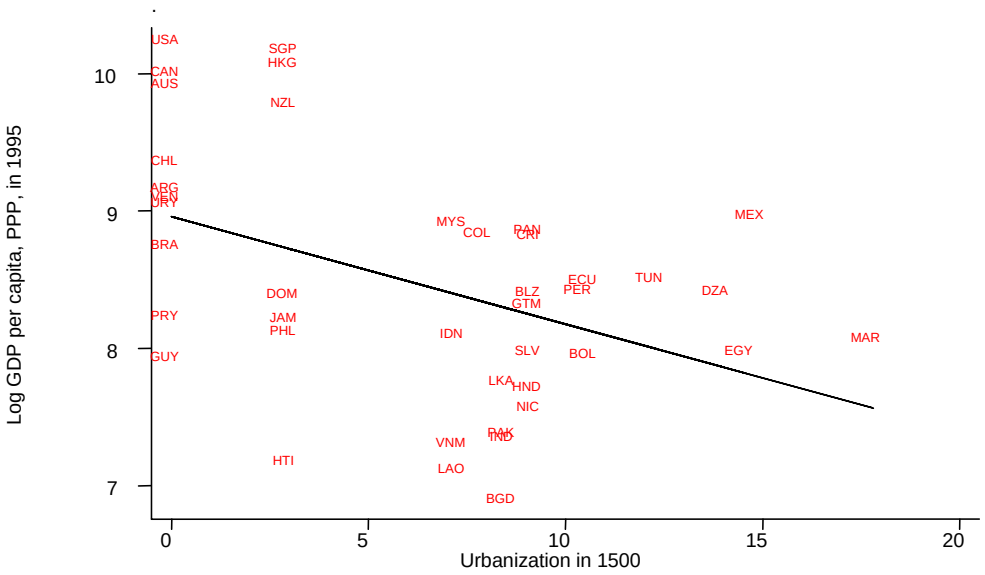

FIGURE 4.4. Reversal of Fortune: urbanization in 1500 versus income per capita in 1995 among the former European colonies.

FIGURE 4.5. Reversal of Fortune: population density in 1500 versus income per capita in 1995 them on the former European colonies.

There is something extraordinary and unusual about this reversal. A wealth of evidence shows that after the initial spread of agriculture there was remarkable persistence in urbanization and population density for all countries, including those that were subsequently

colonized by Europeans. Extending the data on urbanization to earlier periods shows that both among former European colonies and non-colonies, urbanization rates and prosperity persisted for 500 years or longer. Even though there are prominent examples of the decline and fall of empires, such as Ancient Egypt, Athens, Rome, Carthage and Venice, the overall pattern was one of persistence. It is also worth noting that reversal was not the general pattern in the world after 1500. When we look at Europe as a whole or at the entire world excluding the former European colonies, there is no evidence of a similar reversal between 1500 and 1995.

There is therefore no reason to think that what is going on in Figures 4.4 and 4.5 is some sort of natural reversion to the mean. Instead, the Reversal of Fortune among the former European colonies reflects something unusual, something related to the intervention that these countries experienced. The major intervention, of course, was related to the change in institutions. Not only did the Europeans impose a different order in almost all countries they conquered, there were also tremendous differences between the types of institutions they imposed on in the different colonies.[5] These institutional differences among the former colonies are likely at the root of the reversal in economic fortunes. This conclusion is bolstered further when we look at the timing and the nature of the reversal. Acemoglu, Johnson and Robinson (2002) show that the reversal took place largely in the 19th century and appears to be closely connected to industrialization.

These patterns are clearly inconsistent with simple geography based views of relative prosperity. In 1500 it was the countries in the tropics which were relatively prosperous, today it is the reverse. This makes it implausible to base a theory of relative prosperity on the intrinsic poverty of the tropics, climate, disease environments or other fixed characteristics.

Nevertheless, following Diamond (1997), one could propose what Acemoglu, Johnson and Robinson (2002) call a "sophisticated geography hypothesis," that geography matters but in a time-varying manner. For example, Europeans created "latitude specific" technologies, such as heavy metal ploughs, that only worked in temperate latitudes and not with tropical soils. Thus when Europe conquered most of the world after 1492, they introduced specific technologies that functioned in some places (the United States, Argentina, Australia) but not others (Peru, Mexico, West Africa). However, the timing of the reversal in the nineteenth century is inconsistent with the most natural types of sophisticated geography hypotheses. Europeans may have had latitude specific technologies, but the timing implies that these technologies must have been industrial, not agricultural, and it is difficult to see why industrial technologies do not function in the tropics (and in fact, they have functioned quite successfully in tropical Singapore and Hong Kong).

[5]In some instances, including those in Central America and India, the colonial institutions were built on the pre-colonial institutions. In these cases, the issue becomes one of whether Europeans maintained and further developed existing hierarchical institutions, such as those in the Aztec, the Inca or the Mughal Empires, or whether they introduced or imposed political and economic institutions encouraging broad-based participation and investment.

Similar considerations weigh against the culture hypothesis. Although culture is slow-changing the colonial experiment was sufficiently radical to have caused major changes in the cultures of many countries that fell under European rule. In addition, the destruction of many indigenous populations and the immigration from Europe are likely to have created new cultures or at least modified existing cultures in major ways. Nevertheless, the culture hypothesis does not provide a natural explanation for the reversal and has nothing to say on the timing of the reversal. Moreover, as discussed below, econometric models that control for the effect of institutions on income do not find any evidence of an effect of religion or culture on prosperity.

The importance of luck is also limited. The different institutions imposed by the Europeans were not random. They were instead very much related to the conditions they encountered in the colonies. In other words, the types of institutions that were imposed and developed in the former colonies were endogenous outcomes, outcomes of equilibria that we need to study.

4.4.3. The Reversal and the Institutions Hypothesis. Is the Reversal of Fortune consistent with a dominant role for economic institutions in comparative development? The answer is yes. In fact, once we recognize the variation in economic institutions created by colonization, we see that the Reversal of Fortune is exactly what the institutions hypothesis predicts.

The evidence in Acemoglu, Johnson and Robinson (2002a) shows a close connection between initial population density, urbanization, and the creation of good economic institutions. In particular, the evidence points out that, others things equal, the higher the initial population density or the greater initial urbanization, the worse were subsequent institutions, including both institutions right after independence and also institutions today. Figures 4.6 and 4.7 illustrate these relationships using the same measure of current economic institutions used in Figure 4.1, protection against expropriation risk today. They document that the relatively densely settled and highly urbanized colonies ended up with worse institutions, while sparsely-settled and non-urbanized areas received an influx of European migrants and developed institutions protecting the property rights of a broad cross-section of society. European colonialism therefore led to an "institutional reversal," in the sense that the previously-richer and more-densely settled places ended up with worse institutions. The institutional reversal does not mean that institutions were better in the previously more densely-settled areas. It only implies a tendency for the relatively poorer and less densely-settled areas to end up with better institutions than previously-rich and more densely-settled areas.

As discussed in footnote 5 above, it is possible that the Europeans did not actively introduce institutions discouraging economic progress in many of these places, but inherited them from previous civilizations there. The structure of the Mughal, Aztec and Inca empires were already very hierarchical with power concentrated in the hands of narrowly based ruling elites and structured to extract resources from the majority of the population for the benefit of a minority. Often Europeans simply took over these existing institutions. What is important

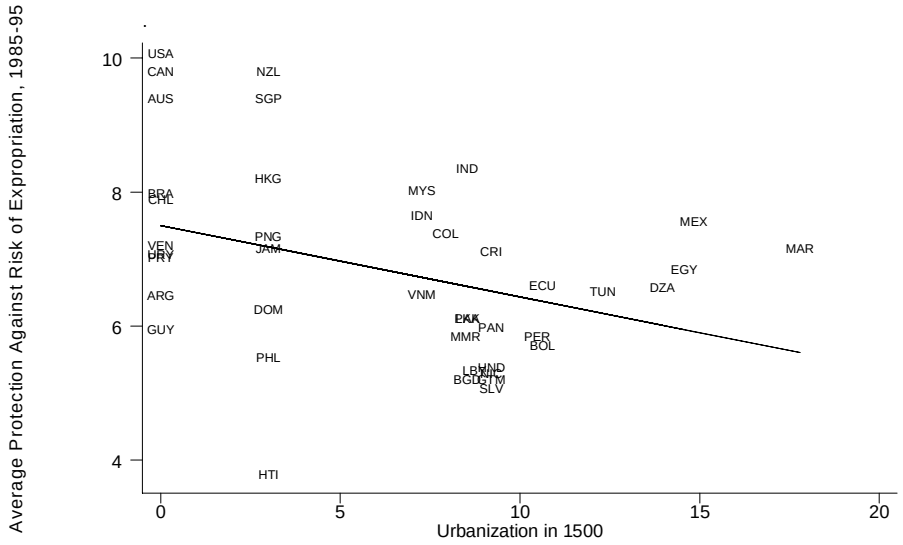

Figure 4.6. The Institutional Reversal: urbanization in 1500 and economic institutions today among the former European colonies.

in any case is that in densely-settled and relatively-developed places it was in the interests of Europeans to have institutions facilitating the extraction of resources, without any respect for the property rights of the majority of the populace. In contrast, in the sparsely-settled areas it was in their interests to develop institutions protecting property rights. These incentives led to an institutional reversal.

The institutional reversal, combined with the institutions hypothesis, predicts the Reversal of Fortune: relatively rich places ended up with relatively worse economic institutions. And if these institutions are important, we should see them become relatively poor over time. This is what the Reversal of Fortune shows.

Moreover, the institutions hypothesis is consistent with the timing of the reversal. Recall that the institutions hypothesis links incentives to invest in physical and human capital and in technology to economic institutions, and argues that economic prosperity results from these investments. Therefore, we expect economic institutions to play a more important role in shaping economic outcomes when there are major new investment opportunities—thus creating greater need for entry by new entrepreneurs and for the process of creative destruction. The opportunity to industrialize was the major investment opportunity of the 19th century. As documented in Chapter 1, countries that are rich today, both among the former European colonies and other countries, are those that industrialized successfully during this critical period.

The explanation for the reversal that emerges from the discussion so far is one in which the economic institutions in various colonies were shaped by Europeans to serve their own (economic) interests. Moreover, because conditions and endowments differed between colonies,

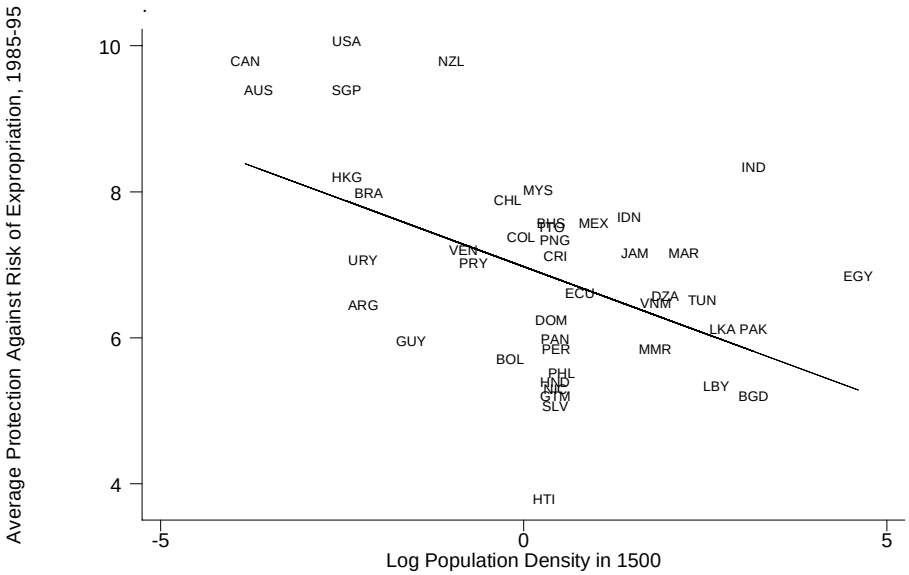

Figure 4.7. The Institutional Reversal: population density in 1500 and economic institutions today among the former European colonies.

Europeans consciously created different economic institutions, which, in many cases, still persist and continue to shape economic performance. Why did Europeans introduce better institutions in previously-poor and unsettled areas than in previously-rich and densely-settled areas? Without going into details, a number of obvious ideas that have emerged from the research in this area can be mentioned.

Europeans were more likely to introduce or maintain economic institutions facilitating the extraction of resources in areas where they would benefit from the extraction of resources. This typically meant areas controlled by a small group of Europeans, as well as areas offering resources to be extracted. These resources included gold and silver, valuable agricultural commodities such as sugar, but most importantly, what is perhaps the most valuable commodity overall, human labor. In places with a large indigenous population, Europeans could exploit the population. This was achieved in various forms, using taxes, tributes or employment as forced labor in mines or plantations. This type of colonization was incompatible with institutions providing economic or civil rights to the majority of the population. Consequently, a more developed civilization and a denser population structure made it more profitable for the Europeans to introduce worse economic institutions.

In contrast, in places with little to extract, and in sparsely-settled places where the Europeans themselves became the majority of the population, it was in their interests to introduce economic institutions protecting their own property rights (and also to attract further settlers).

4.4.4. Settlements, Mortality and Development. The initial conditions of the colonies emphasized so far, indigenous population density and urbanization, are not the only factors affecting Europeans' colonization strategy. In addition, the disease environments differed markedly among the colonies, with obvious consequences on the attractiveness of European settlement. As noted above, when Europeans settled, they established institutions that they themselves had to live under, so whether Europeans could settle or not had a major effect on the subsequent path of institutional development. In other words, the disease environment 200 or more years ago, especially the prevalence of malaria and yellow fever which crucially affected potential European mortality, may have crucially affected the paths of institutional and economic development in the former European colonies. If in addition, the disease environment of the colonial times affects economic outcomes today only through its effect on institutions, then this historical disease environment can be used as an exogenous source of variation in current institutions. From an econometric point of view, this will correspond to a valid instrument to estimate the casual effect of economic institutions on prosperity. Although mortality rates of potential European settlers could be correlated with indigenous mortality, which may determine income today, in practice local populations had developed much greater immunity to malaria and yellow fever. Acemoglu, Johnson and Robinson (2001) present a variety of evidence suggesting that the major effect of European settler mortality is through institutions.

In particular, Acemoglu, Johnson and Robinson (2001) argue that the European colonization strategy was influenced by the feasibility of settlements. In places where the disease environment was not favorable to European settlement, "extractive institutions," used for extracting resources from the indigenous population, were more likely. These institutions did not provide security of property rights to the majority of the population. On the contrary, the majority of the indigenous population were without any economic or political rights and lived under severe oppression. In contrast, in places where the disease environment enabled European settlements, it was possible for the Europeans to settle in large numbers, and in this case, they would be developing institutions under which they themselves would live. Such institutions were more likely to protect property rights, encourage investment and provide broad-based participation in economic and political life. Moreover, the colonial state and institutions persisted to some degree and make it more likely that former European colonies that suffered extractive colonization have worse institutions today. Summarizing this argument schematically:

$$\begin{array}{c}\text{(potential) settler}\\ \text{mortality}\end{array} \Rightarrow \text{settlements} \Rightarrow \begin{array}{c}\text{early}\\ \text{institutions}\end{array} \Rightarrow \begin{array}{c}\text{current}\\ \text{institutions}\end{array} \Rightarrow \begin{array}{c}\text{current}\\ \text{performance}\end{array}$$

Based on this reasoning, Acemoglu, Johnson and Robinson (2001) use the mortality rates expected by the first European settlers in the colonies as an instrument for current institutions in the sample of former European colonies. Their instrumental-variables estimates show a large and robust effect of institutions on economic growth and income per capita. Figures 4.8 and 4.9 provide an overview of the evidence. Figure 4.8 shows the cross-sectional relationship

between income per capita and the measure of economic institutions we encountered in Figure 4.1, protection against expropriation risk. It shows a very strong relationship between the historical mortality risk faced by Europeans and the current extent to which property rights are enforced. A bivariate regression has an R^2 of 0.26. It also shows that there were very large differences in European mortality. Countries such as Australia, New Zealand and the United States were very healthy, and existing evidence suggests that life expectancy in Australia and New Zealand was in fact greater than in Britain. In contrast, Europeans faced extremely high mortality rates in Africa, India and South-East Asia. These differential mortality rates were largely due to tropical diseases such as malaria and yellow fever and at the time it was not understood how these diseases arose nor how they could be prevented or cured.

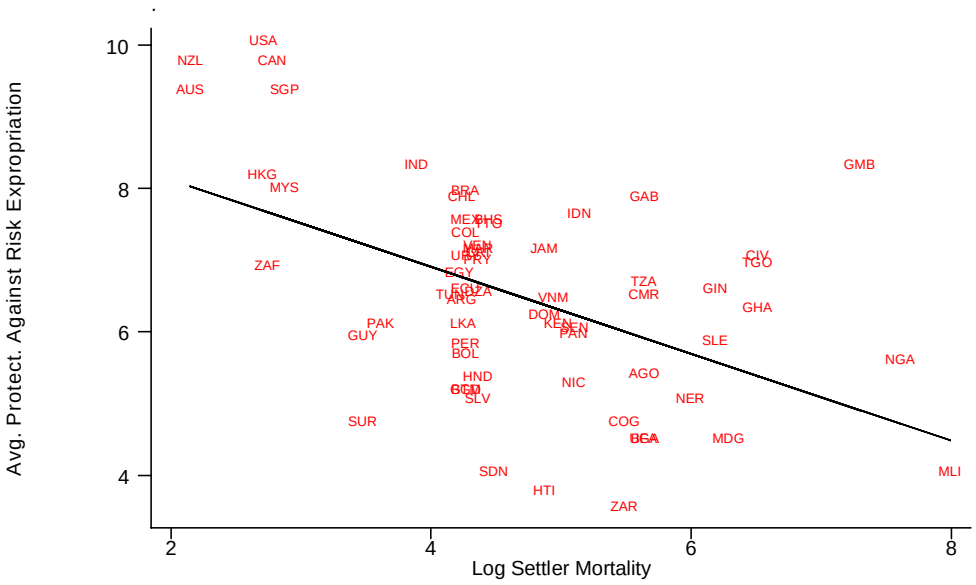

Figure 4.8. The relationship between mortality of potential European settlers and current economic institutions.

Figures 4.8 and 4.9 already show that, if the exclusion restriction, that the mortality rates of potential European settlers should have no effect on current economic outcomes other than through institutions, is valid, then there is a large impact of economic institutions on economic performance. This is documented in detail in Acemoglu, Johnson and Robinson (2001), who present a range of robustness checks confirming this result. Their estimates suggest that most of the gap between rich and poor countries today is due to differences in economic institutions. For example, the evidence suggests that over 75% of the income gap between relatively rich and relatively poor countries can be explained by differences in their economic institutions (as proxied by security of property rights). Equally important, the evidence indicates that once the effect of institutions is estimated via this methodology, there appears to be no effect of geographical variables; neither latitude, nor whether or not

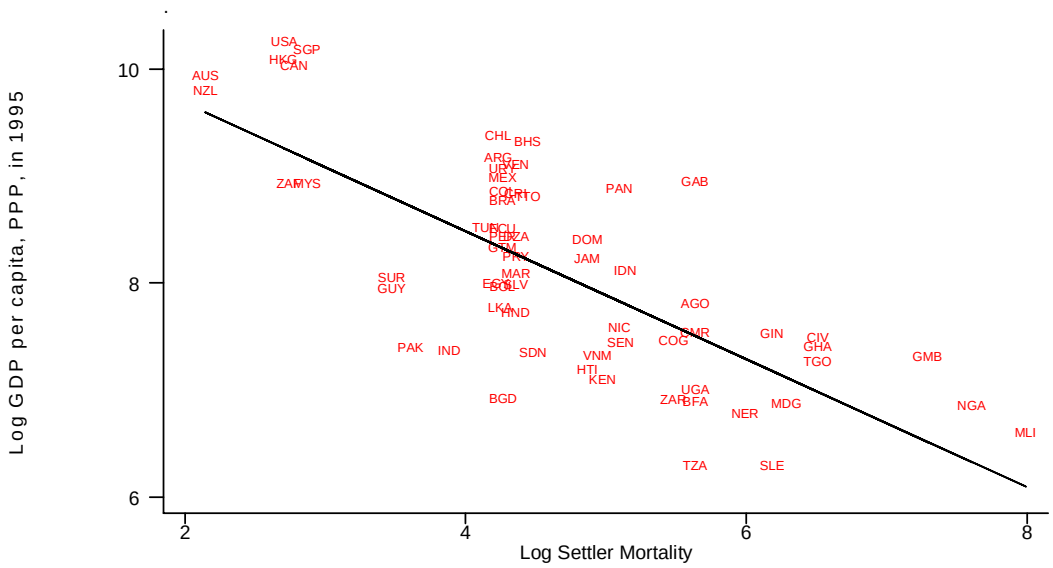

FIGURE 4.9. The relationship between mortality of potential European settlers and GDP per capita in 1995.

a country is land-locked nor the current disease environment appear to have much effect on current economic outcomes. This evidence again suggests that institutional differences across countries are a major determinant of their economic fortunes, while geographic differences are much less important.

These results also provide an interpretation for why Figure 4.2 showed a significant correlation between latitude and income per-capita. It is accounted for by the correlation between latitude and the determinants of European colonization strategies. Europeans did not have immunity to tropical diseases during the colonial period and thus settler colonies tended, other things equal, to be created in temperate latitudes. Thus the historical creation of economic institutions was correlated with latitude. Without considering the role of economic institutions, one would find a spurious relationship between latitude and income per capita. However, once economic institutions are properly controlled for, these relationships go away and there appears to be no causal effect of geography on prosperity today (though geography may have been important historically in shaping economic institutions).

4.4.5. Culture, Colonial Identity and Economic Development. One might think that culture may have played an important role in the colonial experience, since Europeans not only brought new institutions, but also their own "cultures". European culture might have affected the economic development of former European colonies through three different channels. First, as already mentioned above, cultures may be systematically related to the national identity of the colonizing power. For example, the British may have implanted a

"good" Anglo-Saxon culture into colonies such as Australia and the United States, while the Spanish may have condemned Latin America by endowing it with an Iberian culture. Second, Europeans may have had a culture, work ethic or set of beliefs that were conducive to prosperity. Finally, Europeans also brought different religions with different implications for prosperity. Many of these hypotheses have been suggested as explanations for why Latin America, with its Roman Catholic religion and Iberian culture, is poor relative to the Anglo-Saxon Protestant North America.

Yet, the econometric evidence in Acemoglu, Johnson and Robinson (2001) is not consistent with any of these views either. Similar to the evidence related to geographical variables, the econometric strategy discussed above suggests that, once the effect of economic institutions is taken into account, neither the identity of the colonial power, nor the contemporary fraction of Europeans in the population, nor the proportions of the populations of various religions appear to have a direct effect on economic growth and income per capita.

These econometric results are supported by historical examples. Although no Spanish colony has been as successful economically as British colonies such as the United States, many former British colonies, such as those in Africa, India and Bangladesh, are poor today. It is also clear that the British in no way simply re-created British institutions in their colonies. For example, by 1619 the North American colony of Virginia had a representative assembly with universal male suffrage, something that did not arrive in Britain itself until 1919. Another telling example is that of the Puritan colony in Providence Island in the Caribbean. While the Puritan values are often credited with the arrival of democracy and equality of opportunity in Northeastern United States, the Puritan colony in Providence Island quickly became just like any other Caribbean slave colony despite its Puritanical inheritance.

Similarly, even though the 17th century Dutch had perhaps the best domestic economic institutions in the world, their colonies in South-East Asia ended up with institutions designed for the extraction of resources, providing little economic or civil rights to the indigenous population. These colonies consequently experienced slow growth relative to other countries.

Overall, the evidence does not appear to be consistent with a major role of geography, religion or culture transmitted by the identity of the colonizer or the presence of Europeans. Instead, differences in economic institutions appear to be the robust causal factor underlying the differences in income per capita across countries. Institutions therefore appear to be the most important fundamental cause of income differences and long-run growth.

4.5. What Types of Institutions?

As already noted above the notion of institutions used in this chapter and in much of the literature is rather broad. It encompasses different types of social arrangements, laws, regulations, enforcement of property rights and so on. One may, perhaps rightly, complain that we are learning relatively little by emphasizing the importance of such a broad cluster of institutions. It is therefore important to try to understand what types of institutions are more important. This will not only be useful in our empirical analysis of fundamental

causes, but can provide us a better sense of what types of models to develop in order to link fundamental causes to growth mechanics and to ultimate economic outcomes.

There is relatively little work on "unbundling" the broad cluster of institutions in order to understand what specific types of institutions might be more important for economic outcomes. Much of this type of work remains to be done in the future. Here it is useful to briefly mention some recent existing work attempting to distinguish the impact of "contracting institutions" from the influence of "property rights institutions". One of the important roles of institutions is to facilitate contracting between lenders and borrowers or between different firms, to facilitate the functioning of markets and the allocation of resources. Such contracting is only possible if laws, courts and regulations uphold contracts in the appropriate way. Let us refer to institutional arrangements of this sort that support private contracts as *contracting institutions*. The other cluster of institutions emphasized above relates to those constraining government and elite expropriation. Let us refer to these as *property rights institutions*. Although in many situations contracting institutions and property right institutions will be intimately linked, they are nonetheless conceptually different. While contracting institutions regulate "horizontal" relationships in society between regular citizens, property rights institutions are about the protection of citizens against the power of elites, politicians and privileged groups. These two sets of institutions are potentially distinct and can thus have distinct effects.

Acemoglu and Johnson (2005) investigate the relative roles of these two sets of institutions. Their strategy is again to make use of the natural experiment of colonial history. What helps this particular unbundling exercise is that in the sample of former European colonies, the legal system imposed by colonial powers appears to have a strong effect on contracting institutions, but little impact on the available measures of property rights institutions. At the same time, both mortality rates for potential European settlers and population density in 1500 have a large effect on current property rights institutions, and no impact on contracting institutions. Using these different sources of variation in the sample of former European colonies, it is possible to estimate the separate effects of contracting institutions and property rights institutions.

The empirical evidence estimating the different sources of variation in colonial history finds that property rights institutions are much more important for current economic outcomes than contracting institutions. Countries with greater constraints on politicians and elites and more protection against expropriation by these powerful groups appear to have substantially higher long-run growth rates and higher levels of current income. They also have significantly greater investment levels and generate more credit for the private sector. In contrast, the role of contracting institutions is more limited. Once the effects of property rights institutions are controlled for, contracting institutions seem to have no impact on income per capita, the investment to GDP ratio, and the private credit to GDP ratio. Contracting institutions appear to have some effect on stock market development, however.

These results suggest that contracting institutions affect the form of financial intermediation, but have less impact on economic growth and investment. It seems that economies can function in the face of weak contracting institutions without disastrous consequences, but not in the presence of a significant risk of expropriation from the government or other powerful groups. A possible interpretation is that private contracts or other reputation-based mechanisms can, at least in part, alleviate the problems originating from weak contracting institutions. For example, when it is more difficult for lenders to collect on their loans, interest rates increase, banks that can monitor effectively play a more important role, or reputation-based credit relationships may emerge. In contrast, property rights institutions relate to the relationship between the state and citizens. When there are no checks on the state, on politicians, and on elites, private citizens do not have the security of property rights necessary for investment.

Nevertheless, in interpreting the evidence in Acemoglu and Johnson (2005), one should also bear in mind that the sources of variation in income per capita and investment rates identifying the different effects of contracting and property rights institutions relate to very large differences discussed in Chapter 1. It is possible that contracting institutions have relatively small affects, so that they are hard to detect when we look at countries with thirty-fold differences in income per capita. Therefore, this evidence should be interpreted as suggesting that contracting institutions are less important in generating the very large differences in economic development than the property rights institutions, not necessarily as suggesting that contracting institutions do not matter for economic outcomes.

4.6. Disease and Development

The evidence presented above already militates against a major role of geographic factors in economic development. One version of the geography hypothesis deserves further analysis, however. A variety of evidence suggests that unhealthy individuals are less productive and often less successful in acquiring human capital. Could the differences in the disease environments across countries have an important effect on economic development? Could they be a major factor in explaining the very large income differences across countries? A recent paper by David Weil (2006), for example, argues that the framework used in the previous chapter, with physical capital, human capital and technology, should be augmented by including health capital. In other words, we may want to think over production function of the form $F(K, H, Z, A)$, where H denotes efficiency units of labor (human capital as conventionally measured), while Z is "health capital". Weil suggests a methodology for measuring the contribution of health capital to productivity from micro estimates and argues that differences in health capital emerge as an important factor in accounting for cross-country differences in income levels.

The idea that part of the low productivity of less-developed nations is due to the unhealthy state of their workforces has obvious appeal. The micro evidence and the work by David Weil shows that it has some empirical validity as well. But does it imply that geographic factors are

an important fundamental cause of economic growth? Not necessarily. As already mentioned above, the burden of disease is endogenous. Today's unhealthy countries are unhealthy precisely because they are poor and are unable to invest in health care, clean water and other health-improving technologies. After all, much of Europe was very unhealthy and suffering from low life expectancy only 200 years ago. This changed *with* economic growth. In this sense, even if "health capital" is a useful concept and does contribute to accounting for cross-country income differences, it may itself be a proximate cause, affected by other factors, such as institutions or culture.

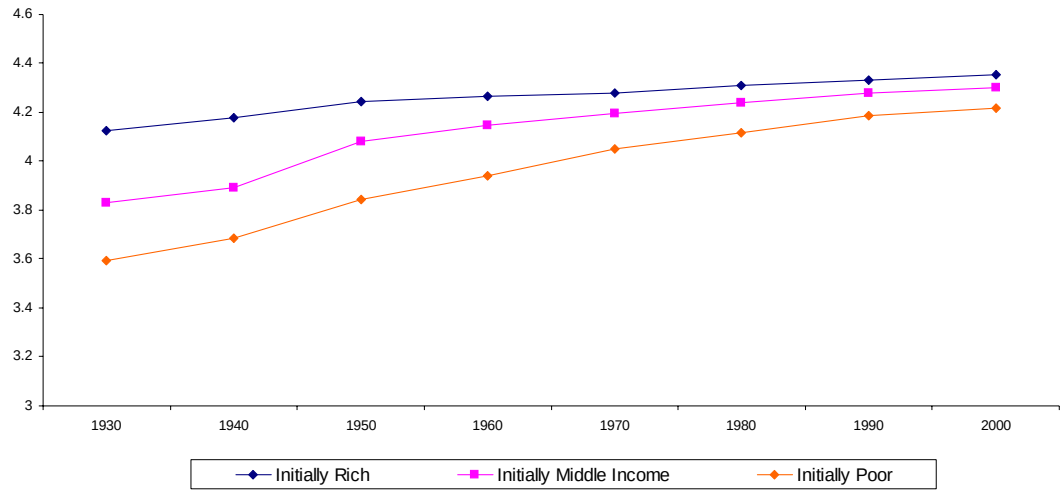

Figure 4.10. Evolution of life expectancy at birth among initially-poor, initially-middle-income and initially-rich countries, 1940-2000.

A recent paper by Acemoglu and Johnson (2007) directly investigates the impact of changes in disease burdens on economic development. They exploit the large improvements in life expectancy, particularly among the relatively poor nations, that took place starting in the 1940s. These health improvements were the direct consequence of significant international health interventions, more effective public health measures, and the introduction of new chemicals and drugs. More important for the purposes of understanding the effect of disease on economic growth, these health improvements were by and large exogenous from the viewpoint of individual nations. Moreover, their impact on specific nations also varied, depending on whether the country in question was affected by the specific diseases for which the cures and the drugs became internationally available. The impact of these health improvements was major, in fact so major that it may deserve to be called the *international epidemiological transition*, since it led to an unprecedented improvement in life expectancy in a large number of countries. Figure 4.10 shows this unprecedented convergence in life

expectancy by plotting life expectancy in countries that were initially (circa 1940) poor, middle income, and rich. It illustrates that while in the 1930s life expectancy was low in many poor and middle-income countries, this transition brought their levels of life expectancy close to those prevailing in richer parts of the world. As a consequence of these developments, health conditions in many parts of the less-developed world today, though still in dire need of improvement, are significantly better than the corresponding health conditions were in the West at the same stage of development.

The international epidemiological transition allows a promising empirical strategy to isolate potentially-exogenous changes in health conditions. The effects of the international epidemiological transition on a country's life expectancy were related to the extent to which its population was initially (circa 1940) affected by various specific diseases, for example, tuberculosis, malaria, and pneumonia, and to the timing of the various health interventions. This reasoning suggests that potentially-exogenous variation in the health conditions of the country can be measured by calculating a measure of predicted mortality, driven by the interaction of baseline cross-country disease prevalence with global intervention dates for specific diseases. Acemoglu and Johnson (2007) show that such measures of predicted mortality have a large and robust effect on changes in life expectancy starting in 1940, but have *no* effect on changes in life expectancy *prior* to this date (that is, before the key interventions). This suggests that the large increases in life expectancy experienced by many countries after 1940 were in fact related to the global health interventions.

Perhaps not surprisingly, Acemoglu and Johnson (2007) find that predicted mortality and the changes in life expectancy that it causes have a fairly large effect on population; a 1% increase in life expectancy is related to an approximately 1.3-1.8% increase in population. However, they find no evidence of a positive effect on GDP per capita. Figure 4.11 provides an aggregated version of this evidence. It shows no convergence in income per capita between initially-poor, initially-middle-income and initially-rich countries.

Why did the very significant increases in life expectancy and health not cause improvements in GDP per capita? The most natural answer to this question comes from neoclassical growth theory (presented in the previous two chapters and in Chapter 8 below). The first-order effect of increased life expectancy is to increase population, which initially reduces capital-to-labor and land-to-labor ratios, depressing income per capita. This initial decline is later compensated by higher output as more people enter the labor force. However, there is no reason to expect either a complete offset of the initial decline in income per capita or a large significant increase, especially when many of the effect countries are heavily vested in agriculture, so that land-to-labor ratios may change permanently. Consequently, small beneficial effects of health on productivity may not be sufficient to offset or reverse the negative effects of population pressures on income per capita.

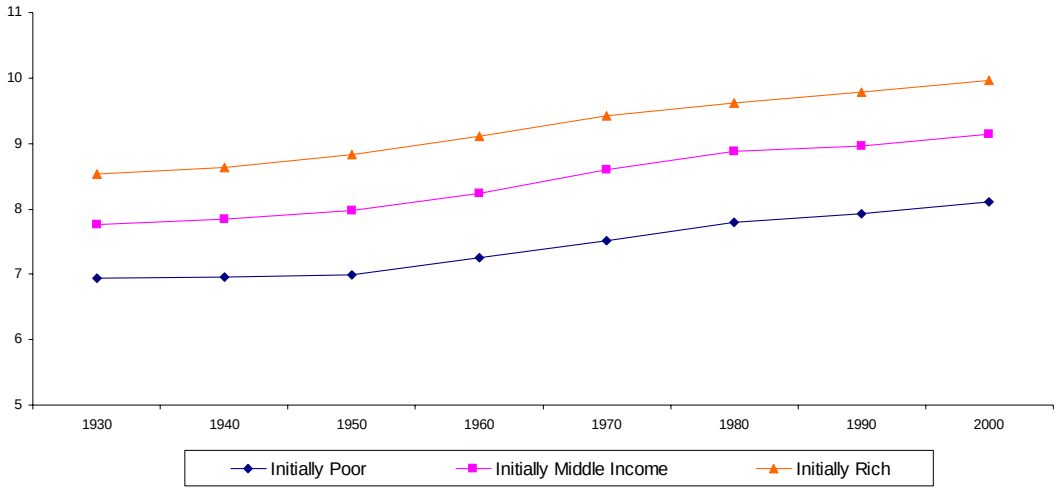

Figure 4.11. Evolution of GDP per capita among initially-poor, initially-middle-income and initially-rich countries, 1940-2000.

4.7. Political Economy of Institutions: First Thoughts

The evidence presented in this chapter suggests that institutions are the most important fundamental cause of economic growth. We must therefore think about why institutions and policies differ across countries in order to understand why some countries are poor and some are rich. I will also argue below that understanding institutional changes holds clues about why the process of world economic growth started 200 years ago or so.

However, an explanation of differences in income across countries and over time in terms of institutional differences is also incomplete. If, as this chapter has documented, some institutions are conducive to rapid economic growth and others to stagnation, why would any society collectively choose institutions that condemn them to stagnation? The answer to this question relates to the nature of collective choices in societies. Institutions and policies, like other collective choices, are not taken for the good of the society at large, but are a result of a political equilibrium. In order to understand such political equilibria, we need to understand the conflicting interests of different individuals and groups in societies, and analyze how they will be mediated by different political institutions. Thus, a proper understanding of how institutions affect economic outcomes and why institutions differ across countries (and why they sometimes change and pave the way for growth miracles) requires models of *political economy*, which explicitly studies how the conflicting interests of different individuals are aggregated into collective choices. Models of political economy also specify why certain individuals and groups may be opposed to economic growth or prefer institutions that eschew growth opportunities.

The discussion in this chapter therefore justifies why a study of political economy has to be part of any detailed investigation of economic growth. Much of the study of economic growth has to be about the structure of models, so that we understand the mechanics of economic growth and the proximate causes of income differences. But part of this broad study must also confront the fundamental causes of economic growth, which relate to policies, institutions and other factors that lead to different investment, accumulation and innovation decisions.

4.8. Taking Stock

This chapter has emphasized the differences between the proximate causes of economic growth, related to physical capital accumulation, human capital and technology, and the fundamental causes, which influence the incentives to invest in these factors of production. I have argued that many of the questions motivating our study of economic growth must lead us to an investigation of the fundamental causes. But an understanding of fundamental causes is most useful when we can link them to the parameters of fully-worked-out model of economic growth to see how they affect the mechanics of growth and what types of predictions they generate.

The institutions hypothesis, which seems to receive support from the evidence presented in the current chapter, calls for a careful theoretical investigation. The institutions view makes sense only when there are groups in society that favor institutions that do not necessarily enhance the growth potential of the economy. They will do so because they will not directly or indirectly benefit from the process of economic growth. Thus it is important to develop a good understanding of the distributional implications of economic growth (for example, how it affects relative prices and relative incomes, and how it may destroy the rents of incumbents). This theoretical understanding of the implications of the growth process then needs to be combined with political economy models of collective decision-making, to investigate under what circumstances groups opposed to economic growth can be powerful enough to maintain non-growth-enhancing institutions in place.

In this chapter, my objective was more limited (since many of the more interesting growth models will be developed later in the book) and I have focused on the broad outlines of a number of alternative fundamental causes of economic growth and had a first look at the long-run empirical evidence relevance to these hypotheses. I argued that approaches emphasizing institutional differences (and differences in policies, laws and regulations) across societies are most promising for understanding both the current growth experiences of countries and the historical process of economic growth. I also emphasized the importance of studying the political economy of institutions, as a way of understanding why institutions differ across societies and lead to divergent economic paths.

4.9. References and Literature

The early part of this chapter builds on Acemoglu, Johnson and Robinson (2006), who discuss the distinction between proximate and fundamental causes and the various different

approaches to the fundamental causes of economic growth. North and Thomas (1973) appear to be the first to implicitly criticize growth theory for focusing on the proximate causes alone and ignoring the fundamental causes of economic growth. Diamond (1997) also draws a distinction between proximate and fundamental explanations.

The importance of population in generating economies of scale was first articulated by Julian Simon (1990). The model presented in Section 4.2 draws on Simon's work and the more recent work by Michael Kremer (1993). Kremer (1993) argues for the importance of economies of scale and increasing returns to population based on the acceleration in the growth rate of world population. Another important argument relating population to technological change is proposed by Esther Boserup (1965) and is based on the idea that increases in population creates scarcity, inducing societies to increase their productivity. Other models that build economies of scale to population and discuss the transition of the world economy from little or no growth to one of rapid economic growth include Hansen and Prescott (2001), Galor and Weil (2001), Galor and Moav (2002) and Jones (2004). Some of these papers also try to reconcile the role of population in generating technological progress with the later demographic transition. Galor (2006) provides an excellent summary of this literature and an extensive discussion. McEvedy and Jones (1978) provide a concise history of world population and relatively reliable information going back to 10,000 B.C. Their data indicate that, as claimed in the text, total population in Asia has been consistently greater than in Western Europe over this time period.

The geography hypothesis has many proponents. In addition to Montesquieu, Machiavelli was an early proponent of the importance of climate and geographic characteristics. Marshall (1890), Kamarck (1976), and Myrdal (1986) are among the economists who have most clearly articulated various different versions of the geography hypothesis. It has more recently been popularized by Sachs (2000, 2001), Bloom and Sachs (1998) and Gallup and Sachs (2001). Diamond (1997) offers a more sophisticated version of the geography hypothesis, where the availability of different types of crops and animals, as well as the axes of communication within continents, influence the timing of settled agriculture and thus the possibility of developing complex societies. Diamond's thesis is therefore based on geographic differences, but also relies on such institutional factors as intervening variables.

Scholars emphasizing the importance of various types of institutions in economic development include John Locke, Adam Smith, John Stuart Mill, Arthur Lewis, Douglass North and Robert Thomas. The recent economics literature includes many models highlighting the importance of property rights, for example, Skaperdas (1992), Tornell and Velasco (1992), Acemoglu (1995), Grossman and Kim (1995, 1996), Hirsleifer (2001) and Dixit (2004). Other models emphasize the importance of policies within a given institutional framework. Well-known examples of this approach include Perotti (1993), Saint-Paul and Verdier (1993), Alesina and Rodrik (1994), Persson and Tabellini (1994), Ades and Verdier (1996), Krusell and Rios-Rull (1999), and Bourguignon and Verdier (2000). There is a much smaller literature on endogenous institutions and the effect of these institutions on economic outcomes.

Surveys of this work can be found in Acemoglu (2007) and Acemoglu and Robinson (2006). The literature on the effect of economic institutions on economic growth is summarized and discussed in greater detail in Acemoglu, Johnson and Robinson (2006), which also provides an overview of the empirical literature on the topic. I return to many of these issues and Part 8 of the book.

The importance of religion for economic development is most forcefully argued in Max Weber's work, for example (1930, 1958). Many other scholars since then have picked up on this idea and have argued about the importance of religion. Prominent examples include Huntington (2001) and Landes (2001). Landes, for example, tries to explain the rise of the West based on cultural and religious variables. This evidence is criticized in Acemoglu, Johnson and Robinson (2005). Barro and McCleary (2003) provide evidence of a positive correlation between the prevalence of religious beliefs and economic growth. One has to be careful in interpreting this evidence as showing a causal effect of religion on economic growth, since religious beliefs are endogenous both to economic outcomes and to other fundamental causes of income differences.

The emphasis on the importance of cultural factors or "social capital" goes back to Banfield (1958), and has recently been popularized by Putnam (1993). The essence of these interpretations appears to be related to the role of culture or social capital in ensuring the selection of better equilibrium. Similar ideas are also advanced in Greif (2006). Many scholars, including Véliz (1994), North, Summerhill and Weingast (2000) and Wiarda (2001), emphasize the importance of cultural factors in explaining the economic backwardness of Latin American countries. Knack and Keefer (1997) and Durlauf and Fafchamps (2003) document positive correlations between measures of social capital and various economic outcomes. None of this work establishes a causal effect of social capital because of the potential endogeneity of social capital and culture. A number of recent papers attempt to overcome these difficulties. Notable contributions here include Guiso, Sapienza and Zingales (2004) and Tabellini (2007).

The discussion of the Puritan colony in the Providence Island is based on Newton (1914) and Kupperman (1993).

The literature on the effect of economic institutions and policies on economic growth is vast. Most growth regressions include some controls for institutions or policies and find them to be significant (see, for example, those reported in Barro and Sala-i-Martin, 2004). One of the first papers looking at the cross-country correlation between property rights measures and economic growth is Knack and Keefer (1995). This literature does not establish causal effect either, since simultaneity and endogeneity concerns are not dealt with. Mauro (1998) and Hall and Jones (1999) present the first instrumental-variable estimates on the effect of institutions (or corruption) on long-run economic development.

The evidence reported here, which exploits differences in colonial experience to create an instrumental-variables strategy, is based on Acemoglu, Johnson and Robinson (2001, 2002). The urbanization and population density data used here are from Acemoglu, Johnson and Robinson (2002), which compiled these based on work by Bairoch (1988), Bairoch, Batou

and Chèvre (1988), Chandler (1987), Eggimann (1999), McEvedy and Jones (1978). Further details and econometric results are presented in Acemoglu, Johnson and Robinson (2002). The data on mortality rates of potential settlers is from Acemoglu, Johnson and Robinson (2001), who compiled the data based on work by Curtin (1989, 1998) and Gutierrez (1986). That paper also provides a large number of robustness checks, documenting the influence of economic institutions on economic growth and showing that other factors, including religion and geography, have little effect on long-run economic development once the effect of institutions is controlled for.

The details of the Korean experiment and historical references are provided in Acemoglu (2003) and Acemoglu, Johnson and Robinson (2006).

The discussion of distinguishing the effects of different types of institutions draws on Acemoglu and Johnson (2005).

The discussion of the effect of disease on development is based on Weil (2006) and especially on Acemoglu and Johnson (2007), which used the econometric strategy described in the text. Figures 4.10 and 4.11 are from Acemoglu and Johnson (2007). In these figures, initially-poor countries are those that are poorer than Spain in 1940, and include China, Bangladesh, India, Pakistan, Myanmar, Thailand, El Salvador, Honduras, Indonesia, Brazil, Sri Lanka, Malaysia, Nicaragua, Korea, Ecuador, and the Philippines, while initially-rich countries are those that are richer than Argentina in 1940 and include Belgium, Netherlands, Sweden, Denmark, Canada, Germany, Australia, New Zealand, Switzerland, the United Kingdom and the United States. Young (2004) investigates the effect of the HIV epidemic in South Africa and reaches a conclusion similar to that reported here, though his analysis relies on a calibration of the neoclassical growth model rather than econometric estimation.

4.10. Exercises

Exercise 4.1. Derive eq.'s (4.3) and (4.4).

Exercise 4.2. Derive eq. (4.7). Explain how and why the behavior implied for technology by this equation differs from (4.4). Do you find the assumptions leading to (4.4) or to (4.7) more plausible?

Exercise 4.3.

(1) Show that the models leading to both (4.4) or to (4.7) imply a constant income per capita throughout.

(2) Modify eq. (4.2) to

$$L(t) = \phi Y(t)^{\beta},$$

for some $\beta \in (0,1)$. Justify this equation and derive the law of motion technology and income per capita under the two scenarios considered in the text. Are the implications of this model more reasonable than those considered in the text?

Exercise 4.4. In his paper "Tropical Underdevelopment", Jeff Sachs notes that differences in income per capita between tropical and temperate zones have widened over the past 150 years. He interprets this pattern as evidence indicating that the "geographical burden" of the tropical areas has been getting worse over the process of recent development. Discuss this

thesis. If you wish, offer alternative explanations. How would you go about testing different approaches? (If possible, suggest original ways, rather than approaches that were already tried).

Part 2

Towards Neoclassical Growth

This part of the book is a preparation for what is going to come next. In some sense, it can be viewed as the "preliminaries" for the rest of the book. The ultimate purpose is to enrich the basic Solow model by introducing well-defined consumer preferences and consumer optimization, and in the process, clarify the relationship between growth theory and general equilibrium theory. This will enable us to open the blackbox of savings and capital accumulation, turning these decisions into forward-looking investment decisions. It will also enable us to make welfare statements about whether the rate of growth of an economy is too slow, too fast or just right from a welfare-maximizing (Pareto optimality) viewpoint. This will then open the way for us to study technology as another forward-looking investment by firms, researchers and individuals. However, much of this will have to wait for Parts 3 and 4 of the book. In the next three chapters, present the material necessary to appreciate what is to come then. The next chapter will set up the problem and make the relationship between models of economic growth and general equilibrium theory more explicit. It will also highlight some of the assumptions implicit in the growth models. The two subsequent chapters develop the mathematical tools for dynamic optimization in discrete and continuous time. To avoid making these chapters purely about mathematics, I will use a variety of economic models of some relevance to growth theory as examples and also include the analysis of the equilibrium and optimal growth. Nevertheless, the material in the next three chapters is significantly more mathematical than anything else in the book (except the Appendices at the end). The reader therefore may wish to first consults the Appendices and/or skip some of the proofs presented in the next three chapters in a first reading.

CHAPTER 5

Foundations of Neoclassical Growth

The Solow growth model is predicated on a constant saving rate. Instead, it would be much more satisfactory to specify the *preference orderings* of individuals, as in standard general equilibrium theory, and derive their decisions from these preferences. This will enable us both to have a better understanding of the factors that affect savings decisions and also to discuss the "optimality" of equilibria—in other words, to pose and answer questions related to whether the (competitive) equilibria of growth models can be "improved upon". The notion of improvement here will be based on the standard concept of Pareto optimality, which asks whether some households can be made better-off without others being made worse-off. Naturally, we can only talk of individuals or households being "better-off" if we have some information about well-defined preference orderings.

5.1. Preliminaries

To prepare for this analysis, let us consider an economy consisting of a unit measure of infinitely-lived households. By a unit measure of households I mean an uncountable number of households, for example, the set of households $\mathcal{H}$ could be represented by the unit interval $[0,1]$. This is an abstraction adopted for simplicity, to emphasize that each household is infinitesimal and will have no effect on aggregates. Nothing in this book hinges on this assumption. If the reader instead finds it more convenient to think of the set of households, $\mathcal{H}$, as a countable set of the form $\mathcal{H} = \{1, 2, ..., M\}$ or $\mathcal{H} = \mathbb{N}$, this can be done without any loss of generality. The advantage of having a unit measure of households is that averages and aggregates are the same, enabling us to economize on notation. It would be even simpler to have $\mathcal{H}$ as a finite set in the form $\{1, 2, ..., M\}$. While this would sufficient in many contexts, overlapping generations models in Chapter 9 require the set of households to be infinite.

Households in this economy may be truly "infinitely lived," or alternatively they may consist of overlapping generations with full (or partial) altruism linking generations within the household. Throughout, I equate households with individuals, and thus ignore all possible sources of conflict or different preferences within the household. In other words, I assume that households have well-defined preference orderings.

As in basic general equilibrium theory, let us suppose that preference orderings can be represented by utility functions. In particular, suppose that each household h has an *instantaneous utility function* given by

$$u_h\left(c_h\left(t\right)\right),$$

where $u_h : \mathbb{R}_+ \to \mathbb{R}$ is increasing and concave and $c_h(t)$ is the consumption of household h. I take the domain of the utility function to be $\mathbb{R}_+$ rather than $\mathbb{R}$, so that negative levels of consumption are not allowed. Even though some well-known economic models allow negative consumption, this is not easy to interpret in general equilibrium or in growth theory, thus this restriction is sensible in most models.

The instantaneous utility function captures the utility that an individual (or household) derives from consumption at time t. It is therefore *not* the same as a utility function specifying a complete preference ordering over all commodities—here consumption levels in all dates. For this reason, the instantaneous utility function is sometimes also referred to as the "felicity function".

There are two major assumptions in writing an instantaneous utility function. First, it imposes that the household does not derive any utility from the consumption of other households, so consumption externalities are ruled out. Second, in writing the instantaneous utility function, I already imposed that overall utility is *time-separable*, that is, instantaneous utility at time t is independent of the consumption levels at past or future dates. This second feature is important in enabling us to develop tractable models of dynamic optimization.

Finally, let us introduce a third assumption and suppose that households discount the future "exponentially"—or "proportionally". In discrete time, and ignoring uncertainty, this implies that household preferences at time $t = 0$ can be represented as

$$\sum_{t=0}^{\infty} \beta_h^t u_h\left(c_h(t)\right), \tag{5.1}$$

where $\beta_h \in (0,1)$ is the discount factor of household h. This functional form implies that the weight given to tomorrow's utility is a fraction β_h of today's utility, and the weight given to the utility the day after tomorrow is a fraction β_h^2 of today's utility, and so on. Exponential discounting and time separability are convenient for us because they naturally ensure "time-consistent" behavior.

A solution $\{x(t)\}_{t=0}^{T}$ (possibly with $T = \infty$) to a dynamic optimization problem is *time-consistent* if the following is true: whenever $\{x(t)\}_{t=0}^{T}$ is an optimal solution starting at time $t = 0$, $\{x(t)\}_{t=t'}^{T}$ is an optimal solution to the continuation dynamic optimization problem starting from time $t = t' \in [0, T]$. If a problem is not time-consistent, it is *time-inconsistent*. Time-consistent problems are much more straightforward to work with and satisfy all of the standard axioms of rational decision-making. Although time-inconsistent preferences may be useful in the modeling of certain behaviors, such as problems of addiction or self-control, time-consistent preferences are ideal for the focus in this book, since they are tractable, relatively flexible and provide a good approximation to reality in the context of aggregative models. It is also worth noting that many classes of preferences that do not feature exponential and time-separable discounting nonetheless lead to time-consistent behavior. Exercise 5.1 discusses issues of time-consistency further and shows how certain other types of utility formulations lead to time-inconsistent behavior, while Exercise 5.2 introduces some common non-time-separable preferences that lead to time-consistent behavior.

There is a natural analog to (5.1) in continuous time, again incorporating exponential discounting, which is introduced and discussed below (see Chapter 7).

The expression in (5.1) ignores uncertainty in the sense that it assumes the sequence of consumption levels for household h, $\{c_h(t)\}_{t=0}^{\infty}$ is known with certainty. If instead this sequence were uncertain, we would need to look at expected utility maximization. Most growth models do not necessitate an analysis of behavior under uncertainty, but a stochastic version of the neoclassical growth model is the workhorse of much of the rest of modern macroeconomics and will be presented in Chapter 17. For now, it suffices to say that in the presence of uncertainty, $u_h(\cdot)$ should be interpreted as a "*Bernoulli utility function*," defined over risky acts, so that the preferences of household h at time $t = 0$ can be represented by the following *von Neumann-Morgenstern* expected utility function:

$$\mathbb{E}_0^h \sum_{t=0}^{\infty} \beta_h^t u_h(c_h(t)),$$

where $\mathbb{E}_0^h$ is the expectation operator with respect to the information set available to household h at time $t = 0$.

The formulation so far indexes individual utility function, $u_h(\cdot)$, and the discount factor, β_h, by "h" to emphasize that these preference parameters are potentially different across households. Households could also differ according to their income processes. For example, each household could have effective labor endowments of $\{e_h(t)\}_{t=0}^{\infty}$, thus a sequence of labor income of $\{e_h(t)\, w(t)\}_{t=0}^{\infty}$, where $w(t)$ is the equilibrium wage rate per unit of effective labor.

Unfortunately, at this level of generality, this problem is not tractable. Even though we can establish some existence of equilibrium results, it would be impossible to go beyond that. Proving the existence of equilibrium in this class of models is of some interest, but our focus is on developing workable models of economic growth that generate insights about the process of growth over time and cross-country income differences. I will therefore follow the standard approach in macroeconomics and assume the existence of a *representative household*.

5.2. The Representative Household

When we say that an economy *admits a representative household*, this means that the preference (demand) side of the economy can be represented *as if* there were a single household making the aggregate consumption and saving decisions (and also the labor supply decisions when these are endogenized) subject to an aggregate budget constraint. The major convenience of the representative household assumption is that rather than modeling of the preference side of the economy resulting from equilibrium interactions of many heterogeneous households, it allows us to model it as a solution to a single maximization problem. Note that, for now, the description concerning a representative household is purely positive—it asks the question of whether the aggregate behavior can be represented *as if* it were generated by a single household. A stronger notion, the "normative" representative household, would also allow to use the utility function of the normative representative household for

welfare comparisons, but is valid under much more stringent conditions. I return to a further discussion of these issues below.

Let us start with the simplest case that will lead to the existence of a representative household. Suppose that each household is identical, that is, each has the same discount factor β, the same sequence of effective labor endowments $\{e(t)\}_{t=0}^{\infty}$ and the same instantaneous utility function

$$u(c_h(t))$$

where $u : \mathbb{R}_+ \to \mathbb{R}$ is increasing and concave and $c_h(t)$ is the consumption of household h. Therefore, there is really a representative household in this case. Consequently, again ignoring uncertainty, the preference side of the economy can be represented as the solution to the following maximization problem starting at time $t = 0$:

$$\max \sum_{t=0}^{\infty} \beta^t u(c(t)), \tag{5.2}$$

where $\beta \in (0,1)$ is the common discount factor of all the households, and $c(t)$ is the consumption level of the representative household.

The economy described so far admits a representative household rather trivially; all households are identical. In this case, the representative household's preferences, (5.2), can be used not only for positive analysis (for example, to determine what the level of savings will be), but also for normative analysis, such as evaluating the optimality of equilibria.[1]

Often, we may not want to assume that the economy is indeed inhabited by a set of identical households, but instead assume that the behavior of the households can be modeled *as if* it were generated by the optimization decision of a representative household. Naturally, this would be more realistic than assuming that all households are identical. Nevertheless, this is not without any costs. First, in this case, the representative household will have positive meaning, but not always a normative meaning (see below). Second, it is not in fact true that most models with heterogeneity lead to a behavior that can be represented as if it were generated by a representative household.

In fact most models do not admit a representative household. To illustrate this, let us consider a simple exchange economy with a finite number of commodities and state an important theorem from general equilibrium theory. In preparation for this theorem, recall that in an exchange economy, we can think of the object of interest as the excess demand functions (or correspondences) for different commodities. Let these be denoted by $\mathbf{x}(p)$ when the vector of prices is p. The demand side of an economy will admit a representative household if these excess demands, $\mathbf{x}(p)$, can be modeled as if they result from the maximization problem of a single consumer.

[1]This exercise would (implicitly) focus on optimal allocations in which all households have the same Pareto weights. One can also imagine allocations in which ex ante identical households receive different weights and different utility levels. Throughout, whenever the economy admits a representative household, I follow the standard practice of focusing on optima with equal Pareto weights.

Theorem 5.1. ***(Debreu-Mantel-Sonnenschein Theorem)*** *Let $\varepsilon > 0$ and $N \in \mathbb{N}$. Consider a set of prices $\mathbf{P}_\varepsilon = \left\{p \in \mathbb{R}^N_+ \text{: } p_j/p_{j'} \geq \varepsilon \text{ for all } j \text{ and } j'\right\}$ and any continuous function $\mathbf{x} : \mathbf{P}_\varepsilon \to \mathbb{R}^N_+$ that satisfies Walras's Law and is homogeneous of degree 0. Then, there exists an exchange economy with N commodities and $H < \infty$ households, where the aggregate demand is given by $\mathbf{x}(p)$ over the set $\mathbf{P}_\varepsilon$.*

Proof. See Debreu (1974) or Mas-Colell, Winston and Green (1995), Proposition 17.E.3. □

Therefore, the fact that excess demands result from aggregating the optimizing behavior of households places few restrictions on the form of these demands. In particular, the excess demand function $\mathbf{x}(p)$ does not necessarily possess a negative-semi-definite Jacobian or satisfy the weak axiom of revealed preference (which are requirements of demands generated by households). This implies that, without imposing further structure, it is impossible to derive the aggregate excess demand, $\mathbf{x}(p)$, from the maximization behavior of a single household. This theorem therefore raises a severe warning against the use of the representative household assumption.

Nevertheless, this result is partly an outcome of very strong income effects. Special but approximately realistic preference functions, as well as restrictions on the distribution of income across individuals, enable us to rule out arbitrary aggregate excess demand functions. To show that the representative household assumption is not as hopeless as Theorem 5.1 suggests, I now present a special and relevant case in which aggregation of individual preferences is possible and enables the modeling of the economy as if the demand side was generated by a representative household.

To prepare for this theorem, consider an economy with a finite number N of commodities and recall that an indirect utility function for household h, $v_h\left(p, y^h\right)$, specifies the household's (ordinal) utility as a function of the price vector $p = (p_1, ..., p_N)$ and the household's income y^h. Naturally, any indirect utility function $v_h\left(p, y^h\right)$ has to be homogeneous of degree 0 in p and y.

Theorem 5.2. ***(Gorman's Aggregation Theorem)*** *Consider an economy with a finite number $N < \infty$ of commodities and a set $\mathcal{H}$ of households. Suppose that the preferences of household $h \in \mathcal{H}$ can be represented by an indirect utility function of the form*

$$v^h\left(p, y^h\right) = a^h(p) + b(p)\, y^h, \tag{5.3}$$

then these preferences can be aggregated and represented by those of a representative household, with indirect utility

$$v(p, y) = a(p) + b(p)\, y,$$

where $a(p) \equiv \int_{h \in \mathcal{H}} a^h(p)\, dh$, and $y \equiv \int_{h \in \mathcal{H}} y^h dh$ is aggregate income.

Proof. See Exercise 5.3. □

This theorem implies that when preferences admit this special quasi-linear form, aggregate behavior can indeed be represented as if it resulted from the maximization of a single household. This class of preferences are referred to as *Gorman preferences* after Terrence Gorman, who was among the first economists studying issues of aggregation and proposed the special class of preferences used in Theorem 5.2. The quasi-linear structure of these preferences limits the extent of income effects and enables the aggregation of individual behavior. Notice that instead of the summation, this theorem used the integral over the set $\mathcal{H}$ to allow for the possibility that the set of households may be a continuum. The integral should be thought of as the "Lebesgue integral," so that when $\mathcal{H}$ is a finite or countable set, $\int_{h\in\mathcal{H}} y^h dh$ is indeed equivalent to the summation $\sum_{h\in\mathcal{H}} y^h$.[2] Note also that this theorem is stated for an economy with a finite number of commodities. This is only for simplicity, and the same result can be generalized to an economy with an infinite or even a continuum of commodities. However, for most of this chapter, I restrict attention to economies with either a finite or a countable number of commodities to simplify notation and avoid technical details.

Note also that for preferences to be represented by an indirect utility function of the Gorman form does not necessarily mean that this utility function will give exactly the indirect utility in (5.3). Since, in the absence of uncertainty, a *monotonic transformation* of the utility function has no effect on behavior (but affects the indirect utility function), all that is required in models without uncertainty is that there exists a monotonic transformation of the indirect utility function that takes the form given in (5.3).

Another attractive feature of Gorman preferences for our purposes is that they contain some commonly-used preferences in macroeconomics. To illustrate this, let us start with the following example:

EXAMPLE 5.1. **(Constant Elasticity of Substitution Preferences)** A very common class of preferences used in industrial organization and macroeconomics are the constant elasticity of substitution (CES) preferences, also referred to as Dixit-Stiglitz preferences after the two economists who first used these preferences. Suppose that each household denoted by $h \in \mathcal{H}$ has total income y^h and preferences defined over $j = 1, ..., N$ goods given by

$$U^h\left(x_1^h, ..., x_N^h\right) = \left[\sum_{j=1}^{N}\left(x_j^h - \xi_j^h\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{5.4}$$

where $\sigma \in (0, \infty)$ and $\xi_j^h \in [-\bar{\xi}, \bar{\xi}]$ is a household specific term, which parameterizes whether the particular good is a necessity for the household. For example, $\xi_j^h > 0$ may mean that household h needs to consume at least a certain amount of good j to survive. The utility function (5.4) is referred to as a CES function for the following reason: define the level of

[2]Throughout the book I avoid the use of measure theory whenever I can, but I will refer to the "Lebesgue integral" a number of times. This is a generalization of the standard Riemann integral. The few references to Lebesgue integrals simply signal that we should think of integration in a context slightly more generally, so that the integrals could be representing expectations or averages even with discrete random variables and discrete distributions (or mixtures of continuous and discrete distributions). References to some introductory treatments of measure theory and the Lebesgue integral are provided at the end of Chapter 16 below.

consumption of each good as $\hat{x}_j^h = x_j^h - \xi_j^h$, the elasticity of substitution between any two $\hat{x}_j^h$ and $\hat{x}_{j'}^h$ (with $j \neq j'$) would be equal to σ.

Each consumer faces a vector of prices $p = (p_1, ..., p_N)$, and we assume that for all h,

$$\sum_{j=1}^{N} p_j \bar{\xi} < y^h,$$

so that the household can afford a bundle such that $\hat{x}_j^h \geq 0$ for all j. In Exercise 5.6, you will be asked to derive the optimal consumption levels for each household and show that their indirect utility function is given by

$$v^h\left(p, y^h\right) = \frac{\left[-\sum_{j=1}^{N} p_j \xi_j^h + y^h\right]}{\left[\sum_{j=1}^{N} p_j^{1-\sigma}\right]^{\frac{1}{1-\sigma}}}, \tag{5.5}$$

which satisfies the Gorman form (and is also homogeneous of degree 0 in p and y). Therefore, this economy admits a representative household with indirect utility:

$$v\left(p, y\right) = \frac{\left[-\sum_{j=1}^{N} p_j \xi_j + y\right]}{\left[\sum_{j=1}^{N} p_j^{1-\sigma}\right]^{\frac{1}{1-\sigma}}},$$

where y is aggregate income given by $y \equiv \int_{h \in \mathcal{H}} y^h dh$ and $\xi_j \equiv \int_{h \in \mathcal{H}} \xi_j^h dh$. It is also straightforward to verify that the utility function leading to this indirect utility function is

$$U\left(x_1, ..., x_N\right) = \left[\sum_{j=1}^{N} \left(x_j - \xi_j\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}. \tag{5.6}$$

Preferences closely related to the CES preferences presented in this example will play a special role not only in aggregation but also in ensuring *balanced growth* in neoclassical growth models.

The converse to Theorem 5.2 is also straightforward; Gorman preferences are not only sufficient for the economy to admit a representative household, but they are also *necessary*. In particular, if preferences do not take the Gorman form (with the same $b(p)$ for all households), then there will exist a specific redistribution of income that will change demands and thus the economy cannot admit a representative household.[3] Since this result is not central to our focus, I will not state it formally and provide a proof.

In addition to the aggregation result in Theorem 5.2, Gorman preferences also imply the existence of a *normative representative household*, in the sense that not only there exists a representative household whose maximization problem will generate the relevant aggregate demands, but also the utility function of this household can be used for welfare analysis. More

[3]Naturally, we can obtain a broader class of preferences for which aggregate behavior can be represented as if it resulted from the maximization problem of a single representative household if we are willing to restrict attention to specific distribution of income (wealth) across households.

specifically, recall that an allocation is *Pareto optimal* if no household can be made strictly better-off without some other household being made worse-off (see Definition 5.2 below).

Theorem 5.3. ***(Existence of a Normative Representative Household)*** *Consider an economy with a finite number $N < \infty$ of commodities and a set $\mathcal{H}$ of households. Suppose that the preferences of each household $h \in \mathcal{H}$ take the Gorman form, $v^h\left(p, y^h\right) = a^h(p) + b(p)\, y^h$, where $p = (p_1, ..., p_N)$ is the price vector.*

1. *Then, any allocation that maximizes the utility of the representative household, $v(p, y) = \sum_{h\in\mathcal{H}} a^h(p) + b(p)\, y$, with $y \equiv \sum_{h\in\mathcal{H}} y^h$, is Pareto optimal.*
2. *Moreover, if $a^h(p) = a^h$ for all p and all $h \in \mathcal{H}$, then any Pareto optimal allocation maximizes the utility of the representative household.*

Proof. I prove this result for an exchange economy. Suppose that the economy has a total endowment vector of $\omega = (\omega_1, ..., \omega_N)$. Then, a Pareto optimal allocation can be represented as:

$$\max_{p,\left\{y^h\right\}_{h\in\mathcal{H}}} \sum_{h\in\mathcal{H}} \alpha^h v^h\left(p, y^h\right) = \max_{p,\left\{y^h\right\}_{h\in\mathcal{H}}} \sum_{h\in\mathcal{H}} \alpha^h \left(a^h(p) + b(p)\, y^h\right) \tag{5.7}$$

subject to

$$\begin{aligned} -\frac{1}{b(p)}\left(\sum_{h\in\mathcal{H}} \frac{\partial a^h(p)}{\partial p_j} + \frac{\partial b(p)}{\partial p_j} y\right) &= \omega_j \text{ for } j = 1, ..., N \\ \sum_{h\in\mathcal{H}} y^h &= y \\ \sum_{j=1}^{N} p_j \omega_j &= y, \\ p_j &\geq 0 \text{ for all } j, \end{aligned}$$

where $\left\{\alpha^h\right\}_{h\in\mathcal{H}}$ are nonnegative Pareto weights with $\sum_{h\in\mathcal{H}} \alpha^h = 1$. The first set of constraints use Roy's identity to express the total demand for good j and set it equal to the supply of good j, which is the endowment ω_j. The second equation defines total income as the sum of the income of the households. The third equation makes sure that total income in the economy is equal to the value of the endowments. The third set of constraints requires that all prices are nonnegative.

Now compare the maximization problem (5.7) to the following problem:

$$\max_{p,\left\{y^h\right\}_{h\in\mathcal{H}}} \sum_{h\in\mathcal{H}} a^h(p) + b(p)\, y \tag{5.8}$$

subject to the same set of constraints. The only difference between the two problems is that in the latter, each household has been assigned the same weight. Let $\mathbf{y} \in \mathbb{R}^{|\mathcal{H}|}$ (note that, here, y is a number, whereas $\mathbf{y} = (y^1, .., y^{|\mathcal{H}|})$ is a vector).

Let $(p^*, \mathbf{y}^*)$ be a solution to (5.8) with $y^h = y^*$ for all $h \in \mathcal{H}$. By definition it is also a solution to (5.7) with $\alpha^h = \alpha$, and therefore it is Pareto optimal, which establishes the first part of the theorem.

To establish the second part, suppose that $a^h(p) = a^h$ for all p and all $h \in \mathcal{H}$. To obtain a contradiction, suppose that $(p_\alpha^{**}, \mathbf{y}_\alpha^{**})$ is a solution to (5.7) for some weights $\{\alpha^h\}_{h\in\mathcal{H}}$ and suppose that it is not a solution to (5.8). Let

$$\alpha^M = \max_{h\in\mathcal{H}} \alpha^h$$

and

$$\mathcal{H}^M = \{h \in \mathcal{H} \,|\alpha^h = \alpha^M\}$$

be the set of households given the maximum Pareto weight. Let $(p^*, \mathbf{y}^*)$ be a solution to (5.8) such that

$$y^{h*} = 0 \text{ for all } h \notin \mathcal{H}^M, \tag{5.9}$$

and $y^* = \sum_{h\in\mathcal{H}^M} y^{h*}$. Note that such a solution exists, since the objective function and the constraint set in the second problem depend only on the vector $(y^1, .., y^{|\mathcal{H}|})$ through $y = \sum_{h\in\mathcal{H}} y^h$.

Since, by definition, $(p_\alpha^{**}, \mathbf{y}_\alpha^{**})$ is in the constraint set of (5.8) and is not a solution,

$$\begin{aligned} \sum_{h\in\mathcal{H}} a^h + b(p^*)\, y^* &> \sum_{h\in\mathcal{H}} a^h + b(p_\alpha^{**})\, y_\alpha^{**} \\ b(p^*)\, y^* &> b(p_\alpha^{**})\, y_\alpha^{**}. \end{aligned} \tag{5.10}$$

The hypothesis that $(p_\alpha^{**}, \mathbf{y}_\alpha^{**})$ is a solution to (5.7) implies that

$$\begin{aligned} \sum_{h\in\mathcal{H}} \alpha^h a^h + \sum_{h\in\mathcal{H}} \alpha^h b(p_\alpha^{**})\, y_\alpha^{h**} &\geq \sum_{h\in\mathcal{H}} \alpha^h a^h + \sum_{h\in\mathcal{H}} \alpha^h b(p^*)\, y^{h*} \\ \sum_{h\in\mathcal{H}} \alpha^h b(p_\alpha^{**})\, y_\alpha^{h**} &\geq \sum_{h\in\mathcal{H}} \alpha^h b(p^*)\, y^{h*}, \end{aligned} \tag{5.11}$$

where y^{h*} and y_α^{h**} denote the hth components of the vectors $\mathbf{y}^*$ and $\mathbf{y}_\alpha^{**}$.

Now, note also that any solution $(p^{**}, \mathbf{y}^{**})$ to (5.7) satisfies $y^{h**} = 0$ for any $h \notin \mathcal{H}^M$. In view of this and the choice of $(p^*, \mathbf{y}^*)$ in (5.9), eq. (5.11) implies

$$\begin{aligned} \alpha^M b(p_\alpha^{**}) \sum_{h\in\mathcal{H}} y_\alpha^{h**} &\geq \alpha^M b(p^*) \sum_{h\in\mathcal{H}} y^{h*} \\ b(p_\alpha^{**})\, y_\alpha^{**} &\geq b(p^*)\, y^*, \end{aligned}$$

which contradicts eq. (5.10), and establishes that, under the stated assumptions, any Pareto optimal allocation maximizes the utility of the representative household. □

5.3. Infinite Planning Horizon

Another important microfoundation for the standard preferences used in growth theory and macroeconomics concerns the planning horizon of individuals. Although some growth models are formulated with finitely-lived households (see, e.g., Chapter 9), most growth and macro models assume that households have an infinite-planning horizon as in eq. (5.2) or eq. (5.26) below. A natural question to ask is whether this is a good approximation to reality. After all, most individuals we know are not infinitely-lived.

There are two reasonable microfoundations for this assumption. The first comes from the "Poisson death model" or the *perpetual youth model*, which will be discussed in greater detail in Chapter 9. The general justification for this approach is that, while individuals are finitely-lived, they are not aware of when they will die. Even somebody who is 100 years old will recognize that he cannot consume all his assets, since there is a fair chance that he will live for another 5 or 10 years. At the simplest level, we can consider a discrete-time model and assume that each individual faces a constant probability of death equal to $\nu > 0$. This is a strong simplifying assumption, since the likelihood of survival to the next age in reality is not a constant, but a function of the age of the individual (a feature best captured by actuarial life tables, which are of great importance to the insurance industry). Nevertheless, this is a good starting point, since it is relatively tractable and also implies that individuals have an expected lifespan of $1/\nu < \infty$ periods, which can be used to get a sense of what the value of ν should be.

Suppose also that each individual has a standard instantaneous utility function $u : \mathbb{R}_+ \to \mathbb{R}$, and a "true" or "pure" discount factor $\hat{\beta}$, meaning that this is the discount factor that he would apply between consumption today and tomorrow if he were sure to live between the two dates. Moreover, let us normalize $u(0) = 0$ to be the utility after death. Now consider an individual who plans to have a consumption sequence $\{c(t)\}_{t=0}^{\infty}$ (conditional on living). Clearly, after the individual dies, the future consumption plans do not matter. Standard arguments imply that this individual would have *expected* utility at time $t = 0$ given by

$$\begin{aligned} U(0) &= u(c(0)) + \hat{\beta}(1-\nu)u(c(1)) + \hat{\beta}\nu u(0) \\ &\quad + \hat{\beta}^2(1-\nu)^2 u(c(2)) + \hat{\beta}^2(1-\nu)\nu u(0) + ... \\ &= \sum_{t=0}^{\infty} \left(\hat{\beta}(1-\nu)\right)^t u(c(t)) \\ &= \sum_{t=0}^{\infty} \beta^t u(c(t)), \end{aligned} \tag{5.12}$$

where the second line collects terms and uses $u(0) = 0$, while the third line defines $\beta \equiv \hat{\beta}(1-\nu)$ as the "effective discount factor" of the individual. With this formulation, the model with finite-lives and random death leads to an individual maximization problem identical to that in the model with infinitely-lived households (though the reasonable values of β in this case would differ; see also Exercise 5.7 for a similar derivation in continuous time). Note also the emphasized adjective "expected" utility here. While until now agents faced no uncertainty, the possibility of death implies that there is a non-trivial (in fact quite important!) uncertainty in individuals' lives. As a result, instead of the standard ordinal utility theory, we have to use the expected utility theory as developed by von Neumann and Morgenstern. In particular, eq. (5.12) is already the expected utility of the individual, since probabilities have been substituted in and there is no need to include explicit expectations.

Throughout, except in the stochastic growth analysis in Chapters 16 and 17, I do not introduce expectations operators. Instead whenever possible, I directly specify the expected utility function of households.

A second justification for the infinite planning horizon comes from intergenerational altruism or from the "bequest" motive. At the simplest level, imagine an individual who lives for one period and has a single offspring (who will also live for a single period and will beget a single offspring and so on). Suppose that this individual not only derives utility from his consumption but also from the bequest he leaves to his offspring. For example,

$$U\left(c\left(t\right),b\left(t\right)\right)=u\left(c\left(t\right)\right)+U^{b}\left(b\left(t\right)\right),$$

where $c\left(t\right)$ is his consumption and $b\left(t\right)$ denotes the bequest left to his offspring. For concreteness, let us suppose that the individual has total income $y\left(t\right)$, so that his budget constraint is

$$c\left(t\right)+b\left(t\right)\leq y\left(t\right).$$

The function $U^{b}\left(\cdot\right)$ contains information about how much the individual values bequests left to his offspring. In general, there may be various reasons why individuals leave bequests (including accidental bequests that will be left by an individual facing random death probability as in the example discussed in the previous paragraph). Nevertheless, a natural benchmark might be one in which the individual is "purely altruistic" so that he cares about the utility of his offspring (with some discount factor).[4] Let the discount factor between generations be β. Also assume that the offspring will have an income of w without the bequest. Then, the utility of the individual can be written as

$$u\left(c\left(t\right)\right)+\beta V\left(b\left(t\right)+w\right),$$

where $V\left(\cdot\right)$ can now be interpreted as the continuation value, equal to the utility that the offspring will obtain from receiving a bequest of $b\left(t\right)$ (together with his own income of w). Naturally, the value of the individual at time t can in turn be written as

$$V\left(y\left(t\right)\right)=\max_{c(t)+b(t)\leq y(t)}\left\{u\left(c\left(t\right)\right)+\beta V\left(b\left(t\right)+w\left(t+1\right)\right)\right\},$$

which defines the current-value of the individual starting with income $y\left(t\right)$ and takes into account what the continuation value will be. The next chapter shows that this is the canonical form of a dynamic programming representation of an infinite-horizon maximization problem. In particular, under some mild technical assumptions, this dynamic programming representation is equivalent to maximizing

$$\sum_{s=0}^{\infty}\beta^{s}u\left(c_{t+s}\right)$$

at time t. Intuitively, while each individual lives for one period, he cares about the utility of his offspring, and realizes that in turn his offspring cares about the utility of his own offspring,

[4]The alternative to "purely altruistic" preferences are those in which a parent receives utility from specific types of bequests or from a subcomponent of the utility of his or her offspring. Models with such "impure altruism" are sometimes quite convenient and will be discussed in Chapters 9 and 21.

and so on. This makes each individual internalize the utility of all future members of the "dynasty". Consequently, fully altruistic behavior within a dynasty (so-called "dynastic" preferences) will also lead to an economy in which decision makers act as if they have an infinite planning horizon.

5.4. The Representative Firm

The previous section discussed how the general equilibrium economy admits a representative household only under special circumstances. The other assumption commonly used in growth models, and already introduced in Chapter 2, is the "representative firm" assumption. In particular, recall from Chapter 2 that the entire production side of the economy was represented by an aggregate production possibilities set, which can be thought of as the production possibilities set or the "production function" of a representative firm. One may think that this representation also requires quite stringent assumptions on the production structure of the economy. This is not the case, however. While not all economies would admit a representative household, the standard assumptions adopted in general equilibrium theory or a dynamic general equilibrium analysis (in particular no production externalities and competitive markets) are sufficient to ensure that the formulation with a representative firm is without loss of any generality.

This result is stated in the next theorem. Let us also introduce the following notation: for two vectors of the same dimension, p and y, $p \cdot y$ denotes their *inner product* (that is, if $p = (p_1, ..., p_N)$ and $y = (y_1, ..., y_N)$, then $p \cdot y = \sum_{j=1}^{N} p_j y_j$). In addition, let

$$Y = \left\{ \sum_{f \in \mathcal{F}} y^f \, : \, y^f \in Y^f \text{ for each } f \in \mathcal{F} \right\}. \tag{5.13}$$

denote the *aggregate production possibilities set* of the economy.

Theorem 5.4. ***(The Representative Firm Theorem)*** *Consider a competitive production economy with $N \in \mathbb{N} \cup \{+\infty\}$ commodities and a countable set $\mathcal{F}$ of firms, each with a convex production possibilities set $Y^f \subset \mathbb{R}^N$. Let $p \in \mathbb{R}^N_+$ be the price vector in this economy and denote the set of profit maximizing net supplies of firm $f \in \mathcal{F}$ by $\hat{Y}^f(p) \subset Y^f$ (so that for any $\hat{y}^f \in \hat{Y}^f(p)$, we have $p \cdot \hat{y}^f \geq p \cdot y^f$ for all $y^f \in Y^f$). Then, there exists a* representative firm *with production possibilities set $Y \subset \mathbb{R}^N$ and a set of profit maximizing net supplies $\hat{Y}(p)$ such that for any $p \in \mathbb{R}^N_+$, $\hat{y} \in \hat{Y}(p)$ if and only if $\hat{y} = \sum_{f \in \mathcal{F}} \hat{y}^f$ for some $\hat{y}^f \in \hat{Y}^f(p)$ for each $f \in \mathcal{F}$.*

Proof. Let Y be the aggregate production possibilities set defined in (5.13).To prove the "if" part of the theorem, fix $p \in \mathbb{R}^N_+$ and construct $\hat{y} = \sum_{f \in \mathcal{F}} \hat{y}^f$ for some $\hat{y}^f \in \hat{Y}^f(p)$ for each $f \in \mathcal{F}$. Suppose, to obtain a contradiction, that $\hat{y} \notin \hat{Y}(p)$, so that there exists y' such that $p \cdot y' > p \cdot \hat{y}$. By definition of the set Y, this implies that there exists $\left\{y^f\right\}_{f \in \mathcal{F}}$ with

$y^f \in Y^f$ such that

$$\begin{aligned} p \cdot \left(\sum_{f \in \mathcal{F}} y^f \right) &> p \cdot \left(\sum_{f \in \mathcal{F}} \hat{y}^f \right) \\ \sum_{f \in \mathcal{F}} p \cdot y^f &> \sum_{f \in \mathcal{F}} p \cdot \hat{y}^f, \end{aligned}$$

so that there exists at least one $f' \in \mathcal{F}$ such that

$$p \cdot y^{f'} > p \cdot \hat{y}^{f'},$$

which contradicts the hypothesis that $\hat{y}^f \in \hat{Y}^f(p)$ for each $f \in \mathcal{F}$ and completes this part of the proof.

To prove the "only if" part of the theorem, let $\hat{y} \in \hat{Y}(p)$ be a profit maximizing choice for the representative firm. Then, since $\hat{Y}(p) \subset Y$, we have that

$$\hat{y} = \sum_{f \in \mathcal{F}} y^f$$

for some $y^f \in Y^f$ for each $f \in \mathcal{F}$. Let $\hat{y}^f \in \hat{Y}^f(p)$. Then,

$$\sum_{f \in \mathcal{F}} p \cdot y^f \leq \sum_{f \in \mathcal{F}} p \cdot \hat{y}^f,$$

which implies that

$$p \cdot \hat{y} \leq p \cdot \sum_{f \in \mathcal{F}} \hat{y}^f. \tag{5.14}$$

Since, by hypothesis, $\sum_{f \in \mathcal{F}} \hat{y}^f \in Y$ and $\hat{y} \in \hat{Y}(p)$,

$$p \cdot \hat{y} \geq p \cdot \sum_{f \in \mathcal{F}} \hat{y}^f.$$

Therefore, inequality (5.14) must hold with equality, so that

$$p \cdot y^f = p \cdot \hat{y}^f,$$

for each $f \in \mathcal{F}$, and thus $y^f \in \hat{Y}^f(p)$. This completes the proof of the theorem. □

This theorem implies that, given the assumptions that there are "no externalities" and that all factors are priced competitively, focusing on the aggregate production possibilities set of the economy or on the representative firm is without loss of any generality. Why is there such a difference between the representative household and representative firm assumptions? The answer is related to income effects. The reason why the representative household assumption is restrictive is that changes in prices create income effects, which affect different households differently. A representative household exists only when these income effects can be ignored, which is what the Gorman preferences guarantee. Since there are no income effects in producer theory, the representative firm assumption is without loss of any generality.

Naturally, the fact that the production side of an economy can be modeled via a representative firm does not mean that heterogeneity among firms is uninteresting or unimportant.

On the contrary, productivity differences across firms and firms' attempts to increase their productivity relative to others are central phenomena in this study of economic growth. Theorem 5.4 simply says that given the production possibilities sets of the firms in the economy, the production side of the economy can be equivalently represented by a single representative firm or an aggregate production possibilities set.

5.5. Problem Formulation

Let us now consider a discrete-time infinite-horizon economy and suppose that the economy admits a representative household. In particular, once again ignoring uncertainty, the representative household has the $t = 0$ objective function

$$\sum_{t=0}^{\infty} \beta^t u\left(c\left(t\right)\right), \tag{5.15}$$

with a discount factor of $\beta \in (0, 1)$.

In continuous time, this utility function of the representative household becomes

$$\int_0^{\infty} \exp\left(-\rho t\right) u\left(c\left(t\right)\right) dt \tag{5.16}$$

where $\rho > 0$ is now the discount rate of the individuals.

Where does the exponential form of the discounting in (5.16) come from? Discounting in the discrete-time case was already referred to as "exponential", so the link should be apparent. To see it more precisely, imagine we are trying to calculate the value of \$1 in T periods, and divide the interval $[0, T]$ into $T/\Delta t$ equally-sized subintervals. Let the interest rate in each subinterval be equal to $r\Delta t$. It is important that the quantity r is multiplied by Δt, otherwise as we vary Δt, we would be changing the interest rate. Using the standard compound interest rate formula, the value of \$1 in T periods at this interest rate is given by

$$v\left(T \mid \Delta t\right) \equiv \left(1 + r\Delta t\right)^{T/\Delta t}.$$

Next, let us approach the continuous time limit by letting $\Delta t \to 0$ to obtain

$$v\left(T\right) \equiv \lim_{\Delta t \to 0} v\left(T \mid \Delta t\right) \equiv \lim_{\Delta t \to 0} \left(1 + r\Delta t\right)^{T/\Delta t}.$$

Since the limit operator is continuous,

$$\begin{aligned} v\left(T\right) &\equiv \exp\left[\lim_{\Delta t \to 0} \ln\left(1 + r\Delta t\right)^{T/\Delta t}\right] \\ &= \exp\left[\lim_{\Delta t \to 0} \frac{T}{\Delta t} \ln\left(1 + r\Delta t\right)\right]. \end{aligned}$$

However, the term in square brackets has a limit of the form $0/0$ (or of the form $\infty \times 0$). Let us next write this as

$$\lim_{\Delta t \to 0} \frac{\ln\left(1 + r\Delta t\right)}{\Delta t/T} = \lim_{\Delta t \to 0} \frac{r/\left(1 + r\Delta t\right)}{1/T} = rT,$$

where the first equality follows from L'Hospital's Rule (Theorem A.22 in Appendix Chapter A). Therefore,

$$v\left(T\right) = \exp\left(rT\right).$$

Conversely, \$1 in T periods from now, is worth $\exp(-rT)$ today. The same reasoning applies to discounting utility, so the utility of consuming $c(t)$ in period t evaluated at time $t=0$ is $\exp(-\rho t)\,u(c(t))$, where ρ denotes the (subjective) discount rate.

5.6. Welfare Theorems

We are ultimately interested in equilibrium growth. But there is a close connection between Pareto optima and competitive equilibria. These connections could not be exploited in the preceding chapters, because household (individual) preferences were not specified. I now introduce the First and Second Welfare Theorems and develop the relevant connections between the theory of economic growth and dynamic general equilibrium models.

Let us start with models that have a finite number of consumers, so that in terms of the notation above, the set $\mathcal{H}$ is finite. Throughout, I allow an infinite number of commodities, since dynamic growth models almost always feature an infinite number of time periods, thus an infinite number of commodities. The results stated in this section have analogs for economies with a continuum of commodities (corresponding to dynamic economies in continuous time), but for the sake of brevity and to reduce technical details, I focus on economies with a countable number of commodities.

Therefore, let the commodities be indexed by $j \in \mathbb{N}$ and $x^h \equiv \left\{x_j^h\right\}_{j=0}^{\infty}$ be the consumption bundle of household h, and $\omega^h \equiv \left\{\omega_j^h\right\}_{j=0}^{\infty}$ be its endowment bundle. In addition, let us assume that feasible x^h's must belong to some *consumption set* $X^h \subset \mathbb{R}_+^{\infty}$. The most relevant interpretation for us is that there is an infinite number of dates, say indexed by t, and at each date $t=0,1,...,$ each household consumes a finite-dimensional vector of products, so that $\tilde{x}_t^h = \left\{\tilde{x}_{1,t}^h,...,\tilde{x}_{K,t}^h\right\} \in \tilde{X}_t^h \subset \mathbb{R}_+^K$ for some $K \in \mathbb{N}$, and $x^h = \left\{\tilde{x}_t^h\right\}_{t=0}^{\infty}$. The consumption sets are introduced to ensure that households do not have negative consumption levels and are thus subsets of $\mathbb{R}_+^{\infty}$ (this can be extended by allowing some components of the vector, corresponding to different types of labor supply, to be negative; this is straightforward and I do not do this to conserve notation).

Let $\mathbf{X} \equiv \prod_{h\in\mathcal{H}} X^h$ be the Cartesian product of these consumption sets, which can be thought of as the aggregate consumption set of the economy. I also use the notation $\mathbf{x} \equiv \left\{x^h\right\}_{h\in\mathcal{H}}$ and $\boldsymbol{\omega} \equiv \left\{\omega^h\right\}_{h\in\mathcal{H}}$ to describe the entire set of consumption allocation and endowments in the economy. Feasibility of a consumption allocation requires that $\mathbf{x} \in \mathbf{X}$.

Each household in $\mathcal{H}$ has a well defined preference ordering over consumption bundles. At the most general level, this preference ordering can be represented by a relationship $\succsim_h$ for household h, such that $x' \succsim_h x$ implies that household h weakly prefers $\mathbf{x}'$ to $\mathbf{x}$. When these preferences satisfy some relatively weak properties (completeness, reflexivity and transitivity), they can equivalently be represented by a real-valued utility function $u^h : X^h \to \mathbb{R}$, such that whenever $x' \succsim_h x$, we have $u^h(x') \geq u^h(x)$. The domain of this function is $X^h \subset \mathbb{R}_+^{\infty}$.

We also assume that u^h is nondecreasing in each of its arguments for each $h \in \mathcal{H}$. Let $\mathbf{u} \equiv \left\{u^h\right\}_{h\in\mathcal{H}}$ be the set of utility functions.[5]

Let us next describe the production side. Suppose that there is a finite number of firms represented by the set $\mathcal{F}$ and that each firm $f \in \mathcal{F}$ is characterized by a production set Y^f, which specifies what levels of output firm f can produce from specified levels of inputs. In other words, $y^f \equiv \left\{y_j^f\right\}_{j=0}^{\infty}$ is a feasible production plan for firm f if $y^f \in Y^f$. For example, if there were only two commodities, labor and a final good, Y^f would include pairs $(-l, y)$ such that with labor input l (hence a negative sign), the firm can produce at most as much as y. As is usual in general equilibrium theory, let us take each Y^f to be a *cone,* so that if $y \in Y^f$, then $\lambda y \in Y^f$ for any $\lambda \in \mathbb{R}_+$. This implies two important features: first, $0 \in Y^f$ for each $f \in \mathcal{F}$; second, each Y^f exhibits constant returns to scale. If there are diminishing returns to scale because of some scarce factors, such as entrepreneurial talent, this is added as an additional factor of production and Y^f is still interpreted as a cone.

Let $\mathbf{Y} \equiv \prod_{f\in\mathcal{F}} Y^f$ represent the aggregate production set in this economy and $\mathbf{y} \equiv \left\{y^f\right\}_{f\in\mathcal{F}}$ such that $y^f \in Y^f$ for all f, or equivalently, $\mathbf{y} \in \mathbf{Y}$.

The final object that needs to be described is the ownership structure of firms. In particular, if firms make profits, they should be distributed to some agents in the economy. We capture this by assuming that there exists a sequence of numbers (profit shares) represented by $\boldsymbol{\theta} \equiv \left\{\theta_f^h\right\}_{f\in\mathcal{F},h\in\mathcal{H}}$ such that $\theta_f^h \geq 0$ for all f and h, and $\sum_{h\in\mathcal{H}} \theta_f^h = 1$ for all $f \in \mathcal{F}$. The number θ_f^h is the share of profits of firm f that will accrue to household h.

An economy $\mathcal{E}$ is described by preferences, endowments, production sets, consumption sets and allocation of shares, that is, $\mathcal{E} \equiv (\mathcal{H}, \mathcal{F}, \mathbf{u}, \boldsymbol{\omega}, \mathbf{Y}, \mathbf{X}, \boldsymbol{\theta})$. An allocation in this economy is $(\mathbf{x}, \mathbf{y})$ such that $\mathbf{x}$ and $\mathbf{y}$ are feasible; $\mathbf{x} \in \mathbf{X}$, $\mathbf{y} \in \mathbf{Y}$, and $\sum_{h\in\mathcal{H}} x_j^h \leq \sum_{h\in\mathcal{H}} \omega_j^h + \sum_{f\in\mathcal{F}} y_j^f$ for all $j \in \mathbb{N}$. The last requirement implies that the total consumption of each commodity has to be less than the sum of its total endowment and net production.

A price system is a sequence $p \equiv \{p_j\}_{j=0}^{\infty}$, such that $p_j \geq 0$ for all j, with one of the commodities chosen as the numeraire and its price normalized to 1. Recall that $p \cdot x$ is again the inner product of p and x, which are now infinite sequences, so that $p \cdot x \equiv \sum_{j=0}^{\infty} p_j x_j$.[6]

A competitive economy refers to an environment without any externalities and where all commodities are traded competitively. In a competitive equilibrium, all firms maximize profits, all consumers maximize their utility given their budget set and all markets clear. More formally:

DEFINITION 5.1. *A **competitive equilibrium** for economy $\mathcal{E} \equiv (\mathcal{H}, \mathcal{F}, \mathbf{u}, \boldsymbol{\omega}, \mathbf{Y}, \mathbf{X}, \boldsymbol{\theta})$ is given by an allocation $\left(\mathbf{x}^* = \left\{x^{h*}\right\}_{h\in\mathcal{H}}, \mathbf{y}^* = \left\{y^{f*}\right\}_{f\in\mathcal{F}}\right)$ and a price system p^* such that*

[5]Notice that despite the same notation, u, utility functions here are not instantaneous utility functions or felicity functions as in Section 5.1, but they are defined over the entire sequence of consumption bundles.

[6]You may note that such an inner product may not always exist in infinite dimensional spaces. This technical detail is dealt with in the proof of Theorem 5.7 below.

(1) *The allocation* $(\mathbf{x}^*, \mathbf{y}^*)$ *is feasible, i.e.,* $x^{h*} \in X^h$ *for all* $h \in \mathcal{H}$, $y^{f*} \in Y^f$ *for all* $f \in \mathcal{F}$ *and*

$$\sum_{h \in \mathcal{H}} x_j^{h*} \leq \sum_{h \in \mathcal{H}} \omega_j^h + \sum_{f \in \mathcal{F}} y_j^{f*} \text{ for all } j \in \mathbb{N}.$$

(2) *For every firm* $f \in \mathcal{F}$, y^{f*} *maximizes profits, i.e.,*

$$p^* \cdot y^{f*} \geq p^* \cdot y \text{ for all } y \in Y^f.$$

(3) *For every consumer* $h \in \mathcal{H}$, x^{h*} *maximizes utility, i.e.,*

$$u^h\left(x^{h*}\right) \geq u^h(x) \text{ for all } x \text{ such that } x \in X^h \text{ and } p^* \cdot x \leq p^* \cdot \left(\omega^h + \sum_{f \in \mathcal{F}} \theta_f^h y^f\right).$$

A major focus of general equilibrium theory is to establish the existence of a competitive equilibrium under reasonable assumptions. When there is a finite number of commodities and standard convexity assumptions are made on preferences and production sets, this is straightforward (in particular, the proof of existence involves simple applications of Theorems A.16, A.17, and A.19 in Appendix Chapter A). When there is an infinite number of commodities, as in infinite-horizon growth models, proving the existence of a competitive equilibrium is somewhat more difficult and requires more sophisticated arguments. Nevertheless, for our focus here proving the existence of a competitive equilibrium under general conditions is not central (since the typical growth models will have sufficient structure to ensure the existence of a competitive equilibrium in a relatively straightforward manner). Instead, the efficiency properties of competitive equilibria, when they exist, and the decentralization of certain desirable (efficient) allocations as competitive equilibria are more important. For this reason, let us recall the standard definition of Pareto optimality.

DEFINITION 5.2. *A feasible allocation* $(\mathbf{x}, \mathbf{y})$ *for economy* $\mathcal{E} \equiv (\mathcal{H}, \mathcal{F}, \mathbf{u}, \boldsymbol{\omega}, \mathbf{Y}, \mathbf{X}, \boldsymbol{\theta})$ *is* ***Pareto optimal*** *if there exists no other feasible allocation* $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ *such that* $\hat{x}^h \in X^h$ *for all* $h \in \mathcal{H}$, $\hat{y}^f \in Y^f$ *for all* $f \in \mathcal{F}$,

$$\sum_{h \in \mathcal{H}} \hat{x}_j^h \leq \sum_{h \in \mathcal{H}} \omega_j^h + \sum_{f \in \mathcal{F}} \hat{y}_j^f \text{ for all } j \in \mathbb{N},$$

and

$$u^h\left(\hat{x}^h\right) \geq u^h\left(x^h\right) \text{ for all } h \in \mathcal{H}$$

with at least one strict inequality in the preceding relationship.

Our next result is the celebrated *First Welfare Theorem* for competitive economies. Before presenting this result, we need the following definition.

DEFINITION 5.3. *Household* $h \in \mathcal{H}$ *is* ***locally non-satiated*** *at* x^h *if* $u^h\left(x^h\right)$ *is strictly increasing in at least one of its arguments at* x^h *and* $u^h\left(x^h\right) < \infty$. *Household* $h \in \mathcal{H}$ *is* ***non-satiated*** *if it is locally non-satiated at all* $x^h \in X^h$.

The latter requirement in this definition is already implied by the fact that $u^h : X^h \to \mathbb{R}$, but it is included for additional emphasis, since it is important for the proof and also because if, in fact, $u^h\left(x^h\right) = \infty$, we could not meaningfully talk of $u^h\left(x^h\right)$ being strictly increasing.

THEOREM 5.5. ***(First Welfare Theorem I)*** *Suppose that* $(\mathbf{x}^*, \mathbf{y}^*, p^*)$ *is a competitive equilibrium of economy* $\mathcal{E} \equiv (\mathcal{H}, \mathcal{F}, \mathbf{u}, \boldsymbol{\omega}, \mathbf{Y}, \mathbf{X}, \boldsymbol{\theta})$ *with* $\mathcal{H}$ *finite. Assume that all households are non-satiated. Then,* $(\mathbf{x}^*, \mathbf{y}^*)$ *is Pareto optimal.*

PROOF. Suppose that $(\mathbf{x}^*, \mathbf{y}^*, p^*)$ is a competitive equilibrium. To obtain a contradiction, suppose that there exists a feasible $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ such that $u^h\left(\hat{x}^h\right) \geq u^h\left(x^{h*}\right)$ for all $h \in \mathcal{H}$ and $u^h\left(\hat{x}^h\right) > u^h\left(x^{h*}\right)$ for all $h \in \mathcal{H}'$, where $\mathcal{H}'$ is a nonempty subset of $\mathcal{H}$.

Since $(\mathbf{x}^*, \mathbf{y}^*, p^*)$ is a competitive equilibrium, it must be the case that for all $h \in \mathcal{H}$,

$$\begin{aligned} p^* \cdot \hat{x}^h &\geq p^* \cdot x^{h*} \\ &= p^* \cdot \left(\omega^h + \sum_{f \in \mathcal{F}} \theta_f^h y^{f*}\right) \end{aligned} \tag{5.17}$$

and for all $h \in \mathcal{H}'$,

$$p^* \cdot \hat{x}^h > p^* \cdot \left(\omega^h + \sum_{f \in \mathcal{F}} \theta_f^h y^{f*}\right). \tag{5.18}$$

The second inequality follows immediately in view of the fact that x^{h*} is the utility maximizing choice for household h, thus if $\hat{x}^h$ is strictly preferred, then it cannot be in the budget set. The first inequality follows with a similar reasoning. Suppose that it did not hold. Then, by the hypothesis of local-satiation, u^h must be strictly increasing in at least one of its arguments, let us say the j'th component of x. Then, construct $\hat{x}^h(\varepsilon)$ such that $\hat{x}_j^h(\varepsilon) = \hat{x}_j^h$ and $\hat{x}_{j'}^h(\varepsilon) = \hat{x}_{j'}^h + \varepsilon$ for $\varepsilon > 0$. For $\varepsilon \to 0$, $\hat{x}^h(\varepsilon)$ is in household h's budget set and yields strictly greater utility than the original consumption bundle $\hat{x}^h$, contradicting the hypothesis that household h was maximizing utility.

Also note that non-satiation implies that $u^h\left(x^h\right) < \infty$, and thus the right-hand sides of (5.17) and (5.18) are finite.

Now summing (5.17) over $\mathcal{H} \backslash \mathcal{H}'$ and (5.18) over $\mathcal{H}'$, and combining the two, we have

$$\begin{aligned} p^* \cdot \sum_{h \in \mathcal{H}} \hat{x}^h &> p^* \cdot \sum_{h \in \mathcal{H}} \left(\omega^h + \sum_{f \in \mathcal{F}} \theta_f^h y^{f*}\right) \\ &= p^* \cdot \left(\sum_{h \in \mathcal{H}} \omega^h + \sum_{f \in \mathcal{F}} y^{f*}\right), \end{aligned} \tag{5.19}$$

where the second line uses the fact that the sums are finite, so that the order of summation can be exchanged, and that by definition of shares $\sum_{h \in \mathcal{H}} \theta_f^h = 1$ for all $f \in \mathcal{F}$ (it also implicitly uses the fact that by profit maximization $p^* \cdot y^{f*} \geq 0$ for all $f \in \mathcal{F}$, so that the

sum of the value of profits over firms exists). Finally, since $\mathbf{y}^*$ is profit-maximizing at prices p^*, we have that

$$p^* \cdot \sum_{f \in \mathcal{F}} y^{f*} \geq p^* \cdot \sum_{f \in \mathcal{F}} y^f \text{ for any } \left\{y^f\right\}_{f \in \mathcal{F}} \text{ with } y^f \in Y^f \text{ for all } f \in \mathcal{F}. \tag{5.20}$$

However, by feasibility of $\hat{x}^h$ (Definition 5.1, Part 1),

$$\sum_{h \in \mathcal{H}} \hat{x}_j^h \leq \sum_{h \in \mathcal{H}} \omega_j^h + \sum_{f \in \mathcal{F}} \hat{y}_j^f \text{ for all } j,$$

and therefore, by taking the inner products of both sides with p^*, and exploiting (5.20) and the fact that $p^* \geq 0$,

$$\begin{aligned} p^* \cdot \sum_{h \in \mathcal{H}} \hat{x}_j^h &\leq p^* \cdot \left(\sum_{h \in \mathcal{H}} \omega_j^h + \sum_{f \in \mathcal{F}} \hat{y}_j^f \right) \\ &\leq p^* \cdot \left(\sum_{h \in \mathcal{H}} \omega_j^h + \sum_{f \in \mathcal{F}} y_j^{f*} \right), \end{aligned}$$

which contradicts (5.19), establishing that any competitive equilibrium allocation $(\mathbf{x}^*, \mathbf{y}^*)$ is Pareto optimal. □

The proof of the First Welfare Theorem is both intuitive and simple. The proof is based on two intuitive ideas. First, if another allocation Pareto dominates the competitive equilibrium, then it must be non-affordable in the competitive equilibrium for at least one household. Second, profit-maximization implies that any competitive equilibrium already maximizes the set of affordable allocations. It is also simple since it only uses the summation of the values of commodities at a given price vector. In particular, it makes no convexity assumption. However, the proof also highlights the importance of the feature that the relevant sums exist and are finite. Otherwise, the last step would lead to the conclusion that "$\infty < \infty$" which may or may not be a contradiction. The fact that these sums exist, in turn, followed from two assumptions: finiteness of the number of individuals and non-satiation. However, as noted before, working with economies that have only a finite number of households (even if there is an infinite number of commodities) is not always sufficient for our purposes. For this reason, the next theorem turns to the version of the First Welfare Theorem with an infinite number of households. For simplicity, here I take $\mathcal{H}$ to be a countably infinite set, e.g., $\mathcal{H} = \mathbb{N}$. The next theorem generalizes the First Welfare Theorem to this case. It makes use of an additional assumption to take care of infinite sums.

THEOREM 5.6. ***(First Welfare Theorem II)*** *Suppose that $(\mathbf{x}^*, \mathbf{y}^*, p^*)$ is a competitive equilibrium of the economy $\mathcal{E} \equiv (\mathcal{H}, \mathcal{F}, \mathbf{u}, \boldsymbol{\omega}, \mathbf{Y}, \mathbf{X}, \boldsymbol{\theta})$ with $\mathcal{H}$ countably infinite. Assume that all households are locally non-satiated at $\mathbf{x}^*$, that the endowment of each good is finite ($\sum_{h \in \mathcal{H}} \omega_j^h < \infty$ for all j) and that $\sum_{j=0}^{\infty} p_j^* < \infty$. Then, $(\mathbf{x}^*, \mathbf{y}^*, p^*)$ is Pareto optimal.*

PROOF. The proof is the same as that of Theorem 5.5, with a major difference. Local non-satiation does not guarantee that the summations are finite (5.19), since the summations are

over an infinite number of households. However, since endowments are finite, the assumption that $\sum_{j=0}^{\infty} p_j^* < \infty$ ensures that the sums in (5.19) are indeed finite and the rest of the proof goes through exactly as in the proof of Theorem 5.5. □

Theorem 5.6 will be particularly useful in the analysis of overlapping generation models in Chapter 9.

Let us next briefly discuss the Second Welfare Theorem, which is the converse of the First Welfare Theorem. It answers the question of whether a Pareto optimal allocation can be decentralized as a competitive equilibrium. Interestingly, for the Second Welfare Theorem whether or not $\mathcal{H}$ is finite is not as important as for the First Welfare Theorem. Nevertheless, the Second Welfare Theorem requires a number of assumptions on preferences and technology, such as the convexity of consumption and production sets and of preferences, and a number of additional technical requirements (which are trivially satisfied when the number of commodities is finite). This is because the Second Welfare Theorem implicitly contains an "existence of equilibrium argument," which runs into problems in the presence of non-convexities. Before stating this theorem, recall that the consumption set of each household $h \in \mathcal{H}$ is $X^h \subset \mathbb{R}_+^\infty$, so a typical element of X^h is $x^h = \left(x_0^h, x_1^h, x_2^h, ...\right)$, where x_t^h can be interpreted as the vector of consumption of individual h at time t. Similarly, a typical element of the production set of firm $f \in \mathcal{F}$, Y^f, is of the form $y^f = \left(y_0^f, y_1^f, y_2^f, ...\right)$. Let us define $x^h[T] = \left(x_0^h, x_1^h, x_2^h, ..., x_T^h, 0, 0, ...\right)$ and $y^f[T] = \left(y_0^f, y_1^f, y_2^f, ..., y_T^f, 0, 0, ...\right)$. In other words, these are truncated sequences that involve zero consumption or zero production after some date T. It can be verified that $\lim_{T\to\infty} x^h[T] = x^h$ and $\lim_{T\to\infty} y^f[T] = y^f$ in the product topology (see Section A.4 in Appendix Chapter A). Finally, since in this case each x (or y) is a K dimensional vector, with a slight abuse of notation, I use $p \cdot x$ for an appropriately-defined inner product, for example, $p \cdot x^h = \sum_{t=0}^{\infty} \sum_{j=1}^{K} p_{j,t} x_{j,t}^h$.

THEOREM 5.7. ***(Second Welfare Theorem)*** *Consider a Pareto optimal allocation* $(\mathbf{x}^*, \mathbf{y}^*)$ *in an economy with endowment vector* ω*, production sets* $\left\{Y^f\right\}_{f\in\mathcal{F}}$*, consumption sets* $\left\{X^h\right\}_{h\in\mathcal{H}}$*, and utility functions* $\left\{u^h(\cdot)\right\}_{h\in\mathcal{H}}$*. Suppose that all production and consumption sets are convex, all production sets are cones, and all utility functions* $\left\{u^h(\cdot)\right\}_{h\in\mathcal{H}}$ *are continuous and quasi-concave and satisfy local non-satiation. Moreover, suppose also that* $0 \in X^h$ *for each* $h \in \mathcal{H}$*, that for each* $x, x' \in X^h$ *with* $u^h(x) > u^h(x')$ *for all* $h \in \mathcal{H}$*, there exists* $\bar{T}$ *(possibly as a function of* x *and* x'*) such that* $u^h(x[T]) > u^h(x')$ *for all* $T \geq \bar{T}$ *and for all* $h \in \mathcal{H}$ *and that for each* $y \in Y^f$*, there exists* $\tilde{T}$ *such that* $y[T] \in Y^f$ *for all* $T \geq \tilde{T}$ *and for all* $f \in \mathcal{F}$*. Then, there exist a price vector* p^* *and an endowment and share allocations* $(\boldsymbol{\omega}^*, \boldsymbol{\theta}^*)$ *such that in economy* $\mathcal{E} \equiv (\mathcal{H}, \mathcal{F}, \mathbf{u}, \boldsymbol{\omega}^*, \mathbf{Y}, \mathbf{X}, \boldsymbol{\theta}^*)$*,*

(1) *the endowment allocation* $\boldsymbol{\omega}^*$ *satisfies* $\omega = \sum_{h\in\mathcal{H}} \omega^{h*}$*;*
(2) *for all* $f \in \mathcal{F}$*,*

$$p^* \cdot y^{f*} \geq p^* \cdot y \text{ for all } y \in Y^f;$$

(3) *for all* $h \in \mathcal{H}$,

$$\text{if } x^h \in X^h \text{ involves } u^h\left(x^h\right) > u^h\left(x^{h*}\right), \text{ then } p^* \cdot x^h \geq p^* \cdot w^{h*},$$

where $w^{h*} \equiv \omega^{h*} + \sum_{f \in \mathcal{F}} \theta_f^{h*} y^{f*}$.

Moreover, if $p^* \cdot \mathbf{w}^* > 0$ *[i.e.,* $p^* \cdot w^{h*} > 0$ *for each* $h \in \mathcal{H}$*], then economy* $\mathcal{E}$ *has a competitive equilibrium* $(\mathbf{x}^*, \mathbf{y}^*, p^*)$.

The proof of this theorem involves the application of the Geometric Hahn-Banach Theorem, Theorem A.28, from Appendix Chapter A. It is somewhat long and involved. For this reason, a sketch of this proof is provided in the next (starred) section. Here notice that if instead of an infinite-dimensional economy, we were dealing with an economy with a finite commodity space, say with K commodities, then the hypothesis in the theorem, that $0 \in X^h$ for each $h \in \mathcal{H}$ and that for each $x, x' \in X^h$ with $u^h(x) > u^h(x')$ for all $h \in \mathcal{H}$, there exists $\bar{T}$ such that $u^h(x[T]) > u^h(x')$ for all $T \geq \bar{T}$ and for all $h \in \mathcal{H}$ (and also that for each $y \in Y^f$, there exists $\tilde{T}$ such that $y[T] \in Y^f$ for all $T \geq \tilde{T}$ and for all $f \in \mathcal{F}$) would be satisfied automatically, by taking $\bar{T} = \tilde{T} = K$. In fact, this condition is not imposed in the statement of the Second Welfare Theorem in economies with a finite number of commodities. Its role in dynamic economies is that changes in allocations that are very far in the future should not have a "large" effect on preferences. This is naturally satisfied when we look at infinite-horizon economies with discounted utility and separable production structure. Intuitively, if a sequence of consumption levels x is strictly preferred to x', then setting the elements of x and x' to 0 in the very far (and thus heavily discounted) future should not change this conclusion (since discounting implies that x could not be strictly preferred to x' because of higher consumption under x in the arbitrarily far future). Similarly, if some production vector y is feasible, the separable production structure implies that $y[T]$, which involves zero production after some date T, must also be feasible. Exercise 5.13 demonstrates these claims more formally. One difficulty in applying this theorem is that u^h may not be defined when x has zero elements (so that the consumption set X^h does not contain 0). Exercise 5.14 shows that the theorem can be generalized to the case in which there exists a strictly positive vector $\varepsilon \in \mathbb{R}^K$ with each element sufficiently small and with $(\varepsilon, ..., \varepsilon) \in X^h$ for all $h \in \mathcal{H}$.

The conditions for the Second Welfare Theorem are more difficult to satisfy than those for the First Welfare Theorem because of the convexity requirements. In many ways, it is also the more important of the two theorems. While the First Welfare Theorem is celebrated as a formalization of Adam Smith's invisible hand, the Second Welfare Theorem establishes the stronger result that any Pareto optimal allocation can be *decentralized* as a competitive equilibrium. An immediate corollary of this is an existence result; since the Pareto optimal allocation can be decentralized as a competitive equilibrium, a competitive equilibrium must exist (at least for the endowments leading to Pareto optimal allocations).

The Second Welfare Theorem motivates many macroeconomists to look for the set of Pareto optimal allocations instead of explicitly characterizing competitive equilibria. This is

especially useful in dynamic models where sometimes competitive equilibria can be quite difficult to characterize or even to specify, while the characterization of Pareto optimal allocations is typically more straightforward.

The real power of the Second Welfare Theorem in dynamic macro models comes when we combine it with models that admit a representative household. Recall that Theorem 5.3 shows an equivalence between Pareto optimal allocations and optimal allocations for the representative household. In certain models, including many—but not all—growth models studied in this book, the combination of a representative household and the Second Welfare Theorem enables us to characterize *the optimal growth allocation* that maximizes the utility of the representative household and assert that this will correspond to a competitive equilibrium.

5.7. Proof of the Second Welfare Theorem, Theorem 5.7*

In this section, I provide a proof of the Second Welfare Theorem. The most important part of the theorem is proved by using the Geometric Hahn-Banach Theorem, Theorem A.28, from Appendix Chapter A

Proof of Theorem 5.7: **(Sketch)** First, I establish that there exists a price vector p^* and an endowment and share allocation $(\boldsymbol{\omega}^*, \boldsymbol{\theta}^*)$ that satisfy conditions 1-3. This has two parts.

(Part 1) This part follows from the Geometric Hahn-Banach Theorem, Theorem A.28. Define the "more preferred" sets for each $h \in \mathcal{H}$:

$$P^h = \left\{ x^h \in X^h : u^h\left(x^h\right) > u^h\left(x^{h*}\right) \right\}.$$

Clearly, each P^h is convex. Let $P = \sum_{h\in\mathcal{H}} P^h$ and $Y' = \sum_{f\in\mathcal{F}} Y^f + \{\omega\}$, where recall that $\omega = \sum_{h\in\mathcal{H}} \omega^{h*}$, so that Y' is the sum of the production sets shifted by the endowment vector. Both P and Y' are convex (since each P^h and each Y^f are convex). Consider the sequences of production plans for each firm to be elements of the vector space ℓ_∞^K, which includes infinite sequences of vectors of the form $y^f = \left(y_0^f, y_1^f, ...\right)$, with each $y_j^f \in \mathbb{R}_+^K$. Moreover, since each production set is a cone, $Y' = \sum_{f\in\mathcal{F}} Y^f + \{\omega\}$ has an interior point (the argument is identical to that of Exercise A.32 in Appendix Chapter A). Moreover, let $x^* = \sum_{h\in\mathcal{H}} x^{h*}$. By feasibility and local non-satiation, $x^* = \sum_{f\in\mathcal{F}} y^{h*} + \omega$. Then, $x^* \in Y'$ and also $x^* \in \overline{P}$ (where recall that $\overline{P}$ is the closure of P).

Next, observe that $P \cap Y' = \varnothing$. Otherwise, there would exist $\tilde{y} \in Y'$, which is also in P. This would imply that if distributed appropriately across the households, $\tilde{y}$ would make all households equally well off and at least one of them would be strictly better-off (for example, by the definition of the set P, there would exist $\left\{\tilde{x}^h\right\}_{h\in\mathcal{H}}$ such that $\sum_{h\in\mathcal{H}} \tilde{x}^h = \tilde{y}$, $\tilde{x}^h \in X^h$, and $u^h\left(\tilde{x}^h\right) \geq u^h\left(x^{h*}\right)$ for all $h \in \mathcal{H}$ with at least one strict inequality). This would contradict the hypothesis that (x^*, y^*) is a Pareto optimum.

Since Y' has an interior point, P and Y' are convex, and $P \cap Y' = \varnothing$, Theorem A.28 implies that there exists a nonzero continuous linear functional ϕ such that

$$\phi(y) \leq \phi(x^*) \leq \phi(x) \text{ for all } y \in Y' \text{ and all } x \in P. \tag{5.21}$$

(Part 2) I next show that the linear functional ϕ can be interpreted as a price vector, in particular, that it does have an inner product representation. Consider the functional

$$\bar{\phi}(x) = \lim_{T\to\infty} \phi(x[T]),$$

where recall that for $x^h = \left(x_0^h, x_1^h, x_2^h, ...\right)$, $x^h[T] = \left(x_0^h, x_1^h, x_2^h, ..., x_T^h, 0, 0, ...\right)$. Moreover, let $\bar{x}_J^h$ be defined as $\bar{x}_J^h = \left(0, 0, ..., x_J^h, 0, ..., \right)$, that is, as the sequence x^h with zeros everywhere except its Jth element. Finally, define

$$\bar{z}_t^{\phi} = \begin{cases} \bar{x}_t & \text{if } \phi(\bar{x}_t) \geq 0 \\ -\bar{x}_t & \text{if } \phi(\bar{x}_t) < 0 \end{cases}$$

and $\bar{z}^{\phi} = \left(\bar{z}_0^{\phi}, \bar{z}_1^{\phi}, ...\right)$. Let $\|x\|$ denote the sup norm of x (short for $\|\cdot\|_{\infty}$) and let $\|\phi\|$ be the norm of the linear operator ϕ (see Appendix Chapter A). Then, by the linearity of ϕ, we have that

$$\begin{aligned} \phi(x[T]) = \sum_{t=0}^{T} \phi(\bar{x}_t) &\leq \sum_{t=0}^{T} |\phi(\bar{x}_t)| \\ &= \phi\left(\bar{z}^{\phi}[T]\right) \\ &\leq \|\phi\| \left\|\bar{z}^{\phi}[T]\right\| \\ &= \|\phi\| \|x[T]\| \\ &\leq \|\phi\| \|x\|, \end{aligned}$$

where the first line holds by definition, the second line uses the definition of $\bar{z}^{\phi}$ introduced above, the third line uses the fact that ϕ is a linear functional, the fourth line exploits the fact that the norm $\|\cdot\|$ does not depend on whether the elements are negative or positive, and the final line uses the fact that $\|x\| \geq \|x[T]\|$. This string of relationships implies that the sequence $\left\{\sum_{t=0}^{T} |\phi(\bar{x}_t)|\right\}_{T=1}^{\infty}$ dominates the sequence $\{\phi(x[T])\}_{T=1}^{\infty}$ and is also bounded by $\|\phi\| \|x\|$. Therefore, $\{\phi(x[T])\}_{T=1}^{\infty}$ converges to some well-defined linear functional $\bar{\phi}(x)$. The last inequality above also implies that $\bar{\phi}(x) \leq \|\phi\| \|x\|$, so $\bar{\phi}$ is a bounded and thus continuous linear functional (Theorem A.27 in Appendix Chapter A).

Next, define $\bar{\phi}_J : \mathbb{R}^K \to \mathbb{R}$ as $\bar{\phi}_J : x_J \mapsto \phi(\bar{x}_J)$. Clearly, $\bar{\phi}_J$ is a linear functional (since ϕ is a linear) and moreover, since the domain of $\bar{\phi}_J$ is a Euclidean space, it has an inner product representation, and in particular, there exists $p_J^{**} \in \mathbb{R}^K$ such that

$$\bar{\phi}_J(x_J) = p_J^{**} \cdot x_J \text{ for all } x_J \in \mathbb{R}^K.$$

This also implies that

$$\begin{aligned} \bar{\phi}(x) = \lim_{T\to\infty} \phi(x[T]) &= \lim_{T\to\infty} \sum_{J=0}^{T} \bar{\phi}_J(x_J) \\ &= \lim_{T\to\infty} \sum_{J=0}^{T} p_J^{**} \cdot x_J, \end{aligned}$$

so that $\bar{\phi}$ is a continuous linear functional within inner productive presentation.

To complete this part of the proof, it needs to be shown that $\bar{\phi}(x) = \sum_{j=0}^{\infty} \bar{\phi}_j(x_j)$ can be used instead of ϕ as the continuous linear functional in (5.21). This result follows from the hypotheses that $0 \in X^h$ for each $h \in \mathcal{H}$, and moreover that (i) for each $x, x' \in X^h$ with $u^h(x) > u^h(x')$ for all $h \in \mathcal{H}$, there exists $\bar{T}$ such that $u^h(x[T]) > u^h(x')$ for all $T \geq \bar{T}$ and for all $h \in \mathcal{H}$, and (ii) for each $y \in Y^f$, there exists $\tilde{T}$ such that $y[T] \in Y^f$ for all $T \geq \tilde{T}$ and for all $f \in \mathcal{F}$. In particular, take x^* to be the equivalent of x' in this hypothesis, let $T' = \max\left\{\bar{T}, \tilde{T}\right\}$ and fix $x \in P$. Since x has the property that $u^h\left(x^h\right) > u^h\left(x^{h*}\right)$ for all $h \in \mathcal{H}$, we also have that $u^h\left(x^h[T]\right) > u^h\left(x^{h*}\right)$ for all $h \in \mathcal{H}$ and $T \geq T'$. Therefore,

$$\phi(x^*) \leq \phi(x[T]) \text{ for all } T \geq T'.$$

Now taking the limit $T \to \infty$,

$$\phi(x^*) \leq \overline{\phi}(x) \text{ for all } x \in P.$$

Next, take the sequences $\{x_n\}_{n=1}^{\infty} \in P$, where $x_n = x^* + 1/n$, and $\{y'_n\}_{n=1}^{\infty} \in Y'$, where $y'_n = (1 - 1/n)x^* + y'/n \in Y'$, with $y' \in \text{Int}Y'$. Consequently,

$$\overline{\phi}(x_n) \leq \phi(x^*) \leq \overline{\phi}\left(y'_n\right).$$

Taking limits as $n \to \infty$ and using the continuity of $\overline{\phi}$,

$$\phi(x^*) = \overline{\phi}(x^*).$$

This implies, $\overline{\phi}(x)$ can be used as the continuous linear functional separating P and Y'. Since $\overline{\phi}_j(x_j)$ is a linear functional on $X_j \subset \mathbb{R}_+^K$, it has an inner product representation, $\overline{\phi}_j(x_j) = p_j^* \cdot x_j$ and therefore, so does $\overline{\phi}(x) = \sum_{j=0}^{\infty} \phi_j(x_j) = p^* \cdot x$. Moreover, since each u^h is nondecreasing in its arguments, $p^* \geq 0$ and thus can be interpreted as a price vector (functional).

Parts 1 and 2 have therefore established that there exists a price vector p^* such that conditions 2 and 3 hold. Condition 1 is satisfied by construction. Condition 2 is sufficient to establish that all firms maximize profits at the price vector p^*. To show that all consumers maximize utility at the price vector p^*, use the hypothesis that $p^* \cdot w^{h*} > 0$ for each $h \in \mathcal{H}$. We know from Condition 3 that if $x^h \in X^h$ involves $u^h\left(x^h\right) > u^h\left(x^{h*}\right)$, then $p^* \cdot x^h \geq p^* \cdot w^{h*}$. This implies that if there exists x^h that is strictly preferred to x^{h*} and satisfies $p^* \cdot x^h \leq p^* \cdot w^{h*}$ (which would amount to the consumer not maximizing utility), then there exists $x^h - \varepsilon$ for $\varepsilon > 0$ and small enough, such that $u^h\left(x^h - \varepsilon\right) > u^h\left(x^{h*}\right)$ and $p^* \cdot \left(x^h - \varepsilon\right) < p^* \cdot w^{h*}$, thus violating Condition 3. Therefore, consumers also maximize utility at the price vector p^*, establishing that $(\mathbf{x}^*, \mathbf{y}^*, p^*)$ is a competitive equilibrium. □

5.8. Sequential Trading

A final issue that is useful to discuss at this point relates to *sequential trading*. Standard general equilibrium models assume that all commodities are traded at a given point in time—and once and for all. That is, once trading takes place at the initial date, there is no more trade or production in the economy. This may be a good approximation to reality when different commodities correspond to different physical goods. However, when different commodities

correspond to the same good in different time periods or in different states of nature, trading at a single point is much less reasonable. Models of economic growth typically assume that trading takes place at different points in time. For example, in the Solow growth model of Chapter 2, we envisaged firms hiring capital and labor at each t. Does the presence of sequential trading make any difference to the insights of general equilibrium analysis? If the answer to this question were yes, then the applicability of the lessons from general equilibrium theory to dynamic macroeconomic models would be limited. Fortunately, in the presence of complete markets, sequential trading gives the same result as trading at a single point in time.

More explicitly, the *Arrow-Debreu equilibrium* of a dynamic general equilibrium model involves all the households trading at a single market at time $t = 0$ and purchasing and selling irrevocable claims to commodities indexed by date and state of nature. This means that at time $t = 0$, households agree on all future trades (including trades of goods that are not yet produced). Sequential trading, on the other hand, corresponds to separate markets opening at each t, and households trading labor, capital and consumption goods in each such market at each period. Clearly, both for mathematical convenience and descriptive realism, we would like to think of macroeconomic models as involving sequential trading, with separate markets at each date.

The key result concerning the comparison of models with trading at a single point in time and those with sequential trading is due to Kenneth Arrow (1964). Arrow's focus was on trading across states of nature. However, his results also imply that with complete markets (and time-consistent preferences), trading at a single point in time and sequential trading are equivalent. The easiest way of seeing this is to consider the Arrow securities already discussed in Chapter 2. *(Basic) Arrow Securities* provide an economical means of transferring resources across different dates and different states of nature. Instead of completing all trades at a single point in time, say at time $t = 0$, households can trade Arrow securities and then use these securities to purchase goods at different dates or after different states of nature have been revealed. While Arrow securities are most useful when there is uncertainty as well as a temporal dimension, for our purposes it is sufficient to focus on the transfer of resources across different dates.

The reason why sequential trading with Arrow securities achieves the same result as trading at a single point in time is simple: by the definition of a competitive equilibrium, households correctly anticipate all the prices that they will be facing at different dates (and in different states of nature) and purchase sufficient Arrow securities to cover the expenses that they will incur once the time to trade comes. In other words, instead of buying claims at time $t = 0$ for $x^h_{j,t'}$ units of commodity $j = 1, ..., N$ at date t' at prices $\left(p_{1,t'}, ..., p_{N,t}\right)$, it is sufficient for household h to have an income of $\sum_{j=1}^{N} p_{j,t'} x^h_{j,t'}$ and know that it can purchase as many units of each commodity as it wishes at time t' at the price vector $\left(p_{1,t'}, ..., p_{N,t'}\right)$.

This result can be stated in a slightly more formal manner. Let us consider a dynamic exchange economy running across periods $t = 0, 1, ..., T$, possibly with $T = \infty$ (with the convention that when $T = \infty$, all the summations are assumed to take finite values). Nothing here depends on the assumption that we are focusing on an exchange economy, but suppressing production simplifies notation. Imagine that there are N goods at each date, denoted by $(x_{1,t}, ..., x_{N,t})$, and let the consumption of good j by household h at time t be denoted by $x^h_{j,t}$. Suppose that these goods are perishable, so that they are indeed consumed at time t. Denote the set of households by $\mathcal{H}$ and suppose that each household $h \in \mathcal{H}$ has a vector of endowment $\left(\omega^h_{1,t}, ..., \omega^h_{N,t}\right)$ at time t, and preferences given by the time-separable function of the form

$$\sum_{t=0}^{T} \beta^t_h u^h \left(x^h_{1,t}, ..., x^h_{N,t}\right),$$

for some $\beta_h \in (0, 1)$. These preferences imply that there are no externalities and preferences are time-consistent. I also assume that all markets are open and competitive.

Let an Arrow-Debreu equilibrium be given by $(\mathbf{p}^*, \mathbf{x}^*)$, where $\mathbf{x}^*$ is the complete list of consumption vectors of each household $h \in \mathcal{H}$, that is,

$$\mathbf{x}^* = (x_{1,0}, ...x_{N,0}, ..., x_{1,T}, ...x_{N,T}),$$

with $x_{j,t} = \left\{x^h_{j,t}\right\}_{h\in\mathcal{H}}$ for each j and t, and $\mathbf{p}^*$ is the vector of complete prices $\mathbf{p}^* = \left(p^*_{1,0}, ..., p^*_{N,0}, ..., p_{1,T}, ..., p_{N,T}\right)$, with one of the prices, say $p^*_{1,0}$, chosen as the numeraire, so that $p^*_{1,0} = 1$. In the Arrow-Debreu equilibrium, each outsold purchases and sells claims on each of the commodities, thus engages in trading only at $t = 0$ and chooses an allocation that satisfies the budget constraint

$$\sum_{t=0}^{T}\sum_{j=1}^{N} p^*_{j,t} x^h_{j,t} \leq \sum_{t=0}^{T}\sum_{j=1}^{N} p^*_{j,t}\omega^h_{j,t} \text{ for each } h \in \mathcal{H}.$$

Market clearing then requires

$$\sum_{h\in\mathcal{H}}\sum_{j=1}^{N} x^h_{j,t} \leq \sum_{h\in\mathcal{H}}\sum_{j=1}^{N} \omega^h_{j,t} \text{ for each } j = 1, ..., N \text{ and } t = 0, 1, ..., T.$$

In the equilibrium with sequential trading, markets for goods dated t open at time t. Instead, there are T bonds—*Arrow securities*—that are in zero net supply and can be traded among the households at time $t = 0$. The bond indexed by t pays one unit of one of the goods, say good $j = 1$ at time t. Let the prices of bonds be denoted by $(q_1, ..., q_T)$, again expressed in units of good $j = 1$ (at time $t = 0$). This implies that a household can purchase a unit of bond t at time 0 by paying q_t units of good 1, and in return, it will receive one unit of good 1 at time t (or conversely can sell short one unit of such a bond) The purchase of bond t by household h is denoted by $b^h_t \in \mathbb{R}$, and since each bond is in zero net supply, market clearing requires that

$$\sum_{h\in\mathcal{H}} b^h_t = 0 \text{ for each } t = 0, 1, ..., T.$$

Notice that this specification assumes that there are only T bonds (Arrow securities). More generally, I could have introduced additional bonds, for example, bonds traded at time $t > 0$ for delivery of good 1 at time $t' > t$. This restriction to only T bonds is without loss of any generality (see Exercise 5.10).

Sequential trading corresponds to each individual using their endowment plus (or minus) the proceeds from the corresponding bonds at each date t. Since there is a market for goods at each t, it turns out to be convenient (and possible) to choose a separate numeraire for each date t, and let us again suppose that this numeraire is good 1, so that $p_{1,t}^{**} = 1$ for all t. Therefore, the budget constraint of household $h \in \mathcal{H}$ at time t, given the equilibrium price vector for goods and bonds, $(\mathbf{p}^{**}, \mathbf{q}^{**})$, can be written as:

$$\sum_{j=1}^{N} p_{j,t}^{**} x_{j,t}^{h} \leq \sum_{j=1}^{N} p_{j,t}^{**} \omega_{j,t}^{h} + q_t^{**} b_t^{h} \text{ for } t = 0, 1, ..., T, \tag{5.22}$$

with the normalization that $q_0^{**} = 1$. Let an equilibrium of the sequential trading economy be denoted by $(\mathbf{p}^{**}, \mathbf{q}^{**}, \mathbf{x}^{**}, \mathbf{b}^{**})$, where once again $\mathbf{p}^{**}$ and $\mathbf{x}^{**}$ denote the entire lists of prices and quantities of consumption by each household, and $\mathbf{q}^{**}$ and $\mathbf{b}^{**}$ denote the vector of bond prices and bond purchases by each household. Given this specification, the following theorem can be established.

THEOREM 5.8. ***(Sequential Trading)*** *For the above-described economy, if* $(\mathbf{p}^{*}, \mathbf{x}^{*})$ *is an Arrow-Debreu equilibrium, then there exists a sequential trading equilibrium* $(\mathbf{p}^{**}, \mathbf{q}^{**}, \mathbf{x}^{**}, \mathbf{b}^{**})$, *such that* $\mathbf{x}^{*} = \mathbf{x}^{**}$, $p_{j,t}^{**} = p_{j,t}^{*}/p_{1,t}^{*}$ *for all* j *and* t *and* $q_t^{**} = p_{1,t}^{*}$ *for all* $t > 0$. *Conversely, if* $(\mathbf{p}^{**}, \mathbf{q}^{**}, \mathbf{x}^{**}, \mathbf{b}^{**})$ *is a sequential trading equilibrium, then there exists an Arrow-Debreu equilibrium* $(\mathbf{p}^{*}, \mathbf{x}^{*})$ *with* $\mathbf{x}^{*} = \mathbf{x}^{**}$, $p_{j,t}^{*} = p_{j,t}^{**} p_{1,t}^{*}$ *for all* j *and* t, *and* $p_{1,t}^{*} = q_t^{**}$ *for all* $t > 0$.

PROOF. See Exercise 5.9. □

This theorem implies that all the results concerning Arrow–Debreu equilibria apply to economies with sequential trading. In most of the models studied in this book (except in the context of endogenous financial markets), the focus will be on economies with sequential trading and assume that there exist Arrow securities to transfer resources across dates. These securities might be riskless bonds in zero net supply, or in models without uncertainty, this role will typically be played by the capital stock. I also follow the approach leading to Theorem 5.8 and normalize the price of one good at each date to 1. This implies that in economies with a single consumption good, like the Solow or the neoclassical growth models, the price of the consumption good in each date will be normalized to 1 and the interest rates will directly give the intertemporal relative prices. This is the justification for focusing on interest rates as the key relative prices in macroeconomic (economic growth) models.

5.9. Optimal Growth

Motivated by the discussion in the previous section let us start with an economy characterized by an aggregate production function, and a representative household. The *optimal growth problem* in this context refers to characterizing the allocation of resources that maximizes the utility of the representative household. For example, if the economy consists of a number of identical households, then this corresponds to the Pareto optimal allocation giving the same (Pareto) weight to all households (recall Definition 5.2 introduced above). Therefore, the optimal growth problem in discrete time with no uncertainty, no population growth and no technological progress can be written as follows:

$$\max_{\{c(t),k(t)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u\left(c\left(t\right)\right) \tag{5.23}$$

subject to

$$k\left(t+1\right) = f(k\left(t\right)) + \left(1-\delta\right)k\left(t\right) - c\left(t\right), \tag{5.24}$$

$k\left(t\right) \geq 0$ and given $k\left(0\right) = k_0 > 0$. Here $u : \mathbb{R}_+ \to \mathbb{R}$ is the utility function of the representative household. The objective function represents the discounted sum of this utility function. The constraint (5.24) is also straightforward to understand; total output per capita produced with capital-labor ratio $k\left(t\right)$, $f\left(k\left(t\right)\right)$, together with a fraction $1-\delta$ of the capital that is undepreciated make up the total resources of the economy at date t. Out of this resources $c\left(t\right)$ is spent as consumption per capita $c\left(t\right)$ and the rest becomes next period's capital-labor ratio, $k\left(t+1\right)$.

The optimal growth problem imposes that the social planner chooses an entire sequence of consumption levels and capital stocks, only subject to the resource constraint, (5.24). There are no additional equilibrium constraints. The initial level of capital stock $k\left(0\right) > 0$ has been specified as one boundary condition. But in contrast to the basic Solow model, the solution to this problem involves two, not one, dynamic (difference or differential) equations, and thus necessitates two boundary conditions. The additional boundary condition will not take the form of an initial condition, but will come from the optimality of a dynamic plan in the form of a *transversality condition*. The relevant transversality conditions for this class of problems will be discussed in the next two chapters.

This maximization problem can be solved in a number of different ways, for example, by setting up an infinite-dimensional Lagrangian. But the most convenient and common way of approaching it is by using *dynamic programming.*

It is also useful to note that even if we wished to bypass the Second Welfare Theorem and directly solve for competitive equilibria, we would have to solve a problem similar to the maximization of (5.23) subject to (5.24). In particular, to characterize the equilibrium, we would need to start with the maximizing behavior of households. Since the economy admits a representative household, we only need to look at the maximization problem of this household. Assuming that the representative household has one unit of labor supplied

inelastically and denoting its assets at time t by $a(t)$, this problem can be written as:

$$\max_{\{c(t),a(t)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c(t))$$

subject to some given $a(0) > 0$ and

$$a(t+1) = r(t)[a(t) - c(t) + w(t)], \tag{5.25}$$

where $r(t)$ is the rate of return on assets and $w(t)$ is the equilibrium wage rate (and thus the wage earnings of the representative household). Market clearing then requires that $a(t) = k(t)$. The constraint, (5.25) is the flow budget constraint, meaning that it links tomorrow's assets to today's assets. Here we need an additional condition so that this flow budget constraint eventually converges (so that $a(t)$ should not go to negative infinity). This can be ensured by imposing a lifetime budget constraint. Since a flow budget constraint in the form of (5.25) is both more intuitive and often more convenient to work with, we will not work with the lifetime budget constraint, but augment the flow budget constraint with a limiting condition, which will be introduced and discussed below.

The formulation of the optimal growth problem in continuous time is very similar and takes the form

$$\max_{[c(t),k(t)]_{t=0}^{\infty}} \int_0^{\infty} \exp(-\rho t) u(c(t)) dt \tag{5.26}$$

subject to

$$\dot{k}(t) = f(k(t)) - c(t) - \delta k(t), \tag{5.27}$$

$k(t) \geq 0$ and given $k(0) = k_0 > 0$. The objective function (5.26) is the direct continuous-time analog of (5.23), and (5.27) gives the resource constraint of the economy, similar to (5.24) in discrete time. Once again, this problem lacks one boundary condition which will come from the transversality condition. The most convenient way of characterizing the solution to this problem is via *optimal control theory.* Dynamic programming and optimal control theory will be discussed briefly in the next two chapters.

5.10. Taking Stock

This chapter introduced the preliminaries necessary for an in-depth study of equilibrium and optimal growth theory. At some level it can be thought of as an "odds and ends" chapter, introducing the reader to the notions of representative household, dynamic optimization, welfare theorems and optimal growth. However, the material here is more than odds and ends, since a good understanding of the general equilibrium foundations of economic growth and the welfare theorems is necessary for what is to come in Part 3 below and later.

The most important take-away messages from this chapter are as follows. First, the set of models in this book are examples of more general dynamic general equilibrium models. It is therefore important to understand which features of the growth models are general (in the sense that they do not depend on the specific simplifying assumptions) and which results depend on the further simplifying assumptions. In this respect, the First and the

Second Welfare Theorems are essential. They show that provided that all product and factor markets are competitive and that there are no externalities in production or consumption (and under some relatively mild technical assumptions), dynamic competitive equilibrium will be Pareto optimal and any Pareto optimal allocation can be decentralized as a dynamic competitive equilibrium. These results will be relevant especially in Part 3, where the focus will be on competitive economies. Importantly, these results will not directly apply in our analysis of technological change, where product markets will be monopolistic, or in our study of economic development, where various market imperfections will play an important role.

Second, the most general class of dynamic general equilibrium models will not be tractable enough for us to derive sharp results about the process of economic growth. For this reason, we we will often adopt a range of simplifying assumptions. The most important of those is the representative household assumption, which enables us to model the demand side of the economy as if it were generated by the optimizing behavior of a single household. We saw how this assumption is generally not satisfied, but also how a certain class of preferences, the Gorman preferences, enable us to model economies as if they admit a representative household.

In addition, this chapter introduced the first formulation of the optimal growth problems in discrete and in continuous time. These will be used as examples in the next two chapters.

5.11. References and Literature

This chapter covered a lot of ground and in most cases, many details were omitted for brevity. Many readers will be familiar with some of the material in this chapter. Mas-Colell, Winston and Green (1995) provide an excellent discussion of issues of aggregation. They present results similar to Theorems 5.2 and 5.3 above. The first converse to Theorem 5.2 was stated and proved for additive utility functions by Pollak (1971), but as the discussion following this theorem indicates, the same result applies without the need to assume additivity. Mas-Colell, Winston and Green also present a version of the Debreu-Mantel-Sonnenschein theorem, with a sketch proof and an excellent discussion.

Some basic concepts from microeconomic theory were assumed in the discussion, and the reader can also find a thorough exposition of these in Mas-Colell, Winston and Green (1995). These include Roy's identity, used in Theorem 5.3, and Walras's Law, the concept of numeraire and expected utility theory of von Neumann and Morgenstern, used throughout the analysis.

The Representative Firm Theorem, Theorem 5.4, presented here is quite straightforward, but I am not aware of any other mention or discussion of this theorem in the literature. It is important to distinguish the subject matter of this theorem from the Cambridge controversy in early growth theory, which revolved around the issue of whether different types of capital goods could be aggregated into a single capital index (see, for example, Wan, 1969). The Representative Firm Theorem says nothing about this issue.

The best reference for existence of competitive equilibrium and the welfare theorems with a finite number of consumers and a finite number of commodities is still Debreu's (1959) *Theory of Value*. This short book introduces all of the mathematical tools necessary for general equilibrium theory and gives a very clean exposition. Equally lucid and more modern are the treatments of the same topics in Mas-Colell, Winston and Green (1995) and Bewley (2006). The reader may also wish to consult Mas-Colell, Winston and Green (1995, Chapter 16) for a proof of the Second Welfare Theorem with a finite number of commodities (Theorem 5.7 above is more general because it covers the case with infinite number of commodities). Both of these books also have an excellent discussion of the necessary restrictions on preferences so that they can be represented by utility functions. Mas-Colell, Winston and Green (1995, Chapter 19) also gives a very clear discussion of the role of Arrow securities and the relationship between trading at a single point in time and sequential trading. The classic reference on Arrow securities is Arrow (1964).

Neither of these two references discuss infinite-dimensional economies. The seminal reference for infinite-dimensional welfare theorems is Debreu (1954). Stokey, Lucas and Prescott (1989, Chapter 15) presents existence and welfare theorems for economies with a finite number of consumers and countably infinite number of commodities. The mathematical prerequisites for their treatment are greater than what has been assumed here, but their treatment is both thorough and straightforward to follow once the reader makes the investment in the necessary mathematical techniques. The most accessible references for the Hahn-Banach Theorem, which is necessary for a proof of Theorem 5.7 in infinite-dimensional spaces, are Kolmogorov and Fomin (1970), Kreyszig (1978) and Luenberger (1969). The latter is also an excellent source for all the mathematical techniques used in Stokey, Lucas and Prescott (1989) and also contains much material useful for appreciating continuous-time optimization. Finally, a version of Theorem 5.6 is presented in Bewley (2006), which contains an excellent discussion of overlapping generations models.

5.12. Exercises

Exercise 5.1. Recall that a solution $\{x(t)\}_{t=0}^{T}$ to a dynamic optimization problem is *time-consistent* if the following is true: whenever $\{x(t)\}_{t=0}^{T}$ is an optimal solution starting at time $t=0$, $\{x(t)\}_{t=t'}^{T}$ is an optimal solution to the continuation dynamic optimization problem starting from time $t=t'\in[0,T]$.

(1) Consider the following optimization problem

$$\max_{\{x(t)\}_{t=0}^{T}} \sum_{t=0}^{T} \beta^t u(x(t))$$
$$\text{subject to}$$
$$x(t)\in[0,\bar{x}]$$
$$G(x(0),...,x(T))\leq 0.$$

Although you do not need to, you may assume that G is continuous and convex, and u is continuous and concave.

Prove that any solution $\{x^*(t)\}_{t=0}^T$ to this problem is time-consistent.

(2) Consider the optimization problem

$$\max_{\{x(t)\}_{t=0}^T} u(x(0)) + \delta \sum_{t=1}^T \beta^t u(x(t))$$
$$\text{subject to}$$
$$x(t) \in [0, \bar{x}]$$
$$G(x(0), ..., x(T)) \leq 0.$$

Suppose that the objective function at time $t = 1$ becomes $u(x(1)) + \delta \sum_{t=2}^T \beta^{t-1} u(x(t))$.

Interpret this objective function (sometimes referred to as "hyperbolic discounting").

(3) Let $\{x^*(t)\}_{t=0}^T$ be a solution to this maximization problem. Assume that the individual chooses $x^*(0)$ at $t = 0$, and then is allowed to reoptimize at $t = 1$, that is, he can now solve the problem

$$\max_{\{x(t)\}_{t=1}^T} u(x(1)) + \delta \sum_{t=2}^T \beta^{t-1} u(x(t))$$
$$\text{subject to}$$
$$x(t) \in [0, \bar{x}]$$
$$G(x^*(0), ..., x(T)) \leq 0.$$

Prove that the solution from $t = 1$ onwards, $\{x^{**}(t)\}_{t=1}^T$ is not necessarily the same as $\{x^*(t)\}_{t=1}^T$.

(4) Explain which standard axioms of preferences in basic general equilibrium theory are violated by the preferences in parts 2 and 3 of this exercise.

EXERCISE 5.2. This exercise asks you to work through an example that illustrates the difference between the coefficient of relative risk aversion and the intertemporal elasticity of substitution. Consider a household with the following non-time-separable preferences over consumption levels at two dates:

$$V(c_1, c_2) = \mathbb{E}\left[\left(\frac{c_1^{1-\theta} - 1}{1-\theta}\right)^{\frac{\alpha-1}{\alpha}} + \beta\left(\frac{c_2^{1-\theta} - 1}{1-\theta}\right)^{\frac{\alpha-1}{\alpha}}\right]^{\frac{\alpha}{\alpha-1}},$$

where $\mathbb{E}$ is the expectations operator. The budget constraint of the household is

$$c_1 + \frac{1}{1+r} c_2 \leq W,$$

where r is the interest rate and W is its total wealth, which may be stochastic.

(1) Let us first suppose that W is nonstochastic and equal to $W_0 > 0$. Characterize the utility maximizing choice of c_1 and c_2.

(2) Compute the intertemporal elasticity of substitution.
(3) Now suppose that W is distributed over the support $\left[\underline{W}, \overline{W}\right]$ with some distribution function $G(W)$, where $0 < \underline{W} < \overline{W} < \infty$. Characterize the utility maximizing choice of c_1 and compute the coefficient of relative risk aversion. Provide conditions under which the coefficient of relative risk aversion is the same as the intertemporal elasticity of substitution. Explain why the two differ and interpret the conditions under which they are the same.

EXERCISE 5.3. Prove Theorem 5.2.

EXERCISE 5.4. Prove a version of Theorem 5.3 when there is also production.

EXERCISE 5.5. * Generalize Theorem 5.3 to an economy with a continuum of commodities.

EXERCISE 5.6. (1) Derive the utility maximizing demands for consumers in Example 5.1 and show that the resulting indirect utility function for each consumer is given by (5.5).
(2) Show that maximization of (5.6) leads to the indirect utility function corresponding to the representative household.
(3) Now suppose that $U^h\left(x_1^h, ..., x_N^h\right) = \sum_{j=1}^N \left(x_j^h - \xi_j^h\right)^{\frac{\sigma-1}{\sigma}}$. Repeat the same computations and verify that the resulting indirect utility function is homogeneous of degree 0 in p and y, but does not satisfy the Gorman form. Show, however, that a monotonic transformation of the indirect utility function satisfies the Gorman form. Is this sufficient to ensure that the economy admits a representative household?

EXERCISE 5.7. Construct a continuous-time version of the model with finite lives and random death. In particular suppose that an individual faces a constant (Poisson) flow rate of death equal to $\nu > 0$ and has a true discount factor equal to ρ. Show that this individual will behave as if he is infinitely lived with an effective discount factor of $\rho + \nu$.

EXERCISE 5.8. (1) Will dynastic preferences as those discussed in Section 5.2 lead to infinite-horizon maximization if the instantaneous utility function of future generations are different ($u_t(\cdot)$ potentially different for each generation t)?
(2) How would the results be different if an individual cares about the continuation utility of his offspring with discount factor β, but also cares about the continuation utility of the offspring of his offspring with a smaller discount factor δ?

EXERCISE 5.9. Prove Theorem 5.8.

EXERCISE 5.10. Consider the sequential trading model discussed above and suppose now that households can trade bonds at time t that deliver one unit of good 0 at time t'. Denote the price of such bonds by $q_{t,t'}$.

(1) Rewrite the budget constraint of household h at time t, (5.22), including these bonds.
(2) Prove an equivalent of Theorem 5.8 in this environment with the extended set of bonds.

EXERCISE 5.11. Consider a two-period economy consisting of two types of households. N_A households have the utility function

$$u\left(c_1^h\right)+\beta_A u\left(c_2^h\right),$$

where c_1^h and c_2^h denotes the consumption of household h into two periods. The remaining N_B households have the utility function

$$u\left(c_1^h\right)+\beta_B u\left(c_2^h\right),$$

with $\beta_B<\beta_A$. The two groups have, respectively, incomes y_A and y_B at date 1, and can save this to the second date at some exogenously given gross interest rate $R>0$. Show that for general $u(\cdot)$, this economy does not admit a representative household.

EXERCISE 5.12. Consider an economy consisting of N households each with utility function at time $t=0$ given by

$$\sum_{t=0}^{\infty}\beta^t u\left(c^h(t)\right),$$

with $\beta\in(0,1)$, where $c^h(t)$ denotes the consumption of household h at time t. The economy starts with an endowment of Y units of the final good and has access to no production technology. This endowment can be saved without depreciating or gaining interest rate between periods.

(1) What are the Arrow-Debreu commodities in this economy?
(2) Characterize the set of Pareto optimal allocations of this economy.
(3) Does Theorem 5.7 apply to this economy?
(4) Now consider an allocation of Y units to the households, $\left\{y^h\right\}_{h=1}^N$, such that $\sum_{h=1}^N y^h=Y$. Given this allocation, find the unique competitive equilibrium price vector and the corresponding consumption allocations.
(5) Are all competitive equilibria Pareto optimal?
(6) Now derive a redistribution scheme for decentralizing the entire set of Pareto optimal allocations?

EXERCISE 5.13. (1) Suppose that utility of household h given by $u(x)\equiv\sum_{t=0}^{\infty}\beta^t v^h\left(x^h(t)\right)$, where $x^h(t)\in X\subset\mathbb{R}_+^K$, $v^h:X\to\mathbb{R}$ is continuous, X is compact, and $\beta<1$. Show that the hypothesis that for any $x,x'\in X^h$ with $u^h(x)>u^h(x')$, there exists $\bar{T}$ (as a function of x and x') such that $u^h(x[T])>u^h(x')$ for all $T\geq\bar{T}$ in Theorem 5.7 is satisfied.

(2) Suppose that the production structure is given by a neoclassical production function, where the production vector at time t is only a function of inputs at time t and capital stock chosen at time $t-1$, and that higher capital stalled contributes to greater production and there is free disposal. Show that the second hypothesis in Theorem 5.7 that for each $y\in Y^f$, there exists $\tilde{T}$ such that $y[T]\in Y^f$ for all $T\geq\tilde{T}$ is satisfied.

EXERCISE 5.14. *

(1) Show that Theorem 5.7 does not cover the one good neoclassical growth model with instantaneous preferences given by $u(c) = \left(c^{1-\theta} - 1\right) / (1-\theta)$ with $\theta \geq 1$.
(2) Prove Theorem 5.7 under the assumption that there exists a strictly positive vector $\varepsilon \in \mathbb{R}_+^K$ with each element sufficiently small and $(\varepsilon, ..., \varepsilon) \in X^h$ for all $h \in \mathcal{H}$.
(3) Show that this modified version of Theorem 5.7 covers the economy in part 1 of this exercise.

CHAPTER 6

Infinite-Horizon Optimization and Dynamic Programming

This chapter will provide a brief introduction to infinite-horizon optimization in discrete time under certainty. The main purpose of the chapter is to introduce the reader to infinite-horizon optimization techniques and to *dynamic programming*. These techniques will be used in the rest of the book. Since dynamic programming has become an important tool in many areas of economics and especially in macroeconomics, a great deal of the emphasis will be on dynamic programming methods.

The material in this chapter is presented in four parts. The first part (Sections 6.1-6.3), introduces the problem and provides a number of results necessary for applications of *stationary* dynamic programming techniques in infinite-dimensional optimization problems. Since understanding how these results are derived is often useful in various applications, the second part, in particular, Sections 6.4 and 6.5, will provide additional tools necessary for a more detailed analysis of dynamic programming and the proofs of the main theorems. The material in these two sections is not necessary for the rest of the book, so that those who only wish to acquire a working knowledge of dynamic programming techniques can skip them. The third part then provides a number of results for nonstationary optimization problems (Section 6.7). Finally, the fourth part of the chapter, consisting of Sections 6.6, 6.8 and 6.9) provides a more detailed discussion of how dynamic programming techniques can be used in applications and also presents a number of results on optimal growth using these tools.

Throughout this chapter, the focus is on discounted maximization problems under certainty, similar to the maximization problems introduced in the previous chapter. Dynamic optimization problems under uncertainty are discussed in Chapter 16.

6.1. Discrete-Time Infinite-Horizon Optimization

The canonical discrete-time infinite-horizon optimization program can be written as

$$\begin{aligned}
&\sup_{\{x(t+1)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \tilde{U}(t, x(t), y(t)) \\
&\text{subject to} \\
&y(t) \in \tilde{G}(t, x(t)), \qquad \text{for all } t \geq 0 \\
&x(t+1) = \tilde{f}(t, x(t), y(t)), \qquad \text{for all } t \geq 0 \\
&x(0) \text{ given.}
\end{aligned}$$

Here $\beta \in [0,1)$ is the discount factor, $t = 0, 1, ...$ denotes time, and $x(t) \in X \subset \mathbb{R}^{K_x}$ and $y(t) \in Y \subset \mathbb{R}^{K_y}$ for some $K_x, K_y \geq 1$. We can think of $x(t)$ as the *state variables* (state vector) of the problem, while $y(t)$ denotes the *control variables* (control vector) at time t. The real-valued function

$$\tilde{U} : \mathbb{Z}_+ \times X \times Y \to \mathbb{R}$$

is the instantaneous payoff function of this problem and I have imposed that overall payoff (objective function) is a discounted sum of instantaneous payoffs ($\mathbb{Z}_+$ denotes the nonnegative integers). $\tilde{G}(t,x)$ is a set-valued mapping or a correspondence (see Appendix Chapter A), also written as

$$\tilde{G} : \mathbb{Z}_+ \times X \rightrightarrows Y.$$

Thus the first constraint specifies what values of the control vector $y(t)$ are allowed given the value $x(t)$ at time t. In addition, $\tilde{f} : \mathbb{Z}_+ \times X \times Y \to X$ specifies the evolution of the state vector as a function of last period's state and control vectors. While this formulation is useful in highlighting the distinction between state and control variables, it is often more convenient to work with a transformation of this problem, where we eliminate $y(t)$ entirely and write the problem as:

$$\begin{aligned} &\textbf{Problem A0} \quad : \\ V^*(0, x(0)) \quad &= \quad \sup_{\{x(t+1)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t U(t, x(t), x(t+1)) \\ &\qquad\qquad \text{subject to} \\ x(t+1) \quad &\in \quad G(t, x(t)), \qquad \text{for all } t \geq 0 \\ &\qquad\qquad x(0) \text{ given,} \end{aligned}$$

where $x(t) \in X \subset \mathbb{R}^K$ ($K = K_x$ in terms of the notation above) and now $x(t)$ corresponds to the *state variables* (state vector) while $x(t+1)$ plays the role of the *control variables* (control vector) at time t. In addition, $U : \mathbb{Z}_+ \times X \times X \to \mathbb{R}$ is the instantaneous payoff function with arguments $x(t)$ and $x(t+1)$, as well as time t (instead of having arguments $x(t)$ and $y(t)$) and $G : \mathbb{Z}_+ \times X \rightrightarrows X$ specifies the constraint correspondence. Finally, I have also defined the *value function* $V^* : \mathbb{Z}_+ \times X \to \mathbb{R}$, which specifies the supremum (highest possible value) that the objective function can approach (starting with some $x(t)$ at time t).

In the problem formulation, I used "sup," to denote the supremum of the problem (rather than "max" for the maximum), since there is no guarantee that the maximal value is attained by any feasible plan. However, in all cases studied in this book the maximal value will be attained, so the reader may wish to substitute "max" for "sup". When the maximal value is attained by some sequence $\{x^*(t+1)\}_{t=0}^{\infty} \in X^{\infty}$, I refer to this as a solution or as an *optimal plan* (where X^{∞} is the countable infinite product of the set X, so that an element of X^{∞} is a sequence with each member in X).[1] Of particular importance in this chapter is the function

[1] I will also sometimes use the term solution to denote the function V^* defined in Problems A0 or A1 (or V in Problem A2 below), but the context will make it clear whether the term solution refers to the optimal plan or to the function V^* or V.

$V^{*}(t,x)$, which can be thought of as the *value function*, meaning the value of pursuing the optimal strategy starting with initial state x at time t. Our objective is to characterize the optimal plan $\{x^{*}(t+1)\}_{t=0}^{\infty}$ and the value function $V^{*}(0,x(0))$.

A more general formulation would involve undiscounted objective function, written as

$$\sup_{\{x(t)\}_{t=0}^{\infty}} U(x(0),x(1),...).$$

This added generality is not particularly useful for most of problems in economic growth and the discounted objective function ensures *time-consistency* as discussed in Exercise 5.1, in the sense that at a later date the decision maker will have no reason to revise the optimal plan she chooses at time $t=0$.

Problem A0 is somewhat abstract. However, it has the advantage of being tractable and general enough to nest many interesting economic applications. The next example shows how the canonical optimal growth problem can be put into this language.

Example 6.1. Recall the optimal growth problem introduced in Section 5.9 of the previous chapter,

$$\max_{\{c(t),k(t)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^{t} u(c(t))$$

subject to

$$k(t+1) = f(k(t)) + (1-\delta)k(t) - c(t),$$

$k(t) \geq 0$ and given $k(0) > 0$, with $u : \mathbb{R}_{+} \to \mathbb{R}$. This problem maps into the general formulation here with a simple one-dimensional state and control variables. In particular, let $x(t) = k(t)$ and $x(t+1) = k(t+1)$. Then, use the constraint to write:

$$c(t) = f(k(t)) - k(t+1) + (1-\delta)k(t),$$

and substitute this into the objective function to obtain:

$$\max_{\{k(t+1)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^{t} u(f(k(t)) - k(t+1) + (1-\delta)k(t))$$

subject to $k(t) \geq 0$. Now it can be verified that this problem is a special case of Problem A0 with $U(t,k(t),k(t+1)) = u(f(k(t)) - k(t+1) + (1-\delta)k(t))$ and with the constraint correspondence $G(t,k(t))$ given by $k(t+1) \in [0, f(k(t)) + (1-\delta)k(t)]$ (since $c(t) \geq 0$).

A notable feature emphasized by Example 6.1 is that once the optimal growth problem is formulated in the language of Problem A0, U and G do not explicitly depend on time. This is a fairly common feature. Many interesting problems in economics can be formulated in a *stationary* form, meaning that the objective function is a discounted sum, and U and G do not explicitly depend on time. In the next section, I start with the study of stationary dynamic optimization problems, returning to the more general case of Problem A0 in Section 6.7.

6.2. Introduction to Stationary Dynamic Programming

Let us consider the stationary form of Problem A0,

$$\begin{aligned} &\textbf{Problem A1} \quad : \\ V^{*}(x(0)) &= \sup_{\{x(t+1)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^{t} U(x(t), x(t+1)) \\ &\quad \text{subject to} \\ x(t+1) &\in G(x(t)), \qquad \text{for all } t \geq 0 \\ &\quad x(0) \text{ given}, \end{aligned}$$

where again $\beta \in [0,1)$, and now the constraint correspondence and the instantaneous payoff functions take the form $G : X \rightrightarrows X$ and $U : X \times X \to \mathbb{R}$. This problem is stationary because U and G do not explicitly depend on time. Consequently, I also wrote the value function without the time argument as $V^{*}(x(0))$.

Problem A1, like Problem A0, corresponds to the *sequence problem*; it involves choosing an infinite sequence $\{x(t)\}_{t=0}^{\infty}$ from some (vector) space of infinite sequences (for example, $\{x(t)\}_{t=0}^{\infty} \in X^{\infty} \subset \ell^{\infty}$, where ℓ^{∞} is the vector space of infinite sequences that are bounded with the sup norm, $\|\cdot\|_{\infty}$, which I denote throughout by the simpler notation $\|\cdot\|$, see Appendix Chapter A). Sequence problems sometimes have nice features, but their solutions are often difficult to characterize both analytically and numerically.

The basic idea of dynamic programming is to turn the sequence problem into a *functional equation*. That is, it is to transform the problem into one of finding a function rather than a sequence. The relevant functional equation can be written as follows:

$$\begin{aligned} &\textbf{Problem A2} \quad : \\ V(x) &= \sup_{y \in G(x)} \{U(x,y) + \beta V(y)\}, \text{ for all } x \in X, \end{aligned} \tag{6.1}$$

where $V : X \to \mathbb{R}$ is a real-valued function. Throughout, I use the notation V to refer to the function defined in Problem A2 and V^{*} for the function defined in Problem A1.

Intuitively, in Problem A2, instead of explicitly choosing the sequence $\{x(t)\}_{t=0}^{\infty}$, in (6.1), we choose a *policy*, which determines what the control vector $x(t+1)$ should be for a given value of the state vector $x(t)$. Since the instantaneous payoff function $U(\cdot,\cdot)$ does not depend on time, there is no reason for this policy to be time-dependent either, and I denote the control vector by y and the state vector by x. Then, the problem is equivalent to choosing $y \in G(x) \subset X$ for a given value of $x \in X$. Mathematically, this corresponds to maximizing $V(x)$ for any $x \in X$. The only subtlety in (6.1) is the presence of the $V(\cdot)$ on the right-hand side, which will be explained below. This is also the reason why (6.1) is also called the *recursive formulation*—the function $V(x)$ appears both on the left and the right-hand sides of eq. (6.1) and is thus defined recursively. The functional equation in Problem A2 is also called the *Bellman equation*, after Richard Bellman, who was the first to introduce the dynamic programming formulation.

At first sight, the recursive formulation might not appear as a great advance over the sequence formulation. After all, functions might be trickier to work with than sequences. Nevertheless, it turns out that the functional equation of dynamic programming is easy to work with in many instances. In applied mathematics and engineering, it is favored because it is computationally convenient. In economics, perhaps the major advantage of the recursive formulation is that it often gives better economic insights, similar to the logic of comparing today to tomorrow. In particular, $U(x,y)$ is the" return for today" and $\beta V(y)$ is the continuation return from the next date onwards, equivalent to the "return for tomorrow". Consequently, in many applications we can use our intuitions from two-period maximization or economic problems. Finally, in some special but important cases, the solution to Problem A2 is simpler to characterize analytically than the corresponding solution of the sequence problem, Problem A1.

In fact, the form of Problem A2 suggests itself naturally from the formulation Problem A1. Suppose Problem A1 has a maximum starting at $x(0)$ attained by the optimal sequence $\{x^*(t)\}_{t=0}^{\infty}$ with $x^*(0) = x(0)$. Then, under some relatively weak technical conditions, we have that

$$\begin{aligned} V^*(x(0)) &= \sum_{t=0}^{\infty} \beta^t U(x^*(t), x^*(t+1)) \\ &= U(x(0), x^*(1)) + \beta \sum_{s=0}^{\infty} \beta^s U(x^*(s+1), x^*(s+2)) \\ &= U(x(0), x^*(1)) + \beta V^*(x^*(1)). \end{aligned}$$

This equation encapsulates the basic idea of dynamic programming, *the Principle of Optimality*, which will be stated more formally in Theorem 6.2.

Essentially, an optimal plan can be broken into two parts, what is optimal to do today, and the optimal continuation path. Dynamic programming exploits this principle and provides us with a set of powerful tools to analyze optimization in discrete-time infinite-horizon problems.

As noted above, the particular advantage of this formulation is that the solution can be represented by a time invariant *policy function* (or policy mapping),

$$\pi : X \to X,$$

determining which value of $x(t+1)$ to choose for a given value of the state variable $x(t)$. In general, there will be two complications: first, a control reaching the optimal value may not exist, which was the reason why I originally used the notation sup; second, the solution to Problem A2 may involve not a policy function, but a *policy correspondence*, $\Pi : X \rightrightarrows X$, because there may be more than one maximizer for a given state variable. Let us ignore these complications for now and present a heuristic exposition. These issues will be dealt with in greater detail below.

Once the value function V is determined, the policy function is given straightforwardly. In particular, by definition it must be the case that if optimal policy is given by a policy

function $\pi(x)$, then

$$V(x) = U(x, \pi(x)) + \beta V(\pi(x)), \text{ for all } x \in X,$$

which is one way of determining the policy function. This equation simply follows from the fact that $\pi(x)$ is the optimal policy, so when $y = \pi(x)$, the right-hand side of (6.1) reaches the maximal value $V(x)$.

The usefulness of the recursive formulation in Problem A2 comes from the fact that there are some powerful tools that not only establish the existence of a solution, but also some of its properties. These tools are also utilized directly in a variety of problems in economic growth, macroeconomics and other areas of economic dynamics.

The next section presents results about the relationship between the solution to the sequence problem, Problem A1, and the recursive formulation, Problem A2. More importantly, the tool also present a number of results on the concavity, monotonicity and differentiability of the value function $V(\cdot)$ (and of the value function $V^*(\cdot)$ in Problem A1).

6.3. Stationary Dynamic Programming Theorems

Let us start with a number of assumptions on Problem A1. Since these assumptions are only relevant for this section, I number them separately from the main assumptions used throughout the book. Consider first a sequence $\{x^*(t)\}_{t=0}^{\infty} \in X^{\infty}$ which attains the supremum in Problem A1. Our main purpose is to ensure that this sequence will satisfy the recursive equation of dynamic programming, written here as by substituting in for the sequence $\{x^*(t)\}_{t=0}^{\infty}$ (instead of the "max" operator),

$$V(x^*(t)) = U(x^*(t), x^*(t+1)) + \beta V(x^*(t+1)), \text{ for all } t = 0, 1, ..., \tag{6.2}$$

and that any solution to (6.2) will also be a solution to Problem A1, in the sense that it will attain its supremum. In other words, we are interested in establishing equivalence results between the solutions to Problem A1 and Problem A2.

To prepare for these results, let us define the set of feasible sequences or *plans* starting with an initial value $x(t)$ as:

$$\Phi(x(t)) = \{\{x(s)\}_{s=t}^{\infty} : x(s+1) \in G(x(s)), \text{ for } s = t, t+1, ...\}.$$

Intuitively, $\Phi(x(t))$ is the set of feasible choices of vectors starting from $x(t)$. Let us denote a typical element of the set $\Phi(x(0))$ by $\mathbf{x} = (x(0), x(1), ...) \in \Phi(x(0))$. Our first assumption is:

ASSUMPTION 6.1. $G(x)$ is nonempty for all $x \in X$; and for all $x(0) \in X$ and $\mathbf{x} \in \Phi(x(0))$, $\lim_{n\to\infty} \sum_{t=0}^{n} \beta^t U(x(t), x(t+1))$ exists and is finite.

This assumption is stronger than what is necessary to establish the results that will follow. In particular, for much of the theory of dynamic programming, it is sufficient that the limit in Assumption 6.1 exists. However, in economic applications, we are not interested in optimization problems where households or firms achieve infinite value. This is for two obvious

reasons. First, when some agents can achieve infinite value, the mathematical problems are typically not well defined. Second, the essence of economics, tradeoffs in the face of scarcity, would be absent in these cases. In cases, where households can achieve infinite value, economic analysis is still possible, by using methods sometimes called "overtaking criteria," whereby different sequences that give infinite utility are compared by looking at whether one of them gives higher utility than the other one at each date after some finite threshold. None of the models in this book require these more general optimality concepts.

Assumption 6.2. X is a compact subset of $\mathbb{R}^K$, G is nonempty, compact-valued and continuous. Moreover, $U : \mathbf{X}_G \to \mathbb{R}$ is continuous, where $\mathbf{X}_G = \{(x, y) \in X \times X : y \in G(x)\}$.

This assumption is also natural. We need to impose that $G(x)$ is compact-valued, since optimization problems with choices from non-compact sets are not well behaved (see Appendix Chapter A). In addition, the assumption that U is continuous leads to little loss of generality for most economic applications. In all the models we will encounter in this book, U will be continuous. The most restrictive assumption here is that the state variable lies in a compact set—that X is compact. Most results in this chapter can be generalized to the case in which X is not compact, though this requires additional notation and somewhat more difficult analysis. The case in which X is not compact is important in the analysis of economic growth, since most interesting models of growth will involve the state variable, e.g., the capital stock, growing steadily. Nevertheless, in many cases, with a convenient normalization, the mathematical problem can be turned in to one in which the state variable lies in a compact set. One class of important problems that cannot be treated without allowing for non-compact X are those with endogenous growth. However, since the methods developed in the next chapter do not require this type of compactness assumption and since I will often use continuous-time methods to study endogenous growth models, I simplify the discussion here by assuming that X is compact.

Note also that since X is compact, $G(x)$ is continuous and compact-valued, $\mathbf{X}_G$ is also compact. Since a continuous function on a compact domain is also bounded (Corollary A.1 in Appendix Chapter A), Assumption 6.2 also implies that U is bounded, which will be important for some of the results below.

Many (but not all) economic problems impose additional structure on preferences and technology, for example, in the form of concavity or monotonicity of the instantaneous payoff function U and convexity of the constraints. These properties enable us to establish additional structural results. I now introduce these assumptions formally.

Assumption 6.3. U is concave; for any $\alpha \in (0, 1)$ and any (x, y), $(x', y') \in \mathbf{X}_G$, we have

$$U\left(\alpha x + (1-\alpha)x', \alpha y + (1-\alpha)y'\right) \geq \alpha U(x, y) + (1-\alpha)U(x', y').$$

Moreover, if $x \neq x'$,

$$U\left(\alpha x + (1-\alpha)x', \alpha y + (1-\alpha)y'\right) > \alpha U(x, y) + (1-\alpha)U(x', y').$$

In addition, G is convex; for any $\alpha \in [0,1]$, and $x, x', y, y' \in X$ such that $y \in G(x)$ and $y' \in G(x')$, we have

$$\alpha y + (1-\alpha)y' \in G\left(\alpha x + (1-\alpha)x'\right).$$

This assumption imposes conditions similar to those used in many economic applications: the constraint set is assumed to be convex and the objective function is concave or strictly concave.

The next assumption puts some more structure on the objective function, in particular it ensures that the objective function is increasing in the state variables (its first K arguments), and that greater levels of the state variables are also attractive from the viewpoint of relaxing the constraints; i.e., a greater x means more choice.

ASSUMPTION 6.4. For each $y \in X$, $U(\cdot, y)$ is strictly increasing in each of its first K arguments, and G is monotone in the sense that $x \leq x'$ implies $G(x) \subset G(x')$.

The final assumption is that of differentiability and is also common in most economic models. This assumption will enable us to work with first-order necessary conditions.

ASSUMPTION 6.5. U is continuously differentiable on the interior of its domain $\mathbf{X}_G$.

Given these assumptions, the following sequence of results can be established. The proofs for these results are provided in Section 6.5.

THEOREM 6.1. ***(Equivalence of Values)*** *Suppose Assumptions 6.1 holds. Then, for any $x \in X$, any solution $V^*(x)$ to Problem A1 is also a solution to Problem A2. Moreover, any solution $V(x)$ to Problem A2 is also a solution to Problem A1, so that $V^*(x) = V(x)$ for all $x \in X$.*

Therefore, both the sequence problem and the recursive formulation achieve the same value. The next theorem establishes the more important result that solutions to Problems A1 and A2 coincide.

THEOREM 6.2. ***(Principle of Optimality)*** *Suppose Assumption 6.1 holds. Let $\mathbf{x}^* \in \Phi(x(0))$ be a feasible plan that attains $V^*(x(0))$ in Problem A1. Then,*

$$V^*(x^*(t)) = U(x^*(t), x^*(t+1)) + \beta V^*(x^*(t+1)) \tag{6.3}$$

for $t = 0, 1, ...$ with $x^*(0) = x(0)$.

Moreover, if any $\mathbf{x}^* \in \Phi(x(0))$ satisfies (6.3), then it attains the optimal value in Problem A1.

This theorem is the major conceptual result in the theory of dynamic programming. It states that the returns from an optimal plan (sequence) $\mathbf{x}^* \in \Phi(x(0))$ can be broken into two parts; the current return, $U(x^*(t), x^*(t+1))$, and the continuation return $\beta V^*(x^*(t+1))$, where the continuation return is given by the discounted value of a problem starting from

the state vector from tomorrow onwards, $x^*(t+1)$. In view of the fact that V^* in Problem A1 and V in Problem A2 are identical (cf. Theorem 6.1), (6.3) also implies

$$V(x^*(t)) = U(x^*(t), x^*(t+1)) + \beta V(x^*(t+1)).$$

Notice also that the second part of Theorem 6.2 is equally important. It states that if any feasible plan $\mathbf{x}^*$ starting with $x(0)$, that is, $\mathbf{x}^* \in \Phi(x(0))$, satisfies (6.3), then $\mathbf{x}^*$ attains $V^*(x(0))$. Therefore, this theorem states that we can go from the solution of the recursive problem to the solution of the original problem and vice versa. Consequently, as long as Assumption 6.1 is satisfied there is no risk of excluding solutions when the problem is formulated recursively.

The next results summarize certain important features of the value function V in Problem A2. These results will be useful in characterizing qualitative features of optimal plans in dynamic optimization problems without explicitly finding the solutions.

THEOREM 6.3. ***(Existence of Solutions)*** *Suppose that Assumptions 6.1 and 6.2 hold. Then, there exists a unique continuous and bounded function $V : X \to \mathbb{R}$ that satisfies (6.1). Moreover, an optimal plan $\mathbf{x}^* \in \Phi(x(0))$ exists for any $x(0) \in X$.*

This theorem establishes two major results. The first is the uniqueness of the value function (and hence of the Bellman equation) in dynamic programming problems. Combined with Theorem 6.1, this result naturally implies that an optimal policy function (correspondence) achieving the supremum V^* in Problem A1 exists and also that like V, V^* is continuous and bounded. The optimal plan or the policy function corresponding to Problem A1 or A2 may not be unique, however, even though the value function is unique. This may be the case when two alternative feasible sequences achieve the same maximal value. As in static optimization problems, non-uniqueness of solutions is a consequence of lack of strict concavity of the objective function. When the conditions are strengthened by including Assumption 6.3, uniqueness of the optimal plan is guaranteed.

THEOREM 6.4. ***(Concavity of the Value Function)*** *Suppose that Assumptions 6.1, 6.2 and 6.3 hold. Then, the unique $V : X \to \mathbb{R}$ that satisfies (6.1) is strictly concave.*

In addition, it is straightforward to verify that if Assumption 6.3 is relaxed so that U is concave, then a weaker version of Theorem 6.4 applies and implies that V is concave. Combining the previous two theorems:

COROLLARY 6.1. *Suppose that Assumptions 6.1, 6.2 and 6.3 hold. Then, there exists a unique optimal plan $\mathbf{x}^* \in \Phi(x(0))$ for any $x(0) \in X$. Moreover, the optimal plan can be expressed as $x^*(t+1) = \pi(x^*(t))$, where $\pi : X \to X$ is a continuous policy function.*

The important result in this corollary is that the "policy function" π is indeed a function, not a correspondence. This is a consequence of the fact that x^* is uniquely determined. This result also implies that the policy mapping π is continuous in the state vector. Moreover, if

there exists a vector of parameters $\mathbf{z}$ continuously affecting either the constraint correspondence Φ or the instantaneous payoff function U, then the same argument establishes that π is also continuous in this vector of parameters. This feature will enable qualitative analysis of dynamic macroeconomic models under a variety of circumstances.

Our next result shows that under Assumption 6.4, the value function V is also strictly increasing.

THEOREM 6.5. ***(Monotonicity of the Value Function)*** *Suppose that Assumptions 6.1, 6.2 and 6.4 hold and let $V : X \to \mathbb{R}$ be the unique solution to (6.1). Then, V is strictly increasing in all of its arguments.*

Finally, our purpose in developing the recursive formulation is to use it for characterizing the solution to dynamic optimization problems. As with static optimization problems, this is often made easier by using differential calculus. The difficulty in using differential calculus with (6.1) is that the right-hand side of this expression includes the value function V, which is endogenously determined. We can only use differential calculus if we know from more primitive arguments that this value function is indeed differentiable. The next theorem ensures that this is the case and also provides an expression for the gradient of the value function, which corresponds to a version of the familiar Envelope Theorem (Theorem A.32 in Appendix Chapter A). Recall that Int$X$ denotes the interior of the set X, $D_x f$ denotes the gradient of the function f with respect to the vector x, and Df denotes the gradient of the function f with respect to all of its arguments (see Appendix Chapter A).

THEOREM 6.6. ***(Differentiability of the Value Function)*** *Suppose that Assumptions 6.1, 6.2, 6.3 and 6.5 hold. Let $\pi(\cdot)$ be the policy function defined above and assume that $x \in IntX$ and $\pi(x) \in IntG(x)$, then $V(\cdot)$ is differentiable at x, with gradient given by*

$$DV(x) = D_x U(x, \pi(x)). \tag{6.4}$$

These results will us to use dynamic programming techniques in a wide variety of dynamic optimization problems. Before doing so, I discuss how these results are proved. The next section introduces a number of mathematical tools from basic functional analysis necessary for proving some of these theorems and Section 6.5 provides the proofs of all the results stated in this section.

6.4. The Contraction Mapping Theorem and Applications*

In this section, I present a number of mathematical results that are necessary for making progress with the dynamic programming formulation. In this sense, the current section is a "digression" from the main story line. Therefore, this section and the next can be skipped without interfering with the study of the rest of the book. Nevertheless, the material in this and the next section are useful for a good understanding of foundations of dynamic programming and should enable the reader to achieve a better understanding of these methods. The reader may also wish to consult Appendix Chapter A before reading this section.

Recall from Appendix Chapter A that (S, d) is a metric space, if S is a space and d is a metric defined over this space with the usual properties. The metric is referred to as "d" since it loosely corresponds to the "distance" between two elements of S. A metric space is more general than a finite-dimensional Euclidean space such as a subset of $\mathbb{R}^K$. But as with the Euclidean space, we are most interested in defining mappings from the metric space into itself. These mappings are referred to as *operators* to distinguish them from real-valued functions and are often denoted by the letter T (thus $T : S \to S$) and the standard notation often involves writing Tz for the image of a point $z \in S$ under T (rather than the more familiar $T(z)$). Instead, the notation $T(Z)$ stands for the image of a subset Z of S under T, that is, $T(Z) = \{x \in S\colon \exists z \in Z \text{ with } Tz = x\}$. I use this standard notation here.

Definition 6.1. *Let (S, d) be a metric space and $T : S \to S$ be an operator mapping S into itself. If for some $\beta \in (0, 1)$,*

$$d(Tz_1, Tz_2) \leq \beta d(z_1, z_2), \text{ for all } z_1, z_2 \in S,$$

*then, T is a **contraction mapping** (with modulus β).*

In other words, a contraction mapping brings elements of the space S "uniformly closer" to each other.

Example 6.2. Let us take a simple interval of the real line as our space, $S = [a, b]$, with usual metric of this space $d(z_1, z_2) = |z_1 - z_2|$. Then, $T : S \to S$ is a contraction if for some $\beta \in (0, 1)$,

$$\frac{|Tz_1 - Tz_2|}{|z_1 - z_2|} \leq \beta < 1, \quad \text{all } z_1, z_2 \in S \text{ with } z_1 \neq z_2.$$

Definition 6.2. *A **fixed point** of T is any element of S satisfying $Tz = z$.*

Recall also that a metric space (S, d) is complete if every Cauchy sequence (whose elements are getting closer) in S converges to an element in S (see Appendix Chapter A). Despite its simplicity, the following theorem is one of the most powerful results in functional analysis.

Theorem 6.7. ***(Contraction Mapping Theorem)** Let (S, d) be a complete metric space and suppose that $T : S \to S$ is a contraction. Then, T has a unique fixed point, $\hat{z}$, i.e., there exists a unique $\hat{z} \in S$ such that*

$$T\hat{z} = \hat{z}.$$

Proof. **(Existence)** Note $T^n z = T(T^{n-1}z)$ for any $n = 1, 2, ...$ (with $T^0 z = z$). Choose $z_0 \in S$, and construct a sequence $\{z_n\}_{n=1}^{\infty}$ with each element in S, such that $z_{n+1} = Tz_n$ so that

$$z_n = T^n z_0.$$

Since T is a contraction,

$$d(z_2, z_1) = d(Tz_1, Tz_0) \leq \beta d(z_1, z_0).$$

Repeating this argument

$$d(z_{n+1}, z_n) \leq \beta^n d(z_1, z_0), \quad n = 1, 2, ... \tag{6.5}$$

Hence, for any $m > n$,

$$\begin{aligned} d(z_m, z_n) &\leq d(z_m, z_{m-1}) + ... + d(z_{n+2}, z_{n+1}) + d(z_{n+1}, z_n) \\ &\leq \left(\beta^{m-1} + ... + \beta^{n+1} + \beta^n\right) d(z_1, z_0) \\ &= \beta^n \left(1 + \beta + ...\beta+^{m-n-1}\right) d(z_1, z_0) \\ &\leq \frac{\beta^n}{1-\beta} d(z_1, z_0), \end{aligned} \tag{6.6}$$

where the first inequality uses the triangle inequality (see Appendix Chapter A). The second inequality uses (6.5). The last inequality uses the fact that $1/(1-\beta) = 1 + \beta + \beta^2 + ... > 1 + \beta + ... + \beta^{m-n-1}$.

The string of inequalities in (6.6) imply that for n and m sufficiently large, z_m and z_n will be approaching each other, so that $\{z_n\}_{n=1}^{\infty}$ is a Cauchy sequence. Since S is complete, every Cauchy sequence in S has an limit point in S, therefore:

$$z_n \to \hat{z} \in S.$$

The next step is to show that $\hat{z}$ is a fixed point. Note that for any $z_0 \in S$ and any $n \in \mathbb{N}$, we have

$$\begin{aligned} d(T\hat{z}, \hat{z}) &\leq d(T\hat{z}, T^n z_0) + d(T^n z_0, \hat{z}) \\ &\leq \beta d(\hat{z}, T^{n-1} z_0) + d(T^n z_0, \hat{z}), \end{aligned}$$

where the first relationship again uses the triangle inequality, and the second inequality utilizes the fact that T is a contraction. Since $z_n \to \hat{z}$, both of the terms on the right tend to zero as $n \to \infty$, which implies that $d(T\hat{z}, \hat{z}) = 0$, and therefore $T\hat{z} = \hat{z}$, establishing that $\hat{z}$ is a fixed point.

(Uniqueness) Suppose, to obtain a contradiction, that there exist $\hat{z}, z \in S$, such that $Tz = z$ and $T\hat{z} = \hat{z}$ with $\hat{z} \neq z$. This implies

$$0 < d(\hat{z}, z) = d(T\hat{z}, Tz) \leq \beta d(\hat{z}, z),$$

which delivers a contradiction in view of the fact that $\beta < 1$ and establishes uniqueness. □

The Contraction Mapping Theorem can be used to prove many well-known results. The next example and Exercise 6.4 show how it can be used to prove existence of unique solutions to differential equations. Exercise 6.5 shows how it can be used to prove the Implicit Function Theorem (Theorem A.26 in Appendix Chapter A).

EXAMPLE 6.3. Consider the following one-dimensional differential equation

$$\dot{x}(t) = f(x(t)), \tag{6.7}$$

with a boundary condition $x(0) = c \in \mathbb{R}$. Suppose that $f : [0, \bar{x}] \to \mathbb{R}$ is Lipschitz continuous in the sense that it is continuous and also for some $M < \infty$, it satisfies the following

boundedness condition, $|f(x'') - f(x')| \leq M|x'' - x'|$ for all $x', x'' \in [0, \bar{x}]$. The Contraction Mapping Theorem, Theorem 6.7, can be used to prove the existence of a continuous function $x^*(t)$ that is the unique solution to this differential equation on any compact interval, in particular on $[0, s]$ for some $s \in [0, \bar{x}]$. To do this, consider the space of continuous functions on $[0, s]$, $\mathbf{C}[0, s]$, and define the following operator, T such that for any $g \in \mathbf{C}[0, s]$,

$$Tg(z) = c + \int_0^z f(g(x))\, dx.$$

Notice that T is a mapping from the space of continuous functions on $[0, s]$ into itself, i.e., $T : \mathbf{C}[0, s] \to \mathbf{C}[0, s]$. Moreover, it can be verified T is a contraction for some s. This follows because for any $z \in [0, s]$, we have

$$\left| \int_0^z f(g(x))\, dx - \int_0^z f(\tilde{g}(x))\, dx \right| \leq \int_0^z M|g(x) - \tilde{g}(x)|\, dx \tag{6.8}$$

by the Lipschitz continuity of $f(\cdot)$. This implies that

$$\|Tg(z) - T\tilde{g}(z)\| \leq M \times s \times \|g - \tilde{g}\|,$$

where recall that $\|\cdot\|$ denotes the sup norm, now defined over the space of functions. Choosing $s < 1/M$ establishes that for s sufficiently small, T is indeed a contraction. Then, from Theorem 6.7, there exists a unique fixed point of T over $\mathbf{C}[0, s]$. This fixed point is the unique solution to the differential equation and it is also continuous. Exercise 6.4 asks you to verify some of these steps and shows how the result can be extended so that it applies to $\mathbf{C}[0, \bar{x}]$.

The main use of the Contraction Mapping Theorem for us is that it can be applied to any metric space and thus to the space of functions. Applying it to eq. (6.1) will establish the existence of a unique value function V in Problem A2, greatly facilitating the analysis of such dynamic models. Before doing this, let us consider another useful result. Recall that if (S, d) is a complete metric space and S' is a closed subset of S, then (S', d) is also a complete metric space.

Theorem 6.8. ***(Applications of Contraction Mappings)*** *Let (S, d) be a complete metric space, $T : S \to S$ be a contraction mapping with $T\hat{z} = \hat{z}$.*

(1) *If S' is a closed subset of S, and $T(S') \subset S'$, then $\hat{z} \in S'$.*
(2) *Moreover, if $T(S') \subset S'' \subset S'$, then $\hat{z} \in S''$.*

Proof. Take $z_0 \in S'$, and construct the sequence $\{T^n z_0\}_{n=1}^{\infty}$. Each element of this sequence is in S' since $T(S') \subset S'$. Theorem 6.7 implies that $T^n z_0 \to \hat{z}$. Since S' is closed, $\hat{z} \in S'$, proving part 1 in the theorem.

We know that $\hat{z} \in S'$. Then, the fact that $T(S') \subset S'' \subset S'$ implies that $\hat{z} = T\hat{z} \in T(S') \subset S''$, establishing part 2. □

The second part of this theorem is very important to prove results such as strict concavity or strict monotonicity of a function. This is because the set (space) of strictly concave functions or the set of the strictly increasing functions are not closed (and complete). Therefore,

the Contraction Mapping Theorem cannot be directly applied to these spaces of functions. The second part of this theorem enables us to circumvent this problem.

The previous two theorems show that the contraction mapping property is both simple and powerful. Nevertheless, beyond some simple cases, such as Example 6.2, it is difficult to check whether an operator is indeed a contraction. This may seem particularly difficult in the case of spaces whose elements correspond to functions, which are those that are relevant in the context of dynamic programming. The next theorem provides us with sufficient conditions for an operator to be a contraction that are typically straightforward to check. For this theorem, let us use the following notation: for a real valued function $f(\cdot)$ and some constant $c \in \mathbb{R}$ we define $(f+c)(x) \equiv f(x)+c$. Then:

Theorem 6.9. ***(Blackwell's Sufficient Conditions For a Contraction)*** *Let $X \subseteq \mathbb{R}^K$, and $\mathbf{B}(X)$ be the space of bounded functions $f : X \to \mathbb{R}$ defined on X equipped with the sup norm $\|\cdot\|$. Suppose that $\mathbf{B}'(X) \subset \mathbf{B}(X)$ and let $T : \mathbf{B}'(X) \to \mathbf{B}'(X)$ be an operator satisfying the following two conditions:*

(1) ***(monotonicity)*** *For any $f, g \in \mathbf{B}'(X)$ and $f(x) \leq g(x)$ for all $x \in X$ implies $(Tf)(x) \leq (Tg)(x)$ for all $x \in X$.*

(2) ***(discounting)*** *There exists $\beta \in (0,1)$ such that*

$$[T(f+c)](x) \leq (Tf)(x) + \beta c, \quad \text{for all } f \in B(X),\ c \geq 0 \text{ and } x \in X.$$

Then, T is a contraction with modulus β on $\mathbf{B}'(X)$.

Proof. By the definition of the sup norm, $\|f-g\| = \max_{x \in X} |f(x) - g(x)|$. Then, for any $f, g \in \mathbf{B}'(X) \subset \mathbf{B}(X)$,

$$\begin{aligned} f(x) &\leq g(x) + \|f-g\| && \text{for any } x \in X, \\ (Tf)(x) &\leq T[g + \|f-g\|](x) && \text{for any } x \in X, \\ (Tf)(x) &\leq (Tg)(x) + \beta \|f-g\| && \text{for any } x \in X, \end{aligned} \tag{6.9}$$

where the second line applies the operator T on both sides and uses monotonicity, and the third line uses discounting (together with the fact that $\|f-g\|$ is simply a number). By the converse argument,

$$\begin{aligned} g(x) &\leq f(x) + \|g-f\| && \text{for any } x \in X, \\ (Tg)(x) &\leq T[f + \|g-f\|](x) && \text{for any } x \in X, \\ (Tg)(x) &\leq (Tf)(x) + \beta \|g-f\| && \text{for any } x \in X. \end{aligned} \tag{6.10}$$

Combining the inequalities (6.9) and (6.10),

$$\|Tf - Tg\| \leq \beta \|f-g\|,$$

proving that T is a contraction on $\mathbf{B}'(X)$. □

We will see that Blackwell's sufficient conditions are straightforward to check in many economic applications, including the models of optimal or equilibrium growth.

6.5. Proofs of the Main Dynamic Programming Theorems*

I now provide proofs of Theorems 6.1-6.6. The first step is a straightforward lemma, which will be useful in these proofs. For a feasible infinite sequence $\mathbf{x} = (x(0), x(1), ...) \in \Phi(x(0))$ starting at $x(0)$, let

$$\bar{\mathbf{U}}(\mathbf{x}) \equiv \sum_{t=0}^{\infty} \beta^t U(x(t), x(t+1))$$

be the value of choosing this potentially non-optimal infinite feasible sequence. In view of Assumption 6.1, $\bar{\mathbf{U}}(\mathbf{x})$ exists and is finite. The next lemma shows that $\bar{\mathbf{U}}(\mathbf{x})$ can be separated into two parts, the current return and the continuation return.

LEMMA 6.1. *Suppose that Assumption 6.1 holds. Then, for any $x(0) \in X$ and any $\mathbf{x} \in \Phi(x(0))$,*

$$\bar{\mathbf{U}}(\mathbf{x}) = U(x(0), x(1)) + \beta \bar{\mathbf{U}}(\mathbf{x}')$$

where $\mathbf{x}' = (x(1), x(2), ...)$.

PROOF. Since under Assumption 6.1 $\bar{\mathbf{U}}(\mathbf{x})$ exists and is finite, it can be written as

$$\begin{aligned}
\bar{\mathbf{U}}(\mathbf{x}) &= \sum_{t=0}^{\infty} \beta^t U(x(t), x(t+1)) \\
&= U(x(0), x(1)) + \beta \sum_{s=0}^{\infty} \beta^s U(x(s+1), x(s+2)) \\
&= U(x(0), x(1)) + \beta \bar{\mathbf{U}}(\mathbf{x}')
\end{aligned}$$

as defined in the lemma. □

Before providing a proof of Theorem 6.1, it is useful to be more explicit about what it means for V and V^* to be solutions to Problems A1 and A2. Let us start with Problem A1. Using the notation introduced in this section, for any $x(0) \in X$,

$$V^*(x(0)) = \sup_{\mathbf{x} \in \Phi(x(0))} \bar{\mathbf{U}}(\mathbf{x}). \tag{6.11}$$

In view of Assumption 6.1, which ensures that all values are bounded, (6.11) implies

$$V^*(x(0)) \geq \bar{\mathbf{U}}(\mathbf{x}) \text{ for all } \mathbf{x} \in \Phi(x(0)), \tag{6.12}$$

since no other feasible sequence of choices can give higher value than the supremum, $V^*(x(0))$. However, if some function $\tilde{V}$ satisfies condition (6.12), so will $\alpha\tilde{V}$ for $\alpha > 1$ (provided that $\tilde{V}(x(0)) \geq 0$). Therefore, this condition is not sufficient. In addition, we also require that

$$\text{for any } \varepsilon > 0, \text{ there exists } \mathbf{x}' \in \Phi(x(0)) \text{ s.t. } V^*(x(0)) \leq \bar{\mathbf{U}}(\mathbf{x}') + \varepsilon, \tag{6.13}$$

The conditions for $V(\cdot)$ to be a solution to Problem A2 are similar. For any $x(0) \in X$,

$$V(x(0)) \geq U(x(0), y) + \beta V(y), \quad \text{all } y \in G(x(0)), \tag{6.14}$$

and

(6.15) for any $\varepsilon > 0$, there exists $y' \in G(x(0))$ s.t. $V(x(0)) \leq U(x(0), y') + \beta V(y) + \varepsilon$.

Proof of Theorem 6.1.If $\beta = 0$, Problems A1 and A2 are identical, thus the result follows immediately. Suppose that $\beta > 0$ and take arbitrary $x(0) \in X$ and $x(1) \in G(x(0))$. By (6.13) for any $\varepsilon > 0$, there exists $\mathbf{x}'_\varepsilon \in \Phi(x(1))$ such that $\bar{\mathbf{U}}(\mathbf{x}'_\varepsilon) \geq V^*(x(1)) - \varepsilon$. Moreover, by (6.12), for any $\mathbf{x} = (x(0), x(1), ...) \in \Phi(x(0))$, and thus in particular for $\mathbf{x}_\varepsilon = (x(0), \mathbf{x}'_\varepsilon)$, $\bar{\mathbf{U}}(\mathbf{x}_\varepsilon) \leq V^*(x(0))$. Lemma 6.1 then implies

$$\begin{aligned} V^*(x(0)) &\geq U(x(0), x(1)) + \beta \bar{\mathbf{U}}(\mathbf{x}'_\varepsilon), \\ &\geq U(x(0), x(1)) + \beta V^*(x(1)) - \beta\varepsilon. \end{aligned}$$

Since $\varepsilon > 0$ is arbitrary, this implies

$$V^*(x(0)) \geq U(x(0), x(1)) + \beta V^*(x(1)),$$

and thus $V^*(\cdot)$ satisfies (6.14).

Next, take an arbitrary $\varepsilon > 0$. By (6.13), there exists $\mathbf{x}'_\varepsilon = (x(0), x'_\varepsilon(1), x'_\varepsilon(2), ...) \in \Phi(x(0))$ such that

$$\bar{\mathbf{U}}(\mathbf{x}'_\varepsilon) \geq V^*(x(0)) - \varepsilon.$$

Now since $\mathbf{x}''_\varepsilon = (x'_\varepsilon(1), x'_\varepsilon(2), ...) \in \Phi(x'_\varepsilon(1))$ and $V^*(x'_\varepsilon(1))$ is the supremum in Problem A1 starting with $x'_\varepsilon(1)$, Lemma 6.1 implies

$$\begin{aligned} U(x(0), x'_\varepsilon(1)) + \beta \bar{U}(\mathbf{x}''_\varepsilon) &\geq V^*(x(0)) - \varepsilon \\ U(x(0), x'_\varepsilon(1)) + \beta V^*(x'_\varepsilon(1)) &\geq V^*(x(0)) - \varepsilon. \end{aligned}$$

The last inequality implies that $V^*(\cdot)$ satisfies (6.15) since $x'_\varepsilon(1) \in G(x(0))$ for any $\varepsilon > 0$. This proves that any solution to Problem A1 satisfies (6.14) and (6.15), and is thus a solution to Problem A2.

To establish the reverse, note that (6.14) implies that for any $x(1) \in G(x(0))$,

$$V(x(0)) \geq U(x(0), x(1)) + \beta V(x(1)).$$

Now substituting recursively for $V(x(1))$, $V(x(2))$, and so on, and defining $\mathbf{x} = (x(0), x(1), ...)$,

$$V(x(0)) \geq \sum_{t=0}^{n} U(x(t), x(t+1)) + \beta^{n+1} V(x(n+1)).$$

Moreover, $\lim_{n\to\infty} \sum_{t=0}^{n} \beta^t U(x(t), x(t+1)) = \bar{\mathbf{U}}(\mathbf{x})$ and $\lim_{n\to\infty} \beta^{n+1} V(x(n+1)) = \lim_{n\to\infty} \left[\beta^{n+1} \lim_{m\to\infty} \sum_{t=n}^{m} \beta^t U(x(t), x(t+1))\right] = 0$ (since Assumption 6.1 ensures that $\lim_{m\to\infty} \sum_{t=n}^{m} \beta^t U(x(t), x(t+1)) = 0$). Therefore,

$$V(x(0)) \geq \bar{\mathbf{U}}(\mathbf{x}),$$

for any $\mathbf{x} \in \Phi(x(0))$, thus $V(\cdot)$ satisfies (6.12).

Next, let $\varepsilon > 0$ be a positive number. From (6.15), for any $\varepsilon' = \varepsilon (1-\beta) > 0$ there exists $x_\varepsilon (1) \in G(x(0))$ such that

$$V(x(0)) \leq U(x(0), x_\varepsilon(1)) + \beta V(x_\varepsilon(1)) + \varepsilon'.$$

Let $x_\varepsilon (t) \in G(x(t-1))$, with $x_\varepsilon (0) = x(0)$, and define $\mathbf{x}_\varepsilon \equiv (x(0), x_\varepsilon(1), x_\varepsilon(2), ...)$. Again substituting recursively for $V(x_\varepsilon(1))$, $V(x_\varepsilon(2))$,...,

$$\begin{aligned} V(x(0)) &\leq \sum_{t=0}^{n} U(x_\varepsilon(t), x_\varepsilon(t+1)) + \beta^{n+1} V(x(n+1)) + \varepsilon' + \varepsilon'\beta + ... + \varepsilon'\beta^n \\ &\leq \bar{\mathbf{U}}(\mathbf{x}_\varepsilon) + \varepsilon, \end{aligned}$$

where the last step follows using the definition of ε (in particular, $\varepsilon = \varepsilon' \sum_{t=0}^{\infty} \beta^t$) and because $\lim_{n\to\infty} \sum_{t=0}^{n} U(x_\varepsilon(t), x_\varepsilon(t+1)) = \bar{\mathbf{U}}(\mathbf{x}_\varepsilon)$. This establishes that $V(\cdot)$ also satisfies (6.13) and completes the proof. □

In economic problems, we are often interested not in the maximal value of the program but in the optimal plans that achieve this maximal value. Recall that the question of whether the optimal path resulting from Problems A1 and A2 are equivalent was addressed by Theorem 6.2. I now provide a proof of this theorem.

PROOF OF THEOREM 6.2. By hypothesis $\mathbf{x}^* \equiv (x(0), x^*(1), x^*(2), ...)$ is a solution to Problem A1, that is, it attains the supremum, $V^*(x(0))$, starting from $x(0)$. Let $\mathbf{x}_t^* \equiv (x^*(t), x^*(t+1), ...)$.

The first step is to show that for any $t \geq 0$, $\mathbf{x}_t^*$ attains the supremum starting from $x^*(t)$, so that

$$\bar{\mathbf{U}}(\mathbf{x}_t^*) = V^*(x^*(t)). \tag{6.16}$$

The proof is by induction. The base step of induction, for $t = 0$, is straightforward, since, by definition, $\mathbf{x}_0^* = \mathbf{x}^*$ attains $V^*(x(0))$.

Next suppose that the statement is true for t, so that (6.16) is true for t, and we establish that it is true for $t+1$. Equation (6.16) implies that

$$\begin{aligned} V^*(x^*(t)) &= \bar{\mathbf{U}}(\mathbf{x}_t^*) \\ &= U(x^*(t), x^*(t+1)) + \beta \bar{\mathbf{U}}(\mathbf{x}_{t+1}^*). \end{aligned} \tag{6.17}$$

Let $\mathbf{x}_{t+1} = (x^*(t+1), x^*(t+2), ...) \in \Phi(x^*(t+1))$ be any feasible plan starting with $x^*(t+1)$. By definition, $\mathbf{x}_t = (x^*(t), \mathbf{x}_{t+1}) \in \Phi(x^*(t))$. Since $V^*(x^*(t))$ is the supremum starting with $x^*(t)$,

$$\begin{aligned} V^*(x^*(t)) &\geq \bar{\mathbf{U}}(\mathbf{x}_t) \\ &= U(x^*(t), x^*(t+1)) + \beta \bar{\mathbf{U}}(\mathbf{x}_{t+1}). \end{aligned}$$

Combining this inequality with (6.17),

$$\bar{\mathbf{U}}(\mathbf{x}_{t+1}^*) \geq \bar{\mathbf{U}}(\mathbf{x}_{t+1})$$

for all $\mathbf{x}_{t+1} \in \Phi(x^*(t+1))$. This establishes that $\mathbf{x}^*_{t+1}$ attains the supremum starting from $x^*(t+1)$ and completes the induction step, proving that eq. (6.16) holds for all $t \geq 0$.

Equation (6.16) then implies that

$$\begin{aligned} V^*(x^*(t)) &= \bar{\mathbf{U}}(\mathbf{x}^*_t) \\ &= U(x^*(t), x^*(t+1)) + \beta \bar{\mathbf{U}}(\mathbf{x}^*_{t+1}) \\ &= U(x^*(t), x^*(t+1)) + \beta V^*(x^*(t+1)), \end{aligned}$$

establishing (6.3) and thus completing the proof of the first part of the theorem.

Now suppose that (6.3) holds for $\mathbf{x}^* \in \Phi(x(0))$. Then, substituting repeatedly for $\mathbf{x}^*$,

$$V^*(x(0)) = \sum_{t=0}^{n} \beta^t U(x^*(t), x^*(t+1)) + \beta^{n+1} V^*(x(n+1)).$$

In view of the fact that $V^*(\cdot)$ is bounded (cf. Assumption 6.1),

$$\begin{aligned} \bar{\mathbf{U}}(\mathbf{x}^*) &= \lim_{n \to \infty} \sum_{t=0}^{n} \beta^t U(x^*(t), x^*(t+1)) \\ &= V^*(x(0)), \end{aligned}$$

thus $\mathbf{x}^*$ attains the optimal value in Problem A1, completing the proof of the second part of the theorem. □

These two theorems establish that under Assumptions 6.1 and 6.2, we can freely interchange Problems A1 and A2. Our next task is to prove that a policy achieving the optimal path exists for both problems. I will provide two alternative proofs for this to show how this conclusion can be reached either by looking at Problem A1 or at Problem A2, and then exploiting their equivalence. The first proof is more abstract and works directly on the sequence problem, Problem A1. This method of proof will be particularly useful when we look at nonstationary problems.

PROOF OF THEOREM 6.3. *(Version 1)* Consider Problem A1. The objective function in Problem A1 is continuous in the product topology in view of Assumptions 6.1 and 6.2 (see Theorem A.12 in Appendix Chapter A). Moreover, the constraint set $\Phi(x(0))$ is a closed subset of X^∞ (infinite product of X). Since X is compact (Assumption 6.2), Tychonoff's Theorem, Theorem A.13, implies that X^∞ is compact in the product topology. A closed subset of a compact set is compact (Fact A.2 in Appendix Chapter A), which implies that $\Phi(x(0))$ is compact. From Weierstrass' Theorem (Theorem A.9) applied to Problem A1, there exists $\mathbf{x} \in \Phi(x(0))$ attaining $V^*(x(0))$. Moreover, the constraint set is a continuous correspondence (again in the product topology), so Berge's Maximum Theorem, Theorem A.16, implies that $V^*(x(0))$ is continuous. Since $x(0) \in X$ and X is compact, this implies that $V^*(x(0))$ is bounded (Corollary A.1 in Appendix Chapter A). □

PROOF OF THEOREM 6.3. *(Version 2)* Let $\mathbf{C}(X)$ be the set of continuous functions defined on X, endowed with the sup norm, $\|f\| = \sup_{x \in X} |f(x)|$. In view of Assumption 6.2,

the relevant set X is compact and therefore all functions in $\mathbf{C}(X)$ are bounded since they are continuous and X is compact (Corollary A.1 in Chapter A). For $V \in \mathbf{C}(X)$, define the operator T as

$$TV(x) = \max_{y \in G(x)} \{U(x,y) + \beta V(y)\}. \tag{6.18}$$

A fixed point of this operator, $V = TV$, will be a solution to Problem A2. We first prove that such a fixed point (solution) exists. The maximization problem on the right-hand side of (6.18) is one of maximizing a continuous function over a compact set, and by Weierstrass's Theorem, Theorem A.9, it has a solution. Consequently, T is well-defined. Moreover, because $G(x)$ is a nonempty and continuous correspondence by Assumption 6.1, and $U(x,y)$ and $V(y)$ are continuous by hypothesis, Berge's Maximum Theorem, Theorem A.16, implies that

$$\max_{y \in G(x)} \{U(x,y) + \beta V(y)\}$$

is continuous in x, thus $TV(x) \in \mathbf{C}(X)$ and T maps $\mathbf{C}(X)$ into itself.

Next, it is also straightforward to verify that T satisfies Blackwell's sufficient conditions for a contraction in Theorem 6.9 (see Exercise 6.6). Therefore, applying Theorem 6.7, a unique fixed point $V \in \mathbf{C}(X)$ to (6.18) exists and this is also the unique solution to Problem A2.

Now consider the maximization in Problem A2. Since U and V are continuous and $G(x)$ is compact-valued, another application of Weierstrass's Theorem implies that $y \in G(x)$ achieving the maximum exists. This defines the set of maximizers $\Pi(x)$ for Problem A2, that is,

$$\Pi(x) = \arg\max_{y \in G(x)} \{U(x,y) + \beta V(y)\}. \tag{6.19}$$

Let $\mathbf{x}^* = (x(0), x^*(1), ...)$ with $x^*(t+1) \in \Pi(x^*(t))$ for all $t \geq 0$. Then, from Theorems 6.1 and 6.2, $\mathbf{x}^*$ is also an optimal plan for Problem A1. □

An additional result that follows from the second version of the theorem (which can also be derived from version 1, but would require more work), concerns the properties of the set of maximizers $\Pi(x)$ (or equivalently the correspondence $\Pi : X \rightrightarrows X$), defined in (6.19). An immediate application of the Berge's Maximum Theorem, Theorem A.16, implies that Π is a upper hemi-continuous and compact-valued correspondence. This observation will be used in the proof of Corollary 6.1. Before turning to this corollary, I provide a proof of Theorem 6.4. This proof also shows how the Contraction Mapping Theorem, Theorem 6.8, is often used in the study of dynamic optimization problems.

PROOF OF THEOREM 6.4. Recall that $\mathbf{C}(X)$ is the set of continuous (and bounded) functions over the compact set X. Let $\mathbf{C}'(X) \subset \mathbf{C}(X)$ be the set of bounded, continuous, (weakly) concave functions on X, and let $\mathbf{C}''(X) \subset \mathbf{C}'(X)$ be the set of strictly concave functions. Clearly, $\mathbf{C}'(X)$ is a closed subset of the complete metric space $\mathbf{C}(X)$, but $\mathbf{C}''(X)$ is not a closed subset. Let T be as defined in (6.18). Since T is a contraction, it has a unique

fixed point in $\mathbf{C}(X)$. By Theorem 6.8, proving that $T[\mathbf{C}'(X)] \subset \mathbf{C}''(X) \subset \mathbf{C}'(X)$ would be sufficient to establish that this unique fixed point is in $\mathbf{C}''(X)$ and hence the value function is strictly concave. Let $V \in \mathbf{C}'(X)$ and for $x' \neq x''$ and $\alpha \in (0,1)$, let

$$x_\alpha \equiv \alpha x' + (1-\alpha)x''.$$

Let $y' \in G(x')$ and $y'' \in G(x'')$ be solutions to Problem A2 with state vectors x' and x''. This implies that

$$\begin{aligned} TV(x') &= U\left(x', y'\right) + \beta V\left(y'\right) \text{ and} \\ TV(x'') &= U\left(x'', y''\right) + \beta V\left(y''\right). \end{aligned} \tag{6.20}$$

In view of Assumption 6.3 (that G is convex) $y_\alpha \equiv \alpha y' + (1-\alpha)\, y'' \in G(x_\alpha)$, so that

$$\begin{aligned} TV(x_\alpha) &\geq U(x_\alpha, y_\alpha) + \beta V(y_\alpha), \\ &> \alpha\left[U\left(x', y'\right) + \beta V\left(y'\right)\right] \\ &\quad +(1-\alpha)[U\left(x'', y''\right) + \beta V\left(y''\right)] \\ &= \alpha TV(x') + (1-\alpha)TV(x''), \end{aligned}$$

where the first line follows by the fact that $y_\alpha \in G(x_\alpha)$ is not necessarily the maximizer starting with state x_α. The second line uses Assumption 6.3 (strict concavity of U in the concavity of V), and the third line is simply the definition introduced in (6.20). This argument implies that for any $V \in \mathbf{C}'(X)$, TV is strictly concave, thus $T[\mathbf{C}'(X)] \subset \mathbf{C}''(X)$. Then, Theorem 6.8 implies that the unique fixed point V^* is in $\mathbf{C}''(X)$, and hence it is strictly concave. □

PROOF OF COROLLARY 6.1. Assumption 6.3 implies that $U(x,y)$ is concave in y, and under this assumption, Theorem 6.4 established that $V(y)$ is strictly concave in y. The sum of a concave function and a strictly concave function is strictly concave, thus the right-hand side of Problem A2 is strictly concave in y. Therefore, combined with the fact that $G(x)$ is convex for each $x \in X$ (again Assumption 6.3), there exists a unique maximizer $y \in G(x)$ for each $x \in X$. This implies that the policy correspondence $\Pi(x)$ is single-valued, thus a function, and can thus be expressed as $\pi(x)$. Since $\Pi(x)$ is upper hemi-continuous as observed above, so is $\pi(x)$. Since a single-valued upper hemi-continuous correspondence is a continuous function, the corollary follows. □

PROOF OF THEOREM 6.5. The proof again follows from Theorem 6.8. Let $\mathbf{C}'(X) \subset \mathbf{C}(X)$ be the set of bounded, continuous, nondecreasing functions on X, and let $\mathbf{C}''(X) \subset \mathbf{C}'(X)$ be the set of strictly increasing functions. Since $\mathbf{C}'(X)$ is a closed subset of the complete metric space $\mathbf{C}(X)$, Theorem 6.8 implies that if $T[\mathbf{C}'(X)] \subset \mathbf{C}''(X)$, then the fixed point to (6.18), i.e., V, is in $\mathbf{C}''(X)$, and therefore, it is a strictly increasing function. To see that this is the case, consider any $V \in \mathbf{C}'(X)$, i.e., any nondecreasing function. In view of Assumption 6.4, $\max_{y \in G(x)} \{U(x,y) + \beta V(y)\}$ is strictly increasing. This establishes that $TV \in \mathbf{C}''(X)$ and completes the proof. □

For Theorem 6.6, I also provide two proofs. The first one is simpler and only establishes that $V(x)$ has a well-defined matrix of partial derivatives, the Jacobian, $DV(x)$. The second proof appeals to a powerful result from convex analysis and establishes that $V(x)$ is differentiable (see Appendix Chapter A for examples of functions that may possess well-defined Jacobians but may not be differentiable). Moreover, the second proof is more widely known and used in the literature. For this proof, recall that $\varepsilon \downarrow 0$ implies that ε is a positive decreasing sequence approaching 0 and $\varepsilon \uparrow 0$ that ε is a negative increasing sequence approaching 0.

PROOF OF THEOREM 6.6. *(Version 1)* From Corollary 6.1, $\Pi(x)$ is single-valued, thus a function that can be represented by $\pi(x)$. By hypothesis $\pi(x) = x' \in \text{Int}G(x)$. Now consider the initial value $x + \tilde{\varepsilon}_K$, where $\tilde{\varepsilon}_K$ is a K-dimensional vector with the first element equal to the number $\varepsilon > 0$, which will be taken to be small, and the remaining elements equal to 0. By definition,

$$\begin{aligned} V^*(x) &= V(x) \\ &= \bar{\mathbf{U}}(\mathbf{x}^*), \end{aligned}$$

where $\mathbf{x}^* = (x, x', ...) \in \Phi(x^*(0))$ is the optimal plan Now consider the following plan $\mathbf{x}_\varepsilon = (x + \tilde{\varepsilon}_K, x', ...)$ starting with initial value $x + \tilde{\varepsilon}_K$. For ε sufficiently small, this plan is feasible, i.e., $\mathbf{x}_\varepsilon \in \Phi(x + \tilde{\varepsilon}_K)$. To see this, observe that $x' \in \text{Int}G(x)$ and from Assumption 6.2 G is continuous, thus $x' \in \text{Int}G(x + \tilde{\varepsilon}_K)$. Using Lemma 6.1 and the fact that $V^*(x) = V(x)$ for all x, the value of this plan can be written as

$$V(x + \tilde{\varepsilon}_K) \geq U\left(x + \tilde{\varepsilon}_K, x'\right) + \beta V\left(x'\right).$$

Therefore,

$$V(x + \tilde{\varepsilon}_K) - V(x) \geq U\left(x + \tilde{\varepsilon}_K, x'\right) - U\left(x, x'\right).$$

Since by Assumption 6.5 U is differentiable,

$$V_1^+(x) \equiv \lim_{\varepsilon \downarrow 0} \frac{V(x + \tilde{\varepsilon}_K) - V(x)}{\varepsilon} \geq U_1\left(x, x'\right), \tag{6.21}$$

where U_1 denotes the partial derivative of the U function with respect to the first element of the x vector. Next, consider the initial value $x - \tilde{\varepsilon}_K$. Again since $x' \in \text{Int}G(x)$ and G is continuous, $(x - \tilde{\varepsilon}_K, x', ...) \in \Phi(x - \tilde{\varepsilon}_K)$. The same argument then establishes that

$$V(x - \tilde{\varepsilon}_K) - V(x) \geq U\left(x - \tilde{\varepsilon}_K, x'\right) - U\left(x, x'\right),$$

and thus dividing by $-\varepsilon$,

$$V_1^-(x) \equiv \lim_{\varepsilon \uparrow 0} \frac{V(x + \tilde{\varepsilon}_K) - V(x)}{\varepsilon} \leq U_1\left(x, x'\right). \tag{6.22}$$

The same argument clearly applies to any partial derivative, thus

$$V_k^-(x) \leq U_k\left(x, x'\right) \leq V_k^+(x)$$

for each $k = 1, 2, ..., K$. Since from Theorem 6.4, V is concave, $V_k^-(x) \geq V_k^+(x)$ for each k. This establishes that $V_k^+(x) = V_k^-(x) = U_k(x, x')$, and thus $DV(x) = D_x U(x, \pi(x))$, completing the proof. □

PROOF OF THEOREM 6.6. *(Version 2)* From Corollary 6.1, $\Pi(x)$ is single-valued, thus a function that can be represented by $\pi(x)$. By hypothesis, $\pi(x) \in \text{Int}G(x)$ and from Assumption 6.2 G is continuous. Therefore, there exists a neighborhood $\mathcal{N}(x)$ of x such that $\pi(x) \in \text{Int}G(x)$, for all $x \in \mathcal{N}(x)$. Define $W(\cdot)$ on $\mathcal{N}(x)$ by

$$W(x) = U(x, \pi(x)) + \beta V(\pi(x)).$$

In view of Assumptions 6.3 and 6.5, the fact that $V(\pi(x))$ is a number (independent of x), and the fact that U is concave and differentiable, $W(\cdot)$ is concave and differentiable. Moreover, since $\pi(x) \in G(x)$ for all $x \in \mathcal{N}(x)$, it follows that

$$W(x) \leq \max_{y \in G(x)} \{U(x, y) + \beta V(y)\} = V(x), \qquad \text{for all } x \in \mathcal{N}(x) \tag{6.23}$$

with equality at x.

Since $V(\cdot)$ is concave, $-V(\cdot)$ is convex, and by a standard result in convex analysis, it possesses subgradients. Moreover, any subgradient $-p$ of $-V$ at x must satisfy

$$p \cdot (x' - x) \geq V(x') - V(x) \geq W(x') - W(x), \qquad \text{for all } x' \in \mathcal{N}(x),$$

where the first inequality uses the definition of a subgradient and the second uses the fact that $W(x') \leq V(x')$, with equality at x as established in (6.23). This implies that every subgradient p of $-V$ is also a subgradient of $-W$. Since W is differentiable at x, its subgradient p must be unique, and another standard result in convex analysis implies that any convex function with a unique subgradient at an interior point x is differentiable at x. This establishes that $-V(\cdot)$, thus $V(\cdot)$, is differentiable as desired. The expression for the gradient (6.4) then follows from the Envelope Theorem, Theorem A.32. □

6.6. Fundamentals of Stationary Dynamic Programming

In this section, I return to the fundamentals of dynamic programming and show how they can be applied in a range of problems. The main result in this section is Theorem 6.10, which shows how dynamic first-order conditions, the *Euler equations*, together with the *transversality condition* are sufficient to characterize solutions to dynamic optimization problems. This theorem is arguably more useful in practice than the main dynamic programming theorems presented above.

6.6.1. Basic Equations. Consider the functional equation corresponding to Problem A2:

$$V(x) = \max_{y \in G(x)} \{U(x, y) + \beta V(y)\}, \text{ for all } x \in X. \tag{6.24}$$

Let us assume throughout that Assumptions 6.1-6.5 hold. Then, from Theorem 6.4, the maximization problem in (6.24) is strictly concave, and from Theorem 6.6, the maximand is

also differentiable. Therefore for any interior solution $y \in \text{Int}G(x)$, the first-order conditions are necessary and sufficient for an optimum (taking $V(\cdot)$ as given). In particular, optimal solutions can be characterized by the following convenient *Euler equations*:

$$D_y U(x, y^*) + \beta DV(y^*) = 0, \tag{6.25}$$

where I use $*$'s to denote optimal values and once again D denotes gradients (recall that, in the general case, x is a vector not a real number, thus $D_x U$ is a vector of partial derivatives and I denote the vector of partial derivatives of the value function V evaluated at y by $DV(y)$). Throughout the rest of the chapter, I adopt the convention that $D_y U$ denotes the gradient vector of U with respect to its last K arguments (even when I write $D_y U(x(t), x(t+1))$), whereas $D_x U$ is the gradient with respect to the first K arguments.

The set of first-order conditions in eq. (6.25) would be sufficient to solve for the optimal policy, y^*, if we knew the form of the $V(\cdot)$ function. Since this function is determined recursively as part of the optimization problem, there is a little more work to do before we obtain the set of equations that can be solved for the optimal policy.

Fortunately, we can use the equivalent of the Envelope Theorem for dynamic programming (recall Theorem A.32 in Appendix Chapter A) and differentiate (6.24) with respect to the state vector, x, to obtain:

$$DV(x) = D_x U(x, y^*). \tag{6.26}$$

The reason why this is the equivalent of the Envelope Theorem is that the term $[D_y U(x, y^*) + \beta DV(y^*)] dy/dx$ (that is the effect of a change in y times the induced change in y in response to the change in x) is absent from the expression. This is because the term $D_y U(x, y^*) + \beta D_x V(y^*) = 0$ from (6.25).

Now using the notation $y^* = \pi(x)$ to denote the optimal policy function (which is single-valued in view of Assumption 6.3 and Corollary 6.1) and the fact that $DV(y) = D_x V(\pi(x))$, we can combine these two equations to write

$$D_y U(x, \pi(x)) + \beta D_x U(\pi(x), \pi(\pi(x))) = 0, \tag{6.27}$$

where recall that $D_x U$ represents the gradient vector of U with respect to its first K arguments, and $D_y U$ represents its gradient with respect to the second set of K arguments. Notice that (6.27) is a functional equation in the unknown function $\pi(\cdot)$ and characterizes the optimal policy function.

These equations become even simpler and more transparent in the case where both x and y are real numbers. In this case, (6.25) becomes:

$$\frac{\partial U(x, y^*)}{\partial y} + \beta V'(y^*) = 0, \tag{6.28}$$

where V' the notes the derivative of the V function with respect to its single argument.

This equation is very intuitive; it requires the sum of the marginal gain today from increasing y and the discounted marginal gain from increasing y on the value of all future returns to be equal to zero. For instance, as in Example 6.1, we can think of U as decreasing

in y and increasing in x; eq. (6.28) would then require the current cost of increasing y to be compensated by higher values tomorrow. In the context of growth, this corresponds to the current cost of reducing consumption to be compensated by higher consumption tomorrow. As with (6.25), the value of higher consumption in (6.28) is expressed in terms of the derivative of the value function, $V'(y^*)$, which is one of the unknowns. Let us now use the one-dimensional version of (6.26) to find an expression for this derivative:

$$V'(x) = \frac{\partial U(x, y^*)}{\partial x}. \tag{6.29}$$

Now in this one-dimensional case, combining (6.29) together with (6.28) yields the following very simple condition:

$$\frac{\partial U(x, \pi(x))}{\partial y} + \beta \frac{\partial U(\pi(x), \pi(\pi(x)))}{\partial x} = 0,$$

where, in line with the notation for gradients, $\partial U/\partial x$ denotes the derivative of U with respect to its first argument and $\partial U/\partial y$ with respect to the second argument.

Alternatively, explicitly including the time arguments, the Euler equation can be written as

$$\frac{\partial U(x(t), x^*(t+1))}{\partial x} + \beta \frac{\partial U(x^*(t+1), x^*(t+2))}{\partial y} = 0. \tag{6.30}$$

However, this Euler equation is not sufficient for optimality. In addition we need the transversality condition. The transversality condition is essential in infinite-dimensional problems, because it makes sure that there are no beneficial simultaneous changes in an infinite number of choice variables. In contrast, in finite-dimensional problems, there is no need for such a condition, since the first-order conditions are sufficient to rule out possible gains when we change many or all of the control variables at the same time. The role that the transversality condition plays in infinite-dimensional optimization problems will become more apparent after we see Theorem 6.10 and after the discussion in the next subsection.

In the general case, the transversality condition takes the form:

$$\lim_{t \to \infty} \beta^t D_x U(x^*(t), x^*(t+1)) \cdot x^*(t) = 0, \tag{6.31}$$

where "$\cdot$" denotes the inner product operator. In the one-dimensional case, we have the simpler transversality condition:

$$\lim_{t \to \infty} \beta^t \frac{\partial U(x^*(t), x^*(t+1))}{\partial x} \cdot x^*(t) = 0. \tag{6.32}$$

In words, this condition requires that the product of the marginal return from the state variable x times the value of this state variable does not increase asymptotically at a rate faster than or equal to $1/\beta$.

The next theorem shows that the transversality condition together with the transformed Euler equations in (6.27) are necessary and sufficient to characterize an optimal solution to Problem A1 and therefore to Problem A2.

THEOREM 6.10. ***(Euler Equations and the Transversality Condition)*** *Let $X \subset \mathbb{R}_{+}^{K}$, and suppose that Assumptions 6.1-6.5 hold. Then, a sequence $\{x^{*}(t+1)\}_{t=0}^{\infty}$ such that $x^{*}(t+1) \in IntG(x^{*}(t))$, $t=0,1,\ldots,$ is optimal for Problem A1 given $x(0)$ if and only if it satisfies (6.27) and (6.31).*

PROOF. **(Sufficiency)** Consider an arbitrary $x(0)$ and $\mathbf{x}^{*} \equiv (x(0), x^{*}(1), ...) \in \Phi(x(0))$ be a feasible (nonnegative) sequence satisfying (6.27) and (6.31). I first show that $\mathbf{x}^{*}$ yields higher value than any other $\mathbf{x} \equiv (x(0), x(1), ...) \in \Phi(x(0))$. For any $\mathbf{x} \in \Phi(x(0))$, define

$$\Delta_{\mathbf{x}} \equiv \lim_{T\to\infty} \sup \sum_{t=0}^{T} \beta^{t}[U(x^{*}(t), x^{*}(t+1)) - U(x(t), x(t+1))]$$

as the lim sup of the difference of the values of the objective function evaluated at the feasible sequences $\mathbf{x}^{*}$ and $\mathbf{x}$ as the time horizon goes to infinity. Here lim sup is used instead of lim, since there is no guarantee that for an arbitrary $\mathbf{x} \in \Phi(x(0))$, the limit exists.

From Assumptions 6.2 and 6.5, U is continuous, concave, and differentiable. Since U is concave, Theorem A.24 and the multivariate equivalent of Corollary A.4 in Appendix Chapter A imply that

$$\begin{aligned}\Delta_{\mathbf{x}} \geq & \lim_{T\to\infty} \sup \sum_{t=0}^{T} \beta^{t}[D_{x}U(x^{*}(t), x^{*}(t+1)) \cdot (x^{*}(t) - x(t)) \\ & + D_{x}U(x^{*}(t), x^{*}(t+1)) \cdot (x^{*}(t+1) - x(t+1))]\end{aligned}$$

for any $\mathbf{x} \in \Phi(x(0))$. Since $x^{*}(0) = x(0)$, $D_{x}U(x^{*}(0), x^{*}(1)) \cdot (x^{*}(0) - x(0)) = 0$. Therefore, we can rearrange terms in this expression to obtain

$$\begin{aligned}&\Delta_{\mathbf{x}} \geq \\ &\lim_{T\to\infty} \sup \sum_{t=0}^{T} \beta^{t}\left[D_{y}U(x^{*}(t), x^{*}(t+1)) + \beta D_{x}U(x^{*}(t+1), x^{*}(t+2))\right] \cdot (x^{*}(t+1) - x(t+1)) \\ &- \lim_{T\to\infty} \inf \beta^{T} D_{x}U(x^{*}(T), x^{*}(T+1)) \cdot x^{*}(T+1) \\ &+ \lim_{T\to\infty} \sup \beta^{T} D_{x}U(x^{*}(T), x^{*}(T+1)) \cdot x(T+1)).\end{aligned}$$

Since $\mathbf{x}^{*}$ satisfies (6.27), the terms in the first line are all equal to zero. Moreover, since it satisfies (6.31), the second line is also equal to zero. Finally, from Assumption 6.4, U is increasing in x, that is, $D_{x}U \geq 0$, and moreover $x \geq 0$ by hypothesis, so the last term is nonnegative, establishing that $\Delta_{\mathbf{x}} \geq 0$ for any $\mathbf{x} \in \Phi(x(0))$. Consequently, $\mathbf{x}^{*}$ yields higher value than any feasible $\mathbf{x} \in \Phi(x(0))$ and is therefore optimal.

(Necessity) Suppose that $\{x^{*}(t+1)\}_{t=0}^{\infty}$, with $x^{*}(t+1) \in \text{Int}G(x^{*}(t))$ for all t constitutes an optimal plan. This implies that $\Delta_{\mathbf{x}}$ defined in the sufficiency part is nonnegative for any $\mathbf{x} \in \Phi(x(0))$. Consider $\mathbf{x} \in \Phi(x(0))$ such that $x(t) = x^{*}(t) - \varepsilon z(t)$, where $z(t) \in \mathbb{R}^{K}$ for each t and ε is a real number. For ε sufficiently small, such an $\mathbf{x}$ in $\Phi(x(0))$ can be found by virtue of the fact that $x^{*}(t+1) \in \text{Int}G(x^{*}(t))$ for all t. Then, by Theorem A.24 in Appendix

Chapter A,

$$\begin{aligned}\Delta_{\mathbf{x}} &= \lim_{T\to\infty}\sup\sum_{t=0}^{T}\beta^t[D_xU(x^*(t),x^*(t+1))\cdot\varepsilon z(t))\\ &\quad +D_xU(x^*(t),x^*(t+1))\cdot\varepsilon z(t+1))]\\ &\quad +\lim_{T\to\infty}\sup\sum_{t=0}^{T}\beta^t o(\varepsilon),\end{aligned}$$

where the last line consists of second-order terms that satisfy $\lim_{\varepsilon\to 0} o(\varepsilon)/\varepsilon = 0$. Now if (6.27) is violated at some t', take $y(t) = 0$ for all $t \neq t'$, ε and $z(t')$ such that $D_xU(x^*(t'),x^*(t'+1))\cdot z(t') < 0$, and $\varepsilon \to 0$. This guarantees that $\Delta_{\mathbf{x}} < 0$, yielding a contradiction to the hypothesis that (6.27) is not satisfied.

Next, suppose that (6.27) is satisfied but (6.31) is violated. Choosing $x(t) = (1-\varepsilon)x^*(t)$ and repeating the same steps as in this efficiency part, we have

$$\begin{aligned}\Delta_{\mathbf{x}} = -\varepsilon\lim_{T\to\infty}\inf\beta^T D_xU(x^*(T),x^*(T+1))\cdot x^*(T+1)\\ +\lim_{T\to\infty}\sup\sum_{t=0}^{T}\beta^t o(\varepsilon),\end{aligned}$$

where the remaining terms have been canceled by using (6.27). If (6.31) is violated, the first-term can be made strictly negative (by choosing ε negative or positive). Since $\lim_{\varepsilon\to 0} o(\varepsilon)/\varepsilon = 0$, as $\varepsilon \to 0$, this implies that $\Delta_{\mathbf{x}} < 0$, contradicting the hypothesis that (6.31) does not hold. This completes the proof. □

Theorem 6.10 shows that the simple form of the transversality condition, (6.31), is both necessary and sufficient sufficient for an interior optimal plan as long as the Euler equations, (6.27), are also satisfied. The Euler equations, (6.27), are also necessary for an interior solution. Theorem 6.10 is therefore often all that we need to characterize solutions to dynamic optimization problems.

I now illustrate how the tools that have been developed so far can be used in the context of the problem of optimal growth, which will be further discussed in Section 6.8.

Example 6.4. Consider the following optimal growth, with log preferences, Cobb-Douglas technology and full depreciation of capital stock

$$\begin{aligned}&\max_{\{c(t),k(t+1)\}_{t=0}^{\infty}}\sum_{t=0}^{\infty}\beta^t\ln c(t)\\ &\text{subject to}\\ &k(t+1) = k(t)^{\alpha} - c(t)\\ &k(0) = k_0 > 0,\end{aligned}$$

where, as usual, $\beta\in(0,1)$, k denotes the capital-labor ratio (capital stock), and the resource constraint follows from the production function $K^{\alpha}L^{1-\alpha}$, written in per capita terms.

This is one of the canonical examples that admit an explicit-form characterization. To derive this characterization, let us follow Example 6.1 and set up the maximization problem in its recursive form as

$$V(x) = \max_{y \geq 0} \{\ln(x^{\alpha} - y) + \beta V(y)\},$$

with x corresponding to today's capital stock and y to tomorrow's capital stock. Our main objective is to find the policy function $y = \pi(x)$, which determines tomorrow's capital stock as a function of today's capital stock. Once this is done, we can easily determine the level of consumption as a function of today's capital stock from the resource constraint.

It can be verified that this problem satisfies Assumptions 6.1-6.5. In particular, using the same argument as in Section 6.8, x and y again can be restricted to be in a compact set. Consequently, Theorems 6.1-6.6 apply. In particular, since $V(\cdot)$ is differentiable, the Euler equation for the one-dimensional case, (6.28), implies

$$\frac{1}{x^{\alpha} - y} = \beta V'(y).$$

The Envelope condition, (6.29), gives:

$$V'(x) = \frac{\alpha x^{\alpha-1}}{x^{\alpha} - y}.$$

Thus using the notation $y = \pi(x)$ and combining these two equations, we have

$$\frac{1}{x^{\alpha} - \pi(x)} = \beta \frac{\alpha \pi(x)^{\alpha-1}}{\pi(x)^{\alpha} - \pi(\pi(x))} \text{ for all } x,$$

which is a functional equation in a single function, $\pi(x)$. There are no straightforward ways of solving functional equations, but in most cases guess-and-verify type methods are most fruitful. For example in this case, let us conjecture that

$$\pi(x) = ax^{\alpha}. \tag{6.33}$$

Substituting for this in the previous expression,

$$\begin{aligned} \frac{1}{x^{\alpha} - ax^{\alpha}} &= \beta \frac{\alpha a^{\alpha-1} x^{\alpha(\alpha-1)}}{a^{\alpha} x^{\alpha^2} - a^{1+\alpha} x^{\alpha^2}}, \\ &= \frac{\beta}{a} \frac{\alpha}{x^{\alpha} - ax^{\alpha}}, \end{aligned}$$

which implies that $a = \beta\alpha$ satisfies this equation. The policy function $\pi(x) = \beta\alpha x^{\alpha}$ also implies that the law of motion of the capital stock is

$$k(t+1) = \beta\alpha k(t)^{\alpha} \tag{6.34}$$

and the optimal consumption level is

$$c(t) = (1 - \beta a) k(t)^{\alpha}.$$

It can then be verified that the capital-labor ratio $k(t)$ converges to a steady state k^*, which is sufficient to ensure the transversality condition (6.32). Corollary 6.1 and Theorem 6.10 then imply that $\pi(x) = \beta\alpha x^{\alpha}$ must be the unique policy function for this problem. Exercise

6.7 continues with some of the details of this example, and also shows how the optimal growth equilibrium involves a sequence of capital-labor ratios converging to a unique steady state.

Finally, let us have a brief look at the intertemporal utility maximization problem of a consumer facing a certain income sequence.

Example 6.5. Consider the problem of an infinitely-lived consumer with instantaneous utility function defined over consumption $u(c)$, where $u : \mathbb{R}_+ \to \mathbb{R}$ is strictly increasing, continuously differentiable and strictly concave. The consumer discounts the future exponentially with the constant discount factor $\beta \in (0,1)$. He also faces a certain (nonnegative) labor income stream of $\{w(t)\}_{t=0}^{\infty}$ and starts life with a given amount of assets $a(0) \in \mathbb{R}$. He receives a constant net rate of interest $r > 0$ on his asset holdings (so that the gross rate of return is $1+r$). To start with, let us suppose that wages are constant, that is, $w(t) = w \in \mathbb{R}_+$.

Then, the utility maximization problem of the consumer can be written as

$$\max_{\{c(t),a(t)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c(t))$$

subject to the flow budget constraint

$$a(t+1) = (1+r)(a(t) + w - c(t)),$$

with initial assets, $a(0)$, given. This maximization problem, without further constraints, is not well defined and in fact does not capture the utility maximization problem of the consumer. In particular, Exercise 6.11 shows that this problem allows the consumer to build up debt without limit (so that $\lim_{t\to\infty} a(t) = -\infty$). This is sometimes referred to as a "Ponzi game" and involves the consumer continuously borrowing and rolling over debt. Moreover, clearly the consumer will prefer consumption paths that do this, since they enable him to increase his consumption at all dates. However, from an economic point of view this is a nonsensical solution and would not be feasible in a market economy, because the financial market (lenders) transacting with the consumer would make very large losses (at least asymptotically). This result emerges because we have not imposed the appropriate budget constraint on the consumer. As discussed in greater detail in Chapter 8, the flow budget constraint is not sufficient to capture the lifetime budget constraint of the consumer (and this is evidenced by the fact that the consumer can have consumption paths with $\lim_{t\to\infty} a(t) = -\infty$). An additional constraint needs to be imposed so that the consumer cannot run such "Ponzi games".

There are three solutions to this problem. The first, which will be discussed in greater detail in Chapter 8 is to impose a "no-Ponzi condition" that rules out such schemes. The second is to assume that the consumer cannot borrow, so $a(t) \geq 0$ for all t. While this prevents the problem, it is somewhat too "draconian". For example, we may wish to allow the consumer to borrow temporarily, as many financial markets do, but not allow him to have an asset position reaching negative infinity. The third solution, which is in many ways the simplest, is to impose the *natural debt limit*. This essentially involves computing what

the maximum debt that the consumer can repay and then impose that his debt level should never fall below this. It is clear that, since consumption is nonnegative, the consumer can never repay more than his total wage income. Given a constant wage rate w and a constant interest rate r, the net present discounted value of the consumer's wage income is w/r. This is a finite number, since $w \in \mathbb{R}_+$ and $r > 0$. Therefore, if his asset holdings, $a(t)$, ever fall below $-w/r$, the consumer will never be able to repay. It is therefore natural to expect the financial markets, even if they have the power to confiscate all of the income of the consumer, never to lend to such a level that $a(t) < -w/r$ for any t. In this light, the most economically natural additional constraint is to impose the natural debt limit and require that

$$a(t) \geq \underline{a} \equiv -\frac{w}{r} \text{ for all } t.$$

Throughout, the book I will often impose this type of natural debt limit on households, consumers and individuals.

The next challenges that even with the natural debt limit, the assets of the consumer, $a(t)$, which correspond to the state variable of this maximization problem, do not necessarily belong to a compact set. Consequently, the theorems developed so far cannot be directly applied to this problem. One way of proceeding is to strengthen these theorems, so that they cover situations in which the feasible set (X in terms of the notation above) is potentially unbounded. While such a strengthening is possible, it requires additional arguments. An alternative approach is to make use of the economic structure of the model. In particular, the following method works in general (but not always, see Exercise 6.12). Let us choose some $\bar{a}$ and limit $a(t)$ to lie in the interval $(\underline{a}, \bar{a})$. Solve the problem and then verify that $a(t)$ is in the interior of this set. In this example, again using the fact that the net present discounted value of the future labor income of the consumer is w/r, an obvious choice is $\bar{a} \equiv a(0) + w/r < \infty$ (since $a(0) \in \mathbb{R}$ and $w/r < \infty$). Exercise 6.12 shows the conditions under which $a(t)$ always lies in the interval $[\underline{a}, \bar{a}]$. This strategy of finding an upper bound for the state variable and thus ensuring that it will lie in a compact set is often used in applications.

Finally, before characterizing the solution to this problem, note that the budget constraint could have been written alternatively as $a(t+1) = (1+r)a(t) + w - c(t)$. The difference between these two alternative budget constraints involes the timing of interest payments. The first one presumes that the consumer starts the period with assets $a(t)$, then receives his labor income, w at time t, and then consumes $c(t)$. Whatever is left is saved for the next date and earns the gross interest rate $(1+r)$. In this formulation, $a(t)$ refers to asset holdings at the beginning of time t. The alternative formulation instead interprets $a(t)$ as asset holdings at the end of time t. The choice between these two formulations has no bearing on the results.

With these assumptions, the recursive formulation of the consumer's maximization problem can be written in a straightforward manner. The state variable is $a(t)$, and consumption,

$c(t)$, can be expressed as

$$c(t) = a(t) + w - (1+r)^{-1} a(t+1).$$

With standard arguments and denoting the current-value of the state variable by a and its future value by a', the recursive form of this dynamic optimization problem can be written as

$$V(a) = \max_{a' \in [\underline{a}, \bar{a}]} \left\{ u\left(a + w - (1+r)^{-1} a'\right) + \beta V\left(a'\right) \right\}.$$

Clearly $u(\cdot)$ is strictly increasing in a, continuously differentiable in a and a' and is strictly concave in a. Moreover, since $u(\cdot)$ is continuously differentiable in $a \in (\underline{a}, \bar{a})$ and the individual's wealth is finite, $V(a(0))$ is also finite. Thus all of the results from our analysis above, in particular Theorems 6.1-6.6, apply and imply that $V(a)$ is differentiable and a continuous solution $a' = \pi(a)$ exists. Moreover, we can use the Euler equation (6.25), or its more specific form (6.28) for one-dimensional problems to characterize the optimal consumption plan. In particular,

$$\begin{aligned} u'\left(a + w - (1+r)^{-1} a'\right) &= \\ u'(c) &= \beta(1+r) V'\left(a'\right). \end{aligned} \tag{6.35}$$

This important equation is often referred to as the "consumption Euler" equation. It states that the marginal utility of current consumption must be equal to the marginal increase in the continuation value multiplied by the product of the discount factor, β, and the gross rate of return to savings, $(1+r)$. It captures the essential economic intuition of dynamic programming approach, which reduces the complex infinite-dimensional optimization problem to one of comparing today to "tomorrow". As usual, the only difficulty here is that tomorrow itself will involve a complicated maximization problem and hence tomorrow's value function and its derivative are endogenous. But here the Envelope condition, (6.29), again comes to our rescue and yields

$$V'\left(a'\right) = u'\left(c'\right),$$

where c' refers to next period's consumption. Using this relationship, the consumption Euler equation becomes

$$u'(c) = \beta(1+r) u'\left(c'\right). \tag{6.36}$$

This form of the consumption Euler equation is more familiar and requires the marginal utility of consumption today to be equal to the marginal utility of consumption tomorrow multiplied by the product of the discount factor and the gross rate of return. Since we have assumed that β and $(1+r)$ are constant, the relationship between today's and tomorrow's consumption never changes. In particular, since $u(\cdot)$ is assumed to be continuously differentiable and strictly concave, $u'(\cdot)$ always exists and is strictly decreasing. Therefore, the intertemporal consumption maximization problem implies the following simple rule:

$$\begin{array}{ll} \text{if } r = \beta^{-1} - 1 & c = c' \text{ and consumption is constant over time} \\ \text{if } r > \beta^{-1} - 1 & c < c' \text{ and consumption increases over time} \\ \text{if } r < \beta^{-1} - 1 & c > c' \text{ and consumption decreases over time.} \end{array} \tag{6.37}$$

The remarkable feature is that these statements have been made without any reference to the initial level of asset holdings $a(0)$ and the wage rate w. It turns out that these only determine the initial level of consumption. The "slope" of the optimal consumption path is independent of the wealth of the individual. Exercise 6.13 asks you to determine the level of initial consumption using the transversality condition and the intertemporal budget constraint, while Exercise 6.12 asks you to verify that whenever $r \leq \beta - 1$, $a(t) \in (\underline{a}, \bar{a})$ for all t (so that the artificial bounds on asset holdings that I imposed have no bearing on the results).

The example here is somewhat restrictive because wages are assumed to be constant over time. What happens if instead there is a time-varying sequence of wages $\{w(t)\}_{t=0}^{\infty}$? What happens if there is a time-varying sequence of interest rates $\{r(t)\}_{t=0}^{\infty}$? Unfortunately, with a time-varying sequence of wages or interest rates, the problem becomes nonstationary and the theorems developed for Problems A1 and A2 above no longer apply in this case. However, many equilibrium problems will have this feature, because individuals will be facing time-varying market prices. This motivates my discussion of nonstationary problems in the next section.

6.6.2. Dynamic Programming Versus the Sequence Problem. Before turning to nonstationary problems, let us compare the dynamic programming formulation to the sequence problem, and motivate the transversality condition using the sequence problem. Let us also suppose that x is one dimensional and that there is a finite horizon T. Then, the problem becomes

$$\max_{\{x(t+1)\}_{t=0}^{T}} \sum_{t=0}^{T} \beta^t U(x(t), x(t+1))$$

subject to $x(t+1) \geq 0$ with $x(0)$ as given. Moreover, let $U(x(T), x(T+1))$ be the last period's utility, with $x(T+1)$ as the state variable left after the last period (this utility could be thought of as the "salvage value" for example).

In this case, we have a finite-dimensional optimization problem and we can simply look at first-order conditions. Moreover, let us again assume that the optimal solution lies in the interior of the constraint set, that is, $x^*(t) > 0$, so that first-order conditions do not need to be expressed as complementary-slackness type conditions. In particular, in this case they take the following simple form, equivalent to the Euler equations above,

$$\text{for any } 0 \leq t \leq T-1, \ \frac{\partial U(x^*(t), x^*(t+1))}{\partial y} + \beta \frac{\partial U(x^*(t+1), x^*(t+2))}{\partial x} = 0,$$

which are identical to the Euler equations for the infinite-horizon case (recall that $\partial U/\partial x$ denotes the derivative of U with respect to its first argument and $\partial U/\partial y$ with respect to the second argument). In addition, for $x(T+1)$, the following boundary condition is also necessary

$$x^*(T+1) \geq 0, \text{ and } \beta^T \frac{\partial U(x^*(T), x^*(T+1))}{\partial y} x^*(T+1) = 0. \tag{6.38}$$

Intuitively, this boundary condition requires that $x^*(T+1)$ should be positive only if an interior value of it maximizes the salvage value at the end. To provide more intuition for this expression, let us return to the formulation of the optimal growth problem in Example 6.1.

EXAMPLE 6.6. Recall that in terms of the optimal growth problem,

$$U(x(t), x(t+1)) = u(f(x(t)) + (1-\delta)x(t) - x(t+1)),$$

with $x(t) = k(t)$ and $x(t+1) = k(t+1)$. Suppose that, in contrast to the infinite-horizon models, the world comes to an end at date T. Then, at the last date T,

$$\frac{\partial U(x^*(T), x^*(T+1))}{\partial y} = -u'(c^*(T+1)) < 0.$$

From (6.38) and the fact that U is increasing in its first argument (Assumption 6.4), an optimal path must have $k^*(T+1) = x^*(T+1) = 0$. Intuitively, there should be no capital left at the end of the world. If any resources were left after the end of the world, utility could be improved by consuming them either at the last date or at some earlier date.

Now, the transversality condition can be derived heuristically as an extension of condition (6.38) to the case where $T = \infty$. Taking this limit,

$$\lim_{T\to\infty} \beta^T \frac{\partial U(x^*(T), x^*(T+1))}{\partial y} x^*(T+1) = 0.$$

Moreover, as $T \to \infty$, the Euler equation implies

$$\frac{\partial U(x^*(T), x^*(T+1))}{\partial y} + \beta \frac{\partial U(x^*(T+1), x^*(T+2))}{\partial x} = 0.$$

Substituting this relationship into the previous equation,

$$-\lim_{T\to\infty} \beta^{T+1} \frac{\partial U(x^*(T+1), x^*(T+2))}{\partial x} x^*(T+1) = 0.$$

Canceling the negative sign, and changing the timing for convenience:

$$\lim_{T\to\infty} \beta^T \frac{\partial U(x^*(T), x^*(T+1))}{\partial x} x^*(T) = 0,$$

which is exactly the transversality condition in (6.32). This derivation also highlights that the transversality condition can equivalently be written as

$$\lim_{T\to\infty} \beta^T \frac{\partial U(x^*(T), x^*(T+1))}{\partial y} x^*(T+1) = 0.$$

This emphasizes that there is no unique representation of the transversality condition, but various different equivalent forms. They all correspond to a "boundary condition at infinity" ruling out variations that change an infinite number of control variables at the same time.

6.7. Nonstationary Infinite-Horizon Optimization

6.7.1. Main Results. Let us now return to Problem A0. The possible nonstationarity in this problem makes it considerably more difficult than Problems A1 and A2. Nevertheless, many important economic problems, for example, utility maximization by a household in a dynamic competitive equilibrium corresponds to such a nonstationary problem. One way of making progress is to introduce additional assumptions on U and G so as to obtain the equivalents of Theorems 6.1-6.6 (see, for example, Exercise 6.14). Nevertheless, existence of solutions and the necessity and sufficiency of Euler equations and the transversality conditions (as in Theorem 6.10) can be established under fairly general conditions. This is the approach I use in this section. In particular, let us again define the set of feasible sequences or *plans* starting with an initial value $x(t)$ at time t as:

$$\Phi(t, x(t)) = \{\{x(s)\}_{s=t}^{\infty} : x(s+1) \in G(t, x(s)), \text{ for } s = t, t+1, ...\}.$$

Also denote an element of this set by $\mathbf{x}[t] = (x(t), x(t+1), ...) \in \Phi(t, x(t))$. The key assumptions will be as follows (and they are numbered with reference to Assumptions 6.1-6.5 above):

ASSUMPTION 6.1N. $G(t,x)$ is nonempty for all $x \in X$ and $t \in \mathbb{Z}_+$ and $U(t,x,y)$ is uniformly bounded (from above), that is, there exists $M < \infty$ such that $U(t,x,y) \leq M$ for all $t \in \mathbb{Z}_+$, $x \in X$ and $y \in G(t,x)$.

This assumption is again stronger than necessary to establish the results that will follow, but it simplifies the exposition.

ASSUMPTION 6.2N. X is a compact subset of $\mathbb{R}^K$, G is nonempty, compact-valued and continuous. Moreover, $U : \mathbf{X}_G \to \mathbb{R}$ is continuous in x and y, where $\mathbf{X}_G = \{(t,x,y) \in X \times X : y \in G(t,x)\}$.

ASSUMPTION 6.3N. U is concave; for any $\alpha \in (0,1)$ and any (t,x,y), $(t,x',y') \in \mathbf{X}_G$, we have

$$U\left(t, \alpha x + (1-\alpha)x', \alpha y + (1-\alpha)y'\right) \geq \alpha U(t,x,y) + (1-\alpha)U(t,x',y').$$

Moreover, if $x \neq x'$,

$$U\left(t, \alpha x + (1-\alpha)x', \alpha y + (1-\alpha)y'\right) > \alpha U(t,x,y) + (1-\alpha)U(t,x',y').$$

In addition, G is convex; for any $\alpha \in [0,1]$, and $x, x', y, y' \in X$ such that $y \in G(t,x)$ and $y' \in G(t,x')$, we have

$$\alpha y + (1-\alpha)y' \in G\left(t, \alpha x + (1-\alpha)x'\right).$$

ASSUMPTION 6.4N. For each $t \in \mathbb{Z}_+$ and $y \in X$, $U(t,x,y)$ is strictly increasing in each component of x, and G is monotone in x in the sense that $x \leq x'$ implies $G(t,x) \subset G(t,x')$ for any $t \in \mathbb{Z}_+$.

ASSUMPTION 6.5N. U is continuously differentiable in x and y on $\text{Int}\mathbf{X}_G$ (where $\text{Int}\mathbf{X}_G$ refers to the interior of the set $\mathbf{X}_G$ with respect to x and y).

The two key results are the following. The proofs of both theorems are given in the next subsection.

THEOREM 6.11. ***(Existence of Solutions)*** *Suppose that Assumptions 6.1N and 6.2N hold. Then, there exists a unique function* $V^* : \mathbb{Z}_+ \times X \to \mathbb{R}$ *that is a solution to Problem A0.* V^* *is continuous in* x *and bounded. Moreover, an optimal plan* $\mathbf{x}^*[0] \in \Phi(0, x(0))$ *exists for any* $x(0) \in X$.

THEOREM 6.12. ***(Euler Equations and the Transversality Condition)*** *Let* $X \subset \mathbb{R}^K_+$, *and suppose that Assumptions 6.1N-6.5N hold. Then, a sequence* $\{x^*(t+1)\}_{t=0}^{\infty}$, *with* $x^*(t+1) \in IntG(t, x^*(t))$, $t = 0, 1, \ldots$, *is optimal for Problem A0 given* $x(0)$ *if and only if it satisfies the Euler equation*

$$D_y U(t, x^*(t), x^*(t+1)) + \beta D_x U(t+1, x^*(t+1), x^*(t+2)) = 0, \tag{6.39}$$

and the transversality condition

$$\lim_{t\to\infty} \beta^t D_{x(t)} U(t, x^*(t), x^*(t+1)) \cdot x^*(t) = 0. \tag{6.40}$$

In equation (6.39), as in stationary problems, $D_y U$ denotes the vector of partial derivatives of U with respect to the control vector (its last K arguments), whereas $D_x U$ denotes the vector of partial derivatives with respect to the state vector. In (6.40), "$\cdot$" again denotes the inner product of the two vectors.

These two theorems provide us with the necessary tools to tackle nonstationary discrete-time infinite-horizon optimization problems. In particular, Theorem 6.11 ensures that the solution exists and Theorem 6.12 shows that, as in the stationary case, we can simply use the Euler equations and the transversality condition to characterize this solution (as long as it is interior).

6.7.2. Proofs of Theorems 6.11 and 6.12*. .

PROOF OF THEOREM 6.11. Since U is uniformly bounded (Assumption 6.1N) and continuous (Assumption 6.2N), Theorem A.12 in Appendix Chapter A implies that the objective function in Problem A0 is continuous in $\mathbf{x}[0]$ in the product topology. Moreover, the constraint set $\Phi(0, x(0))$ is a closed subset of X^∞ (infinite product of X). Since X is compact (Assumption 6.2N), Tychonoff's Theorem, Theorem A.13, implies that X^∞ is compact in the product topology. A closed subset of a compact set is compact (Fact A.2 in Appendix Chapter A), which implies that $\Phi(0, x(0))$ is compact. From Weierstrass' Theorem (Theorem A.9) applied to Problem A0, there exists $\mathbf{x}^*[0] \in \Phi(0, x(0))$ attaining $V^*(0, x(0))$. Moreover, the constraint set is a continuous correspondence (again in the product topology), so Berge's Maximum Theorem, Theorem A.16, implies that $V^*(0, x(0))$ is continuous in $x(0)$. Since $x(0) \in X$ and X is compact, this implies that $V^*(0, x(0))$ is bounded (Corollary A.1 in Appendix Chapter A). □

PROOF OF THEOREM 6.12. The proof is very similar to that of Theorem 6.10.

(Sufficiency) Consider an arbitrary $x(0)$ and $\mathbf{x}^*[0] \equiv (x(0), x^*(1), ...) \in \Phi(0, x(0))$ be a feasible (nonnegative) sequence satisfying (6.39) and (6.40). For any $\mathbf{x}[0] \equiv (x(0), x(1), ...) \in \Phi(0, x(0))$, define

$$\Delta_{\mathbf{x}} \equiv \lim_{T\to\infty} \sup \sum_{t=0}^{T} \beta^t [U(t, x^*(t), x^*(t+1)) - U(t, x(t), x(t+1))]$$

as the difference of the objective function between the feasible sequences $\mathbf{x}^*$ and $\mathbf{x}$.

From Assumptions 6.2N and 6.5N, U is continuous, concave and differentiable. By definition of a concave function,

$$\begin{aligned}\Delta_{\mathbf{x}} \geq & \lim_{T\to\infty} \sup \sum_{t=0}^{T} \beta^t [D_x U(t, x^*(t), x^*(t+1)) \cdot (x^*(t) - x(t)) \\ & + D_y U(t, x^*(t), x^*(t+1)) \cdot (x^*(t+1) - x(t+1))]\end{aligned}$$

for any $\mathbf{x}[0] \in \Phi(0, x(0))$. Using the fact that $x^*(0) = x(0)$ and rearranging terms,

$$\begin{aligned}&\Delta_{\mathbf{x}} \geq \\ &\lim_{T\to\infty} \sup \sum_{t=0}^{T} \beta^t [D_y U(t, x^*(t), x^*(t+1)) + \beta D_x U(t+1, x^*(t+1), x^*(t+2))] \cdot (x^*(t+1) - x(t+1)) \\ &- \lim_{T\to\infty} \inf \beta^T D_x U(T, x^*(T), x^*(T+1)) \cdot x^*(T+1) \\ &+ \lim_{T\to\infty} \sup \beta^T D_x U(T, x^*(T), x^*(T+1)) \cdot x(T+1)).\end{aligned}$$

Since $\mathbf{x}^*[0]$ satisfies (6.39), the terms in first line are all equal to zero. Moreover, since it satisfies (6.40), the second line is also equal to zero. Finally, from Assumption 6.4N, U is increasing in x, that is, $D_x U \geq 0$, and moreover $x \geq 0$ by hypothesis, so the last term is nonnegative, establishing that $\Delta_{\mathbf{x}} \geq 0$ for any $\mathbf{x}[\mathbf{0}] \in \Phi(0, x(0))$. Consequently, $\mathbf{x}^*[0]$ yields higher value than any feasible $\mathbf{x}[0] \in \Phi(0, x(0))$ and is therefore optimal.

(Necessity) Suppose that $\{x^*(t+1)\}_{t=0}^{\infty}$, with $x^*(t+1) \in \text{Int}G(t, x^*(t))$ for all t constitutes an optimal plan. This implies that $\Delta_{\mathbf{x}}$ defined in the sufficiency part is nonnegative for any $\mathbf{x}[0] \in \Phi(0, x(0))$. Consider $\mathbf{x}[0] \in \Phi(x(0))$ such that $x(t) = x^*(t) - \varepsilon z(t)$, where $z(t) \in \mathbb{R}^K$ and ε is a real number. For ε sufficiently small, such an $\mathbf{x}[0]$ in $\Phi(0, x(0))$ can be found by virtue of the fact that $x^*(t+1) \in \text{Int}G(t, x^*(t))$ for all t. Then, by Theorem A.24 in Appendix Chapter A,

$$\begin{aligned}\Delta_{\mathbf{x}} = & \lim_{T\to\infty} \sup \sum_{t=0}^{T} \beta^t [D_x U(t, x^*(t), x^*(t+1)) \cdot \varepsilon z(t)) \\ & + D_x U(t, x^*(t), x^*(t+1)) \cdot \varepsilon z(t+1))] \\ & + \lim_{T\to\infty} \sup \sum_{t=0}^{T} \beta^t o(\varepsilon),\end{aligned}$$

where the last line consists of all the second-order terms, which satisfy $\lim_{\varepsilon\to 0} o(\varepsilon)/\varepsilon = 0$. Now if (6.39) is violated at some t', take $z(t) = 0$ for all $t \neq t'$, ε and $z(t')$ such that

$D_x U(x^*(t'), x^*(t'+1)) \cdot z(t') < 0$, and $\varepsilon \to 0$. This guarantees that $\Delta_{\mathbf{x}} < 0$, yielding a contradiction to the hypothesis that (6.39) is not satisfied.

Next, suppose that (6.39) is satisfied but (6.40) is violated. Choosing $x(t) = (1-\varepsilon) x^*(t)$ and repeating the same steps as in this efficiency part, we have

$$\Delta_{\mathbf{x}} = -\varepsilon \lim_{T\to\infty} \inf \beta^T D_x U(T, x^*(T), x^*(T+1)) \cdot x^*(T+1) + \lim_{T\to\infty} \sup \sum_{t=0}^{T} \beta^t o(\varepsilon),$$

where the remaining terms have been canceled by using (6.39). If (6.40) is violated, the first-term can be made strictly negative (by choosing ε negative or positive). Since $\lim_{\varepsilon\to 0} o(\varepsilon)/\varepsilon = 0$, as $\varepsilon \to 0$, this implies that $\Delta_{\mathbf{x}} < 0$, contradicting the hypothesis that (6.40) does not hold and completing the proof. □

6.7.3. Application. As an application of nonstationary infinite-horizon optimization, let us consider the consumption problem in a market environment in which wages and interest rates are potentially changing over time.

EXAMPLE 6.5 (CONTINUED). The consumer is again infinitely lived, with instantaneous utility function defined over consumption $u(c)$, where $u : \mathbb{R}_+ \to \mathbb{R}$ is strictly increasing, continuously differentiable and strictly concave. The discount factor is $\beta \in (0,1)$. Now both the sequence of labor income, $\{w(t)\}_{t=0}^{\infty}$, and interest rates, $\{r(t)\}_{t=0}^{\infty}$, are potentially time varying, though there is still no uncertainty. The consumer starts life with a given amount of assets $a(0)$. His utility maximization problem can then be written as

$$\max_{\{c(t),a(t)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c(t))$$

subject to the flow budget constraint

$$a(t+1) = (1+r(t))(a(t) + w(t) - c(t)),$$

with $a(0)$ given. As in the case with constant wage and interest rates, we need to introduce an additional constraint to complement the flow budget constraint. Let us again impose the natural debt limit. Exercise 6.13 shows that the natural debt limit at time t now takes the form

$$a(t) \geq -\sum_{s=0}^{\infty} \left\{ \left(\prod_{j=0}^{s-1} \frac{1}{1+r(t+j)} \right) w(t+s) \right\}, \tag{6.41}$$

(with the convention that $\prod_{j=0}^{-1} b(j) = 1$ for any $b(j)$). Moreover, let us assume that the right-hand side of (6.41) is bounded above by some $\bar{W} < \infty$ for all t. In this light, a natural hypothesis is that $a(t)$ should always lie in a set of the form $\left[-\bar{W}, a(0) + \bar{W}\right]$ (see Exercise

6.13). Repeating the same arguments as in Example 6.5, the consumer maximization problem can be written as

$$\max_{\{a(t+1)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u\left(a(t) + w(t)\left(1 + r(t+1)\right)^{-1} a(t+1)\right)$$

for given $a(0)$. The Euler equation in this case is

$$u'(c(t)) = \beta(1 + r(t))\, u'(c(t+1)), \tag{6.42}$$

for all t. Thus, instead of (6.37), we have the following more general consumption rule:

if $r(t) = \beta^{-1} - 1$	$c(t) = c(t+1)$ and consumption is constant between t and $t+1$
if $r(t) > \beta^{-1} - 1$	$c(t) = c(t+1)$ and consumption increases between t and $t+1$
if $r(t) < \beta^{-1} - 1$	$c(t) = c(t+1)$ and consumption decreases between t and $t+1$.

In some ways, this result is even more remarkable than the one in (6.37), since the "slope" of the optimal consumption path is not only independent of $a(0)$, but it is also independent of current income and in fact of the entire sequence of income levels $\{w(t)\}_{t=0}^{\infty}$.

6.8. Optimal Growth in Discrete Time

Let us now return to the optimal growth problem, introduced in Section 5.9, and use the main results from stationary dynamic programming to obtain a characterization of the optimal growth path of a neoclassical economy. Example 6.4 already showed how this can be done in the special case with logarithmic utility, Cobb-Douglas production function and full depreciation. In this section, the results are more general and are stated for the canonical optimal growth model introduced in Chapter 5.

Recall the optimal growth problem for a one-sector economy admitting a representative household with instantaneous utility function u and discount factor $\beta \in (0,1)$. This can be written as

$$\max_{\{c(t),k(t)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c(t)) \tag{6.43}$$

subject to

$$k(t+1) = f(k(t)) + (1-\delta)k(t) - c(t) \text{ and } k(t) \geq 0, \tag{6.44}$$

with the initial capital stock given by $k(0) > 0$.

The standard assumptions on the production function, Assumptions 1 and 2, are still in effect. In addition, let us impose:

Assumption 3′. $u : [\varepsilon, \infty) \to \mathbb{R}$ is continuously differentiable and strictly concave for some $\varepsilon > 0$ sufficiently small.

This is considerably stronger than necessary. In fact, concavity or even continuity is enough for most of the results. But this assumption helps us avoid inessential technical details. This assumption is referred to as Assumption 3′ to distinguish it from the very closely related Assumption 3 that will be introduced and used in Chapter 8 and thereafter.

The first step is to write the optimal growth problem as a (stationary) dynamic programming problem. This can be done along the lines of the above formulations. In particular, let the choice variable be next date's capital stock, denoted by s. Then, the resource constraint (6.44) implies that current consumption is given by $c = f(k) + (1-\delta)k - s$, and thus the optimal growth problem can be written in the following recursive form:

$$V(k) = \max_{s \in G(k)} \{u(f(k) + (1-\delta)k - s) + \beta V(s)\} \tag{6.45}$$

where $G(k)$ is the constraint correspondence, given by the interval $[0, f(k) + (1-\delta)k]$, which imposes that consumption cannot be negative and that the capital stock cannot be negative.

It can be verified that under Assumptions 1, 2 and 3′, the optimal growth problem satisfies Assumptions 6.1-6.5 of the dynamic programming problems. The only non-obvious feature is that the level of consumption and capital stock belong to a compact set. To verify that this is the case, note that the economy can never settle into a level of capital-labor ratio greater than $\bar{k}$, defined by

$$\delta \bar{k} = f(\bar{k}),$$

since this is the capital-labor ratio that would sustain itself when consumption is set equal to 0. If the economy starts with $k(0) < \bar{k}$, it can never exceed $\bar{k}$. If it starts with $k(0) > \bar{k}$, it can never exceed $k(0)$. Therefore, without loss of any generality, we can restrict consumption and capital stock to lie in the compact set $\left[0, \vec{k}\right]$, where

$$\vec{k} \equiv f\left(\max\left\{k(0), \bar{k}\right\}\right) + (1-\delta)\max\left\{k(0), \bar{k}\right\}.$$

Consequently, Theorems 6.1-6.6 can be directly applied to this problem.

PROPOSITION 6.1. *Given Assumptions 1, 2 and 3′, the optimal growth model as specified in (6.43) and (6.44) has a solution characterized by the unique value function $V(k)$ and consumption function $c(k)$. The capital stock of the next period is given by $s(k) = f(k) + (1-\delta)k - c(k)$. Moreover, $V(k)$ is strictly increasing and strictly concave and $s(k)$ is nondecreasing in k.*

PROOF. That the value function (6.45) is a solution to (6.43) and (6.44) follows from Theorems 6.1 and 6.2. That $V(k)$ exists follows from Theorem 6.3, and the fact that it is increasing and strictly concave, with the policy correspondence being a policy function, follows from Theorem 6.4 and Corollary 6.1.

Thus we only have to show that $s(k)$ is nondecreasing. This can be proved by contradiction using a "revealed preference" type argument. Suppose, to arrive at a contradiction, that $s(k)$ is decreasing, i.e., there exists k and $k' > k$ such that $s(k) > s(k')$. Since $k' > k$, $s(k)$ is feasible when the capital stock is k'. Moreover, since, by hypothesis, $s(k) > s(k')$, $s(k')$ is feasible at capital stock k.

By optimality and feasibility,

$$\begin{aligned} V(k) &= u(f(k)+(1-\delta)k-s(k))+\beta V(s(k)) \\ &\geq u\left(f(k)+(1-\delta)k-s\left(k'\right)\right)+\beta V\left(s\left(k'\right)\right) \\ V\left(k'\right) &= u\left(f\left(k'\right)+(1-\delta)k'-s\left(k'\right)\right)+\beta V\left(s\left(k'\right)\right) \\ &\geq u\left(f\left(k'\right)+(1-\delta)k'-s(k)\right)+\beta V(s(k)). \end{aligned}$$

Combining and rearranging these,

$$\begin{aligned} u(f(k)+(1-\delta)k-s(k))-u\left(f(k)+(1-\delta)k-s\left(k'\right)\right) &\geq \beta\left[V\left(s\left(k'\right)\right)-V(s(k))\right] \\ &\geq u\left(f\left(k'\right)+(1-\delta)k'-s(k)\right) \\ &\quad -u\left(f\left(k'\right)+(1-\delta)k'-s\left(k'\right)\right). \end{aligned}$$

Denoting $z \equiv f(k)+(1-\delta)k$ and $x \equiv s(k)$ and similarly for z' and x',

$$u\left(z-x'\right)-u(z-x) \leq u\left(z'-x'\right)-u\left(z'-x\right). \tag{6.46}$$

But clearly,

$$\left(z-x'\right)-(z-x)=\left(z'-x'\right)-\left(z'-x\right),$$

which combined with the fact that $z' > z$ (since $k' > k$) and $x > x'$ by hypothesis, and that u is strictly concave and increasing implies

$$u\left(z-x'\right)-u(z-x) > u\left(z'-x'\right)-u\left(z'-x\right),$$

contradicting (6.46). This establishes that $s(k)$ must be nondecreasing everywhere. □

In addition, Assumption 2 implies that savings and consumption levels have to be interior, thus Theorem 6.6 applies and immediately establishes:

Proposition 6.2. *Given Assumptions 1, 2 and 3′, the value function $V(k)$ defined above is differentiable.*

Let us next exploit Theorem 6.10, and look at the Euler equation and the transversality condition of this dynamic optimization problem. The Euler equation from (6.45) takes the simple form:

$$u'(c)=\beta V'(s)$$

where s denotes the next date's capital stock. Applying the Envelope condition,

$$V'(k)=\left[f'(k)+(1-\delta)\right]u'(c).$$

Consequently, we obtain the familiar condition

$$u'(c(t))=\beta\left[f'(k(t+1))+(1-\delta)\right]u'(c(t+1)). \tag{6.47}$$

The transversality condition also follows straightforwardly from the more general transversality condition, (6.32). In particular, since the instantaneous payoff function, U, now corresponds to $u(f(k)+(1-\delta)k-s)$ as a function of the current state variable k and tomorrow's state variable s, the transversality condition takes the form

$$\lim_{t\to\infty}\left[\beta^t\left(f'(k(t))+(1-\delta)\right)u'(c(t))\,k(t)\right]=0. \tag{6.48}$$

As before, a *steady state* (of the optimal growth problem) is as an allocation in which the capital-labor ratio and consumption do not depend on time, so again denoting this allocation by *, the steady state capital-labor ratio must satisfy

$$\beta \left[f'(k^*) + (1-\delta) \right] = 1. \tag{6.49}$$

This is a remarkable result; it shows that the steady-state capital-labor ratio does not depend on household preferences except via the discount factor. In particular, technology, the depreciation rate and the discount factor fully characterize the steady-state capital-labor ratio. An analog of this result applies in the continuous-time neoclassical model as well.

In addition, since $f(\cdot)$ is strictly concave, k^* is uniquely defined. Finally, since $c(t) = c^*$ and $k(t) = k^*$ in steady state, $\beta [f'(k^*) + (1-\delta)] = 1$, and $\beta < 1$, the transversality condition (6.48) is automatically satisfied. This analysis leads to the following important proposition.

PROPOSITION 6.3. *In the neoclassical optimal growth model specified in (6.43) and (6.44) with Assumptions 1, 2 and 3′, there exists a unique steady-state capital-labor ratio k^* given by (6.49), and starting from any initial $k(0) > 0$, the economy monotonically converges to this unique steady state, that is, if $k(0) < k^*$, then $k(t) \uparrow k^*$ and if $k(0) > k^*$, then $k(t) \downarrow k^*$.*

PROOF. Uniqueness and existence were established above. To establish monotone convergence, we start with arbitrary initial capital stock $k(0)$ and observe that $k(t+1) = s(k(t))$ for all $t \geq 0$, where $s(\cdot)$ was defined and shown to be nondecreasing in Proposition 6.1. It must be the case that either $k(1) = s(k(0)) \geq k(0)$ or $k(1) = s(k(0)) < k(0)$.

Consider the first case. Since $s(\cdot)$ is nondecreasing and $k(2) = s(k(1))$, we must have $k(2) \geq k(1)$. By induction, $k(t) = s(k(t-1)) \geq k(t-1) = s(k(t-2))$. Moreover, by definition $k(t) \in \left[0, \vec{k}\right]$. Therefore, in this case $\{k(t)\}_{t=0}^{\infty}$ is a nondecreasing sequence in a compact set starting with $k(0) > 0$, thus it necessarily converges to some limit $k(\infty) > 0$, which by definition satisfies $k(\infty) = s(k(\infty))$. Since k^* is the unique steady state (corresponding to positive capital-labor ratio), this implies that $k(\infty) = k^*$, and thus $k(t) \to k^*$. Moreover, since $\{k(t)\}_{t=0}^{\infty}$ is nondecreasing, it must be the case that $k(t) \uparrow k^*$, and thus this corresponds to the case where $k(0) \leq k^*$.

Next consider the case in which $k(1) = s(k(0)) < k(0)$. The same argument as above applied in reverse now establishes that $\{k(t)\}_{t=0}^{\infty}$ is a nonincreasing sequence in the compact set $\left[0, \vec{k}\right]$, thus it converges to a uniquely limit point $k(\infty)$. In this case, there are two candidate values for $k(\infty)$, $k(\infty) = 0$ or $k(\infty) = k^*$. The former is not possible, since, as Exercise 6.19 shows, Assumption 2 implies that $s(\varepsilon) > \varepsilon$ for ε sufficiently small. Thus $k(\infty) = k^*$. Since $\{k(t)\}_{t=0}^{\infty}$ is nonincreasing, in this case we must have $k(0) > k^*$ and thus $\{k(t)\}_{t=0}^{\infty} \downarrow k^*$, completing the proof. □

Consequently, in the optimal growth model there exists a unique steady state and the economy monotonically converges to the unique steady state, for example by accumulating more and more capital (if it starts with a too low capital-labor ratio).

In addition, consumption also monotonically increases (or decreases) along the path of adjustments to the unique-steady state:

PROPOSITION 6.4. *$c(k)$ defined in Proposition 6.1 is nondecreasing. Moreover, if $k_0 < k^*$, then the equilibrium consumption sequence $c(t) \uparrow c^*$ and if $k_0 > k^*$, then $c(t) \downarrow c^*$, where c^* is given by*

$$c^* = f(k^*) - \delta k^*.$$

PROOF. See Exercise 6.17. □

This discussion illustrates that the optimal growth model is very tractable and shares many commonalities with the Solow growth model, for example, a unique steady state and global monotonic convergence. There is no immediate counterpart of a saving rate, since the amount of savings depends on the utility function and changes over time, though the discount factor is closely related to the saving rate.

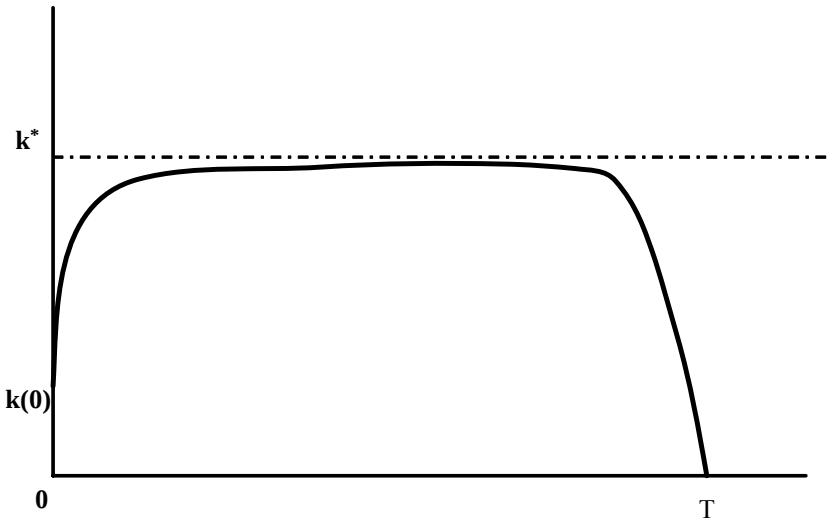

FIGURE 6.1. Turnpike dynamics in a finite-horizon (T-periods) neoclassical growth model starting with initial capital-labor ratio $k(0)$.

The convergence behavior of the optimal growth model is both important and remarkable in its simplicity. Such convergence results, which were first studied in the context of finite-horizon economies, are sometimes referred to as "*Turnpike Theorems*". To understand the meaning of this term, suppose that the economy ends at some date $T > 0$. How do optimal growth and capital accumulation look like in this economy? The early literature on optimal growth showed that as $T \to \infty$, the optimal capital-labor ratio sequence $\{k(t)\}_{t=0}^{T}$ would become arbitrarily close to k^* as defined by (6.49), but then in the last few periods, it would sharply decline to zero to satisfy the transversality condition (recall the discussion of the finite-horizon transversality condition in Section 6.6). The path of the capital-labor ratio thus resembles a turnpike approaching a highway as shown in Figure 6.1 (see Exercise 6.18).

6.9. Competitive Equilibrium Growth

Our main interest is not optimal growth, but equilibrium growth. A detailed analysis of competitive equilibrium growth will be presented in Chapter 8. For now a brief discussion of how the competitive equilibrium can be obtained from the optimal growth problem is sufficient. The Second Welfare Theorem, Theorem 5.7 of the previous chapter, implies that the optimal growth path characterized in the previous section also corresponds to an equilibrium growth path (in the sense that, it can be decentralized as a competitive equilibrium). In fact, since we have focused on an economy admitting a representative household, the most straightforward competitive allocation would be a symmetric one, where all households, each with the instantaneous utility function $u(c)$, make the same decisions and receive the same allocations. I now discuss this symmetric competitive equilibrium briefly.

Rather than appealing to the Second Welfare Theorem, a direct argument can also be used to show the equivalence of optimal and competitive growth problems. Suppose that each household starts with an endowment of capital stock K_0, meaning that the initial endowments are also symmetric (recall that there is a mass 1 of households and the total initial endowment of capital of the economy is K_0). The other side of the economy is populated by a large number of competitive firms, which are modeled using the aggregate production function.

The definition of a competitive equilibrium in this economy is standard.

Definition 6.3. *A **competitive equilibrium** consists of paths of consumption, capital stock, wage rates and rental rates of capital, $\{C(t), K(t), w(t), R(t)\}_{t=0}^{\infty}$, such that the representative household maximizes its utility given initial capital stock K_0 and prices $\{w(t), R(t)\}_{t=0}^{\infty}$, and the path of prices $\{w(t), R(t)\}_{t=0}^{\infty}$ is such that given the path of capital stock and labor $\{K(t), L(t)\}_{t=0}^{\infty}$ all markets clear.*

Households rent their capital to firms. As in the basic Solow model, they receive the competitive rental price of

$$R(t) = f'(k(t)),$$

and thus face a gross rate of return equal to

$$1 + r(t+1) = f'(k(t)) + (1-\delta) \tag{6.50}$$

for renting one unit of capital at time t in terms of date $t+1$ goods. Notice that the gross rate of return on assets is defined as $1+r$ (and this enables me to use r for the net interest rate both in the discrete-time and the continuous-time models). In addition, to capital income, households in this economy receive wage income for supplying their labor at the market wage of $w(t) = f(k(t)) - k(t) f'(k(t))$.

Now consider the maximization problem of the representative household:

$$\max_{\{c(t),a(t)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c(t))$$

subject to the flow budget constraint

$$a(t+1) = (1 + r(t+1)) a(t) - c(t) + w(t), \tag{6.51}$$

with $a(0) > 0$, where $a(t)$ denotes asset holdings at time t and as before, $w(t)$ is the wage income of the household (since labor supply is normalized to 1). The timing underlying the flow budget constraint (6.51) is that the household rents his capital or asset holdings, $a(t)$, to firms to be used as capital at time $t+1$. Out of the proceeds, it consumes and whatever is left, together with its wage earnings, $w(t)$, make up its asset holdings at the next date, $a(t+1)$. Once again, we impose the natural debt limit, which takes the form of equation (6.41) in the continuation of Example 6.5, and market clearing again implies $a(t) = k(t)$.

With an argument identical to that in Example 6.5, the Euler equation for the consumer maximization problem yields

$$u'(c(t)) = (1 + r(t+1)) \beta u'(c(t+1)). \tag{6.52}$$

Imposing steady state implies that $c(t) = c(t+1)$. Therefore, in steady state,

$$(1 + r(t+1)) \beta = 1.$$

Next, market clearing immediately implies that $1 + r(t+1)$ is given by (6.50), so the capital-labor ratio of the competitive equilibrium is given by

$$\beta \left[f'(k(t+1)) + (1-\delta) \right] = 1,$$

The steady state is given by

$$\beta \left[f'(k^*) + (1-\delta) \right] = 1.$$

These two equations are identical to eq.'s (6.49) and (6.50), which characterize the solution to the optimum growth problem. A similar argument establishes that the entire competitive equilibrium path is identical to the optimal growth path. Specifically, substituting for $1 + r(t+1)$ from (6.50) into (6.52),

$$u'(c(t)) = \beta \left[f'(k(t+1)) + (1-\delta) \right] u'(c(t+1)), \tag{6.53}$$

which is identical to (6.47). This condition also implies that given the same initial condition, the trajectory of capital-labor ratio in the competitive equilibrium will be identical to the behavior of the capital-labor ratio in the optimal growth path (see Exercise 6.21). This is, of course, exactly what should be expected given the Second (and First) Welfare Theorems.

6.10. Computation

All of the results presented here have been about existence of solutions and characterization of the form of the value functions, solutions and the properties of the policy functions or optimal plans. Dynamic programming techniques are also widely used in explicit (numerical) computations (Exercise 6.3 below provides one useful starting point in this respect). In particular, the recursive formulation of dynamic programming problems also presents an effective computational approach. This is particularly useful, since as hinted at by the discussion in Example 6.4, only certain very special dynamic optimization problems yield closed-form solutions. Therefore, economists, like engineers, must often use computational tools in order to obtain various qualitative and quantitative insights about solutions to optimization and equilibrium, and dynamic programming is the main starting point of these approaches.

Space restrictions preclude me from providing a discussion of various computational tools and how dynamic programming methods are used in numerical analysis. This should not be interpreted as downplaying the importance of computation in the study of economic growth and the usefulness of dynamic programming approaches in computation. The reader is encouraged to consult Judd (1998) for an excellent and thorough discussion of computational issues in economics and the role of dynamic programming. Ljungqvist and Sargent (2005) also provide a brief introduction to use of computational methods in macroeconomics.

6.11. Taking Stock

This chapter has been concerned with basic dynamic programming techniques for discrete-time infinite-dimensional problems. These techniques are not only essential for the study of economic growth, but are widely used in many diverse areas of macroeconomics and economic dynamics more generally. A good understanding of these techniques is essential for an appreciation of the *mechanics* of economic growth. In particular, they will shed light on how different models of economic growth work, how they can be improved and how they can be taken to the data. For this reason, this chapter is part of the main body of the text, rather than being relegated to the Appendices at the end.

This chapter also presented a number of applications of dynamic programming, including a preliminary analysis of the one-sector optimal growth problem. The reader will have already noted the parallels between this model and the basic Solow model discussed in Chapter 2. These parallels will be developed further in Chapter 8. I also briefly discussed the decentralization of the optimal growth path and the problem of utility maximization in a dynamic competitive equilibrium.

It is important to emphasize that the treatment in this chapter has assumed away a number of difficult technical issues. First, the focus has been on discounted problems, which are simpler than undiscounted problems. In economics, very few situations involve undiscounted objective functions ($\beta = 1$ rather than $\beta \in (0, 1)$). More important, throughout I have assumed that payoffs are bounded and the state vector x belongs to a compact subset of the Euclidean space, X. This rules out many interesting problems, such as endogenous growth, where the state vector grows over time. Almost all of the results presented here have equivalents for these cases, but these require somewhat more advanced treatments.

6.12. References and Literature

At some level the main idea of dynamic programming, the Principle of Optimality, is straightforward. Nevertheless, it is also a powerful concept and this will be best appreciated once a number of its implications are derived. The basic ideas of dynamic programming, including the Principle of Optimality, were introduced by Richard Bellman, in his famous monograph, Bellman (1957). Most of the basic results about finite and infinite-dimensional dynamic programming problems are contained in this monograph. Many of these ideas are

also contained in Shapley's (1953) study of stochastic games. Shapley analyzed the characterization of equilibrium points of zero-sum stochastic games. His formulation of these games anticipated what later became known as Markov Decision Problems, which are closely related to dynamic programming problems. Moreover, Shapley used ideas similar to the Principle of Optimality and the Contraction Mapping Theorem to show the existence of a unique solution to these dynamic zero-sum games. A more detailed treatment of Markov Decision Problems can be found in Puterman (1994), who also discusses the relationship between Shapley's (1953) work, the general theory of Markov Decision Problems and dynamic programming.

To the best of my knowledge, Karlin (1955) was the first one to provide a simple formal proof of the Principle of Optimality, which is similar to the one presented here. Denardo (1967) developed the use of the contraction mappings in the theory of dynamic programming. Puterman (1994) contain a more detailed analysis of discounted stochastic dynamic programming problems. Blackwell (1965) introduced the Blackwell's sufficient conditions for a contraction mapping and applied them in the context of stochastic discounted dynamic programming problems. The result on the differentiability of the value function was first proved in Benveniste and Scheinkman (1979). The second version of the proof of Theorem 6.6 follows their approach closely. The first version of the proof extends the earlier proof by Mirman and Zilcha (1975), which was only for the neoclassical growth model.

The most complete treatment of discounted stationary dynamic programming problems is in Stokey, Lucas and Prescott (1989). My treatment here is heavily influenced by theirs and borrows much from their insights. Relative to their treatment, some of the proofs have been simplified and I have limited the analysis to the case with compact sets and bounded payoff functions. The reader can find generalizations of Theorems 6.1-6.6 to certain problems with unbounded returns and choice sets in Stokey, Lucas and Prescott (1989, Chapter 4). Differently from Stokey, Lucas and Prescott, I also discuss and provide necessary and sufficient conditions for nonstationary infinite-horizon optimization problems.

The issue of the necessity and sufficiency of the transversality condition discussed in Theorem 6.10 is discussed briefly in Stokey, Lucas and Prescott (1989) and in greater detail in Ekeland and Scheinkman (1986) and Kamaihigashi (2001). Kamaihigashi (2001) provides weaker and more general versions of the transversality condition, though the version provided here is sufficient for most economic applications.

A much simpler but insightful exposition of dynamic programming is in Sundaram (1996), which also has a proof of Proposition 6.1 similar to the one given here.

Some useful references on the Contraction Mapping Theorem and its applications include Denardo (1967), Kolmogorov and Fomin (1970), Kreyszig (1978) and the eminently readable Bryant (1985), which contains applications of the Contraction Mapping Theorem to prove existence and uniqueness of solutions to differential equations and the Implicit Function Theorem.

6.13. Exercises

Exercise 6.1. Consider the formulation of the discrete-time optimal growth model as in Example 6.1. Show that with this formulation and Assumptions 1 and 2 from Chapter 2, the discrete-time optimal growth model satisfies Assumptions 6.1-6.5.

Exercise 6.2. * Prove that if for some $n \in \mathbb{N}$ T^n is a contraction over a complete metric space (S, d), then T has a unique fixed point in S.

Exercise 6.3. * Suppose that T is a contraction over the metric space (S, d) with modulus $\beta \in (0, 1)$. Prove that for any $z, z' \in S$ and $n \in \mathbb{N}$,

$$d\left(T^n z, z'\right) \leq \beta^n d\left(z, z'\right).$$

Discuss how this result can be useful in numerical computations.

Exercise 6.4. *

(1) Prove the claims made in Example 6.3 and show that the differential equation in (6.7) has a unique continuous solution.

(2) Recall eq. (6.8) from Example 6.3. Now apply the same argument to Tg and $T\tilde{g}$ and prove that

$$\left\|T^2 g - T^2 \tilde{g}\right\| \leq M^2 \times \frac{s^2}{2} \times \|g - \tilde{g}\|.$$

(3) Applying this argument recursively, prove that for any $n \in \mathbb{N}$, we have

$$\|T^n g - T^n \tilde{g}\| \leq M^n \times \frac{s^n}{n!} \times \|g - \tilde{g}\|.$$

(4) Using the previous inequality, the fact that for any $B < \infty$, $B^n/n! \to 0$ as $n \to 0$ and the result in Exercise 6.2, prove that the differential equation in (6.7) has a unique continuous solution on the compact interval $[0, s]$ for any $s \in \mathbb{R}_+$.

Exercise 6.5. * Recall the Implicit Function Theorem, Theorem A.26 in Appendix Chapter A. Here is a slightly simplified version of it: consider the function $\phi(y, x)$ such that that $\phi : \mathbb{R} \times [a, b] \to \mathbb{R}$ is continuously differentiable with bounded first derivatives. In particular, there exists $0 < m < M < \infty$ such that

$$m \leq \frac{\partial \phi(y, x)}{\partial y} \leq M$$

for all x and y. Then, the Implicit Function Theorem states that there exists a continuously differentiable function $y : [a, b] \to \mathbb{R}$ such that

$$\phi(y(x), x) = 0 \text{ for all } x \in [a, b].$$

Provide a proof for this theorem using the Contraction Mapping Theorem, Theorem 6.7 along the following lines:

(1) Let $\mathbf{C}^1([a, b])$ be the space of continuously differentiable functions defined on$[a, b]$. Then, for every $y \in \mathbf{C}^1([a, b])$, construct the operator

$$Ty = y(x) - \frac{\phi(y(x), x)}{M} \text{ for } x \in [a, b].$$

Show that $T : \mathbf{C}^1([a,b]) \to \mathbf{C}^1([a,b])$ and is a contraction.

(2) Applying Theorem 6.7 derive the Implicit Function Theorem.

Exercise 6.6. * Prove that T defined in (6.18) is a contraction.

Exercise 6.7. Let us return to Example 6.4.

(1) Prove that the law of motion of capital stock given by (6.34) monotonically converges to a unique steady state value of k^* starting with any $k_0 > 0$. What happens to the level of consumption along the transition path?

(2) Now suppose that instead of (6.34), you hypothesize that

$$\pi(x) = ax^{\alpha} + bx + c.$$

Verify that the same steps will lead to the conclusion that $b = c = 0$ and $a = \beta a$.

(3) Now let us characterize the explicit solution by guessing and verifying the form of the value function. In particular, make the following guess: $V(x) = A \ln x$, and using this together with the first-order conditions derive the explicit form solution.

Exercise 6.8. Consider the following discrete-time optimal growth model with full depreciation:

$$\max_{\{c(t),k(t)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left(c(t) - \frac{a}{2} c(t)^2 \right)$$

subject to

$$k(t+1) = Ak(t) - c(t)$$

and $k(0) = k_0$. Assume that $k(t) \in \left[0, \bar{k}\right]$ and $a < \bar{k}^{-1}$, so that the utility function is always increasing in consumption.

(1) Formulate this maximization problem as a dynamic programming problem.

(2) Argue without solving this problem that there will exist a unique value function $V(k)$ and a unique policy rule $c = \pi(k)$ determining the level of consumption as a function of the level of capital stock.

(3) Solve explicitly for $V(k)$ and $\pi(k)$ [Hint: guess the form of the value function $V(k)$, and use this together with the Bellman and Euler equations; verify that this guess satisfies these equations, and argue that this must be the unique solution].

Exercise 6.9. Consider Problem A1 or A2 with $x \in X \subset \mathbb{R}$ and suppose that Assumptions 6.1-6.3 and 6.5 hold and that $\partial^2 U(x,y)/\partial x \partial y \geq 0$. Show that the optimal policy function $y = \pi(x)$ is nondecreasing.

Exercise 6.10. Show that in Theorem 6.10, a sequence $\{x'(t)\}_{t=0}^{\infty}$ that satisfies the Euler equations, but not the transversality condition could yield a suboptimal plan.

Exercise 6.11. Consider the consumer utility maximization problem in Example 6.5, without the borrowing limits (that is, only with the flow budget constraint). Show that for any path of consumption $\{c(t)\}_{t=0}^{\infty}$ that satisfies the flow budget constraint starting with an initial asset level $a(0)$, so does the path of consumption $\{c'(t)\}_{t=0}^{\infty}$, where $c'(t) = c(t) + \gamma$ for $\gamma > 0$. Using this result, show that the consumer can reach infinite utility by choosing a

consumption path that leads to $\lim_{t\to\infty} a(t) = -\infty$. Explain why such paths can not be feasible in market economies.

Exercise 6.12. Consider Example 6.5, with $w(t) = w$ and $r(t) = r$ for all t.

(1) Show that when $r \leq \beta - 1$, $a(t) \in [\underline{a}, \bar{a}]$ for all t, where recall that $\bar{a} \equiv a(0) + w/r$ and $\underline{a} = -w/r$.
(2) Show that when $r > \beta - 1$, there does not exist $\bar{a} < \infty$ such that $a(t) \in [\underline{a}, \bar{a}]$ for all t.

Exercise 6.13. Consider the continuation of Example 6.5.

(1) Show that the natural debt limit at time t takes the form (6.41).
(2) Suppose that
$$\bar{W} = \max_{t=0,1,...} \left\{ \sum_{s=0}^{\infty} \left(\prod_{j=0}^{s-1} \frac{1}{1+r(t+j)} \right) w(t+s) \right\}$$
and that $\bar{W} < \infty$. Provide conditions such that $a(t) \in \left[-\bar{W}, a(0) + \bar{W}\right]$ for all t.
(3) Using the transversality condition together with $a(0)$, $\{w(t)\}_{t=0}^{\infty}$ and $\{r(t)\}_{t=0}^{\infty}$, find an expression implicitly determining the initial level of consumption, $c(0)$. What happens to this level of consumption when $a(0)$ increases?
(4) Suppose that $r(t) = r$ for all t. Consider a change in the earnings profile to a new sequence $\{\tilde{w}(t)\}_{t=0}^{\infty}$ such that for some $T < \infty$, $w(t) < \tilde{w}(t)$ for all $t < T$, $w(t) \geq \tilde{w}(t)$ for all $t \geq T$, and $\sum_{t=0}^{\infty} (1+r)^{-t} w(t) = \sum_{t=0}^{\infty} (1+r)^{-t} \tilde{w}(t)$. What is the effect of this on the initial consumption level and the consumption path? Provide a detailed economic intuition for this result.

Exercise 6.14. * Consider the nonstationary infinite-horizon optimization problem considered in Section 6.7. Provide conditions on G and U such that the operator defined by
$$TV(t,x) = \max_{y \in G(t,x)} \{U(t,x,y) + \beta V(t+1,y)\}$$
for $V \in \mathbf{C}(\mathbb{Z}_+ \times X)$ is a contraction. Using these conditions state and sketch the proofs of the appropriate equivalents of Theorems 6.2-6.6.

Exercise 6.15. Consider the following discrete-time optimal growth model
$$\max_{\{c(t),k(t)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} u(c(t))$$
subject to
$$k(t+1) = k(t) - c(t)$$
and
$$k(0) = k_0 \in (0, \infty).$$
Assume that $u(\cdot)$ is a strictly increasing, strictly concave and bounded function. Prove that there exists no optimal solution to this problem. Explain why.

Exercise 6.16. Consider the following discrete-time optimal growth model with full depreciation:

$$\max_{\{c(t),k(t)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u\left(c\left(t\right)\right)$$

subject to

$$k\left(t+1\right) = f\left(k\left(t\right)\right) - c\left(t\right)$$

and

$$k\left(0\right) = k_0 > 0.$$

Assume that $u\left(\cdot\right)$ is strictly concave and increasing for $c \geq 0$, and $f\left(\cdot\right)$ is concave and increasing.

(1) Formulate this maximization problem as a dynamic programming problem.
(2) Prove that there exists unique value function $V\left(k\right)$ and a unique policy rule $c = \pi\left(k\right)$, and that $V\left(k\right)$ is continuous and strictly concave and $\pi\left(k\right)$ is continuous and increasing.
(3) When will $V\left(k\right)$ be differentiable?
(4) Assuming that $V\left(k\right)$ and all the other functions are differentiable, characterize the Euler equation that determines the optimal path of consumption and capital accumulation.
(5) Is this Euler equation enough to determine the path of k and c? If not, what other condition do we need to impose? Write down this condition and explain intuitively why it makes sense.

Exercise 6.17. Prove that, as claimed in Proposition 6.4, in the basic discrete-time optimal growth model, the optimal consumption plan $c\left(k\right)$ is nondecreasing, and when the economy starts with $k_0 < k^*$, the unique equilibrium involves $c\left(t\right) \uparrow c^*$.

Exercise 6.18. * Consider the finite-horizon optimal growth model described at the end of Section 6.8. Let the optimal capital-labor ratio sequence of the economy with horizon T be denoted by $\left\{k^T\left(t\right)\right\}_{t=0}^{T}$ with $k^T\left(0\right) = k_0 > 0$. Show that for every $\varepsilon > 0$, there exists $T < \infty$ and $t' < T$ such that $\left|k^T\left(t'\right) - k^*\right| < \varepsilon$. Show that $k^T\left(T\right) = 0$. Then, assuming that k_0 is sufficiently small, show that the optimal capital-labor ratio sequence looks as in Figure 6.1.

Exercise 6.19. Prove that as claimed in the proof of Proposition 6.3, Assumption 2 implies that $s\left(\varepsilon\right) > \varepsilon$ for ε sufficiently small. Provide an intuition for this result.

Exercise 6.20. * Provide a proof of Proposition 6.1 without the differentiability assumption on the utility function $u\left(\cdot\right)$ imposed in Assumption 3′.

Exercise 6.21. Prove that the optimal growth path starting with capital-labor ratio k_0, which satisfies (6.47) is identical to the competitive equilibrium starting with capital-labor ratio and satisfying the same condition (or equivalently, eq. (6.53)).

CHAPTER 7

Review of the Theory of Optimal Control

The previous chapter introduced the basic tools of dynamic optimization in discrete time. I now review a number of basic results in dynamic optimization in continuous time—particularly the so-called *optimal control* approach. Both dynamic optimization in discrete time and in continuous time are useful tools for macroeconomics and other areas of dynamic economic analysis. One approach is not superior to another; instead, certain problems become simpler in discrete time while certain others are naturally formulated in continuous time.

Continuous-time optimization introduces a number of new mathematical issues. The main reason is that even with a finite horizon, the maximization is with respect to an infinite-dimensional object (in fact an entire function, $y : [t_0, t_1] \to \mathbb{R}$). This requires a brief review of some basic ideas from the *calculus of variations* and from the theory of optimal control. The main tools and ideas that are necessary for this book are straightforward and are presented here in the simplest way. Nevertheless, a reader who simply wishes to apply these tools may decide skim most of this chapter, focusing on the main theorems, especially Theorems 7.13 and 7.14, and their application to the canonical continuous-time optimal growth problem in Section 7.7.

In the rest of this chapter, I first review the finite-horizon continuous-time maximization problem and provide the simplest treatment of this problem (which is more similar to the calculus of variations than to optimal control). I then present the more powerful theorems from the theory of optimal control as developed by Pontryagin and co-authors.

The canonical continuous-time optimization problem can be written as

$$\max_{\mathbf{x}(t),\mathbf{y}(t)} W\left(\mathbf{x}\left(t\right),\mathbf{y}\left(t\right)\right) \equiv \int_0^{t_1} f\left(t,\mathbf{x}\left(t\right),\mathbf{y}\left(t\right)\right) dt$$

subject to

$$\dot{\mathbf{x}}\left(t\right) = G\left(t,\mathbf{x}\left(t\right),\mathbf{y}\left(t\right)\right)$$

and

$$\mathbf{y}\left(t\right) \in \mathcal{Y}\left(t\right), \mathbf{x}\left(t\right) \in \mathcal{X}\left(t\right) \text{ for all } t, \text{ and } \mathbf{x}\left(0\right) = \mathbf{x}_0,$$

where for each t, $\mathbf{x}\left(t\right)$ and $\mathbf{y}\left(t\right)$ are finite-dimensional vectors (i.e., $\mathcal{X}\left(t\right) \subset \mathbb{R}^{K_x}$ and $\mathcal{Y}\left(t\right) \subset \mathbb{R}^{K_y}$, where $K_x, K_y \in \mathbb{N}$). In addition, $f : \mathbb{R}\times\mathbb{R}^{K_x}\times\mathbb{R}^{K_y} \to \mathbb{R}$ and $G : \mathbb{R}\times\mathbb{R}^{K_x}\times\mathbb{R}^{K_y} \to \mathbb{R}^{K_x}$. Here $\mathbf{x}$ is the *state* variable. Its behavior is governed by a vector-valued differential equation (i.e., a set of differential equations) given the behavior of the *control* variables $\mathbf{y}\left(t\right)$. The end of the planning horizon t_1 can be equal to infinity. The function $W\left(\mathbf{x}\left(t\right),\mathbf{y}\left(t\right)\right)$ denotes the

value of the objective function when controls are given by $\mathbf{y}(t)$ and the resulting behavior of the state variable is summarized by $\mathbf{x}(t)$. I also refer to f as the *objective function* (or the payoff function) and to g as the *constraint function*.

This problem formulation is general enough to incorporate discounting and time-varying constraints, since both the instantaneous payoff function f and the constraint function g depend directly on time in an arbitrary fashion. I start with the finite-horizon case and then treat the infinite-horizon maximization problem, focusing particularly on the case where there is exponential discounting.

7.1. Variational Arguments

Consider the following finite-horizon continuous-time problem

$$\max_{x(t),y(t),x_1} W(x(t),y(t)) \equiv \int_0^{t_1} f(t,x(t),y(t))\,dt \tag{7.1}$$

subject to

$$\dot{x}(t) = g(t,x(t),y(t)) \tag{7.2}$$

and

$$y(t) \in \mathcal{Y}(t) \text{ and } x(t) \in \mathcal{X}(t) \text{ for all } t,\ x(0) = x_0 \text{ and } x(t_1) = x_1. \tag{7.3}$$

Here the state variable $x(t) \in \mathcal{X}(t) \subset \mathbb{R}$ is one-dimensional and its behavior is governed by the differential equation (7.2). The control variable $y(t)$ must belong to the set $\mathcal{Y}(t) \subset \mathbb{R}$. Consequently, now, $f : \mathbb{R}\times\mathbb{R}\times\mathbb{R} \to \mathbb{R}$ and $g : \mathbb{R}\times\mathbb{R}\times\mathbb{R} \to \mathbb{R}$.

I assume that $\mathcal{Y}(t)$ is nonempty and convex. A pair of functions $(x(t),y(t))$ that jointly satisfy (7.2) and (7.3) are referred to as an *admissible pair*.[1] Throughout, as in the previous chapter, suppose that the value of the objective function is finite, that is, $W(x(t),y(t)) < \infty$ for any admissible pair $(x(t),y(t))$.

Let us first suppose that $t_1 < \infty$, so that we have a finite-horizon optimization problem. Notice that there is also a terminal value constraint $x(t_1) = x_1$, but x_1 is included as an additional choice variable. This implies that the terminal value of the state variable x is *free*. In the context of finite-horizon economic problems, the formulation where x_1 is *not* a choice variable may be simpler (see Example 7.1), but it is more natural to start with the case where the terminal value x_1 is free.

[1] More precisely, an admissible pair is $(x(t),y(t))$ such that $x(t)$ is *absolutely continuous*, $y(t)$ is (Lebesgue) *measurable* and $(x(t),y(t))$ satisfies (7.2) almost everywhere. Absolute continuity is defined in Appendix Chapter A and is slightly stronger than continuity. In particular, a function defined by $F(x) = \int_0^x f(s)\,ds$ for $x \in [0,X]$ is absolutely continuous on $[0,X]$ if $f(s)$ is piecewise continuous on $[0,X]$ (or less stringently, if $f(s)$ is measurable). Fact A.16 in Appendix Chapter A shows that differentiability (almost everywhere) implies absolute continuity. Since everywhere in the text, except in Section 7.6, $x(t)$ is taken to be differentiable and thus explicit reference to absolute continuity is suppressed. The requirement that $y(t)$ is measurable is not very restrictive either. Nevertheless, since I do not introduce concepts from measure theory explicitly, I will not provide a formal definition for this. Until Section 7.6 below, there is not much generality lost in suppressing the measurability requirement on $y(t)$ and referring to $x(t)$ as continuous.

In addition, to simplify the exposition, throughout I assume that f and g are continuously differentiable functions of x, y and t (that is, for all $(x(t), y(t)) \in \mathcal{X}(t) \times \mathcal{Y}(t)$ and for all t), and I simply state this as "f and g are continuously differentiable".

The challenge in characterizing the optimal solution to this problem lies in two features:

(1) We are choosing a function $y : [0, t_1] \to \mathcal{Y}$ rather than a vector or a finite-dimensional object.

(2) The constraint takes the form of a differential equation, rather than a set of inequalities or equalities.

These features make it difficult for us to know what type of optimal policy to look for. For example, y may be a highly discontinuous function. It may also hit the boundary of the feasible set—thus corresponding to a "corner solution". Fortunately, in most economic problems there will be enough structure to make optimal solutions continuous functions. Moreover, in most macroeconomic and growth applications, the Inada conditions (Assumption 2) make sure that the optimal solutions to the relevant dynamic optimization problems lie in the interior of the feasible set. These features considerably simplify the characterization of the optimal solution. In fact, when y is a continuous function of time and lies in the interior of the feasible set, it can be characterized by using variational arguments similar to those developed by Euler, Lagrange and others in the context of the theory of calculus of variations. Since these tools are not only simpler but also more intuitive, I start with these variational arguments.

The *variational principle* of the calculus of variations simplifies the above maximization problem by first assuming that a continuous solution (function) $\hat{y}$ that lies everywhere in the interior of the set $\mathcal{Y}$ (with corresponding state variable $\hat{x}$ everywhere in the interior of $\mathcal{X}$) exists, and then characterizes what features this solution must have in order to reach an optimum (for the relationship of the results here to the calculus of variations, see Exercise 7.3).

More formally let us assume that $(\hat{x}(t), \hat{y}(t))$ is an admissible pair such that $\hat{y}(\cdot)$ is continuous on $[0, t_1]$, $\hat{y}(t) \in \text{Int}\mathcal{Y}(t)$ $\hat{x}(t) \in \text{Int}\mathcal{X}(t)$ for all t, and we have

$$W(\hat{x}(t), \hat{y}(t)) \geq W(x(t), y(t))$$

for any other admissible pair $(x(t), y(t))$.

The important and stringent assumption here is that $(\hat{x}(t), \hat{y}(t))$ is an optimal solution that never hits the boundary and that does not involve any discontinuities. Even though this will be a feature of optimal controls in most economic applications, in purely mathematical terms it is a strong assumption. Recall, for example, that the previous chapter did not make such an assumption and instead started with a result on the existence of solutions and then proceeded to characterizing the properties of this solution (such as continuity and differentiability of the value function). However, the problem of continuous-time optimization is sufficiently difficult that proving existence of solutions is not a trivial matter. I return to a discussion of this issue below, but for now I follow the standard practice and assume

that an *interior continuous solution*, with $\hat{y}(t) \in \text{Int}\mathcal{Y}(t)$ and $\hat{x}(t) \in \text{Int}\mathcal{X}(t)$, exists (where $\hat{y}(t) \in \text{Int}\mathcal{Y}(t)$ means that $\hat{y}(t)$ is in the interior of the set $\mathcal{Y}(t)$ for all t, and likewise for $\hat{x}(t) \in \text{Int}\mathcal{X}(t)$). Note also that since the behavior of the state variable x is given by the differential equation (7.2), when $y(t)$ is continuous, $\dot{x}(t)$ will also be continuous, so that $x(t)$ is in fact continuously differentiable. When $y(t)$ is piecewise continuous, $x(t)$ will be, correspondingly, piecewise smooth.

I now exploit these features to derive *necessary* conditions for an optimal path of this form. To do this, consider the following *variation*

$$y(t,\varepsilon) \equiv \hat{y}(t) + \varepsilon\eta(t),$$

where $\eta(t)$ is an arbitrary *fixed* continuous function and $\varepsilon \in \mathbb{R}$ is a real number. I refer to this as a variation, because given $\eta(t)$, varying ε gives different processes for the controls. The problem, of course, is that some of these may be infeasible, i.e., $y(t,\varepsilon) \notin \mathcal{Y}(t)$ for some t. However, since $\hat{y}(t) \in \text{Int}\mathcal{Y}(t)$, and a continuous function over a compact set $[0,t_1]$ is bounded, for any fixed $\eta(\cdot)$ function, we can always find $\varepsilon'_\eta > 0$ such that

$$y(t,\varepsilon) \equiv \hat{y}(t) + \varepsilon\eta(t) \in \text{Int}\mathcal{Y}(t)$$

for all $\varepsilon \in \left[-\varepsilon'_\eta, \varepsilon'_\eta\right]$, so that $y(t,\varepsilon)$ constitutes a *feasible variation*. Consequently, variational arguments can be used for sufficiently small ε's. The fact that we have to look at small ε's is not a drawback for deriving necessary conditions for optimality. In analogy with standard calculus, necessary conditions require that there should be no small change in controls that increase the value of the objective function.

To prepare for these arguments, let us fix an arbitrary $\eta(\cdot)$, and define $x(t,\varepsilon)$ as the path of the state variable corresponding to the path of control variable $y(t,\varepsilon)$. This implies that $x(t,\varepsilon)$ is given by:

$$\dot{x}(t,\varepsilon) \equiv g(t, x(t,\varepsilon), y(t,\varepsilon)) \text{ for all } t \in [0,t_1], \text{ with } x(0,\varepsilon) = x_0. \tag{7.4}$$

Since $x(t) \in \text{Int}\mathcal{X}(t)$ for all t, for ε sufficiently small, in particular, for $\varepsilon \in [-\varepsilon_\eta, \varepsilon_\eta] \subset \left[-\varepsilon'_\eta, \varepsilon'_\eta\right]$ for some $\varepsilon_\eta \leq \varepsilon'_\eta$, we also have that $x(\varepsilon,t) \in \mathcal{X}(t)$. In view of this, for $\varepsilon \in [-\varepsilon_\eta, \varepsilon_\eta]$, define:

$$\begin{aligned} \mathcal{W}(\varepsilon) &\equiv W(x(t,\varepsilon), y(t,\varepsilon)) \\ &= \int_0^{t_1} f(t, x(t,\varepsilon), y(t,\varepsilon))\, dt. \end{aligned} \tag{7.5}$$

By the fact that $\hat{y}(t)$ is optimal, and that for $\varepsilon \in [-\varepsilon_\eta, \varepsilon_\eta]$, $y(t,\varepsilon)$ and $x(t,\varepsilon)$ are feasible,

$$\mathcal{W}(\varepsilon) \leq \mathcal{W}(0) \text{ for all } \varepsilon \in [-\varepsilon_\eta, \varepsilon_\eta].$$

Next, rewrite the eq. (7.4), so that

$$g(t, x(t,\varepsilon), y(t,\varepsilon)) - \dot{x}(t,\varepsilon) \equiv 0$$

for all $t \in [0,t_1]$. This implies that for *any* function $\lambda : [0,t_1] \to \mathbb{R}$,

$$\int_0^{t_1} \lambda(t) \left[g(t, x(t,\varepsilon), y(t,\varepsilon)) - \dot{x}(t,\varepsilon)\right] dt = 0, \tag{7.6}$$

since the term in square brackets is identically equal to zero. In what follows, suppose that the function $\lambda(\cdot)$ is continuously differentiable. This function, when chosen suitably, will be the *costate* variable, with a similar interpretation to the Lagrange multipliers in standard (constrained) optimization problems. As with Lagrange multipliers, this will not be true for any $\lambda(\cdot)$ function, but only for $\lambda(\cdot)$ that is chosen appropriately to play the role of the costate variable.

Adding (7.6) to (7.5),

$$\mathcal{W}(\varepsilon) \equiv \int_0^{t_1} \{f(t, x(t,\varepsilon), y(t,\varepsilon)) + \lambda(t)[g(t, x(t,\varepsilon), y(t,\varepsilon)) - \dot{x}(t,\varepsilon)]\} dt. \tag{7.7}$$

To evaluate (7.7), let us first consider the integral $\int_0^{t_1} \lambda(t)\dot{x}(t,\varepsilon)dt$. Integrating this expression by parts (see Theorem B.3 in Appendix Chapter B),

$$\int_0^{t_1} \lambda(t)\dot{x}(t,\varepsilon)dt = \lambda(t_1)x(t_1,\varepsilon) - \lambda(0)x_0 - \int_0^{t_1} \dot{\lambda}(t)x(t,\varepsilon)dt.$$

Substituting this expression back into (7.7):

$$\begin{aligned}\mathcal{W}(\varepsilon) \equiv & \int_0^{t_1} \left[f(t, x(t,\varepsilon), y(t,\varepsilon)) + \lambda(t) g(t, x(t,\varepsilon), y(t,\varepsilon)) + \dot{\lambda}(t)x(t,\varepsilon)\right] dt \\ & -\lambda(t_1)x(t_1,\varepsilon) + \lambda(0)x_0.\end{aligned}$$

Recall that f and g are continuously differentiable, and $y(t,\varepsilon)$ is continuously differentiable in ε by construction, which also implies that $x(t,\varepsilon)$ is continuously differentiable in ε. Let us denote the partial derivatives of x and y with respect to ε by x_ε and y_ε, and the partial derivatives of f and g by f_t, f_x, f_y, and so on. Differentiating the previous expression with respect to ε (making use of Leibniz's Rule, Theorem B.4 in Appendix Chapter B),

$$\begin{aligned}\mathcal{W}'(\varepsilon) \equiv & \int_0^{t_1} \left[f_x(t, x(t,\varepsilon), y(t,\varepsilon)) + \lambda(t) g_x(t, x(t,\varepsilon), y(t,\varepsilon)) + \dot{\lambda}(t)\right] x_\varepsilon(t,\varepsilon) dt \\ & + \int_0^{t_1} [f_y(t, x(t,\varepsilon), y(t,\varepsilon)) + \lambda(t) g_y(t, x(t,\varepsilon), y(t,\varepsilon))]\eta(t) dt \\ & -\lambda(t_1) x_\varepsilon(t_1,\varepsilon).\end{aligned}$$

Let us next evaluate this derivative at $\varepsilon = 0$ to obtain:

$$\begin{aligned}\mathcal{W}'(0) \equiv & \int_0^{t_1} \left[f_x(t, \hat{x}(t), \hat{y}(t)) + \lambda(t) g_x(t, \hat{x}(t), \hat{y}(t)) + \dot{\lambda}(t)\right] x_\varepsilon(t,0) dt \\ & + \int_0^{t_1} [f_y(t, \hat{x}(t), \hat{y}(t)) + \lambda(t) g_y(t, \hat{x}(t), \hat{y}(t))]\eta(t) dt \\ & -\lambda(t_1) x_\varepsilon(t_1,0).\end{aligned}$$

where, as above, $\hat{x}(t) = x(t, \varepsilon = 0)$ denotes the path of the state variable corresponding to the optimal plan, $\hat{y}(t)$. As with standard finite-dimensional optimization, if there exists some function $\eta(t)$ for which $\mathcal{W}'(0) \neq 0$, this means that $W(x(t), y(t))$ can be increased and thus the pair $(\hat{x}(t), \hat{y}(t))$ could not be an optimal solution. Consequently, optimality requires that

$$\mathcal{W}'(0) \equiv 0 \text{ for all } \eta(t). \tag{7.8}$$

Recall that the expression for $\mathcal{W}'(0)$ applies for any continuously differentiable $\lambda(t)$ function. Clearly, not all such functions $\lambda(\cdot)$ will play the role of a costate variable. Instead, as it is the case with Lagrange multipliers, the function $\lambda(\cdot)$ has to be chosen appropriately. In particular, let us consider the function $\lambda(t)$ that is a solution to the differential equation

$$\dot{\lambda}(t) = -\left[f_x\left(t, \hat{x}(t), \hat{y}(t)\right) + \lambda(t) g_x\left(t, \hat{x}(t), \hat{y}(t)\right)\right], \tag{7.9}$$

with boundary condition $\lambda(t_1) = 0$. Since f_x and g_x are continuous functions (by the assumption that f and g are continuously differentiable), they are bounded for $t \in [0, t_1]$. Theorem B.8 in Appendix Chapter B then implies that this differential equation has a solution (for given interior solution $(\hat{x}(t), \hat{y}(t))$). Since $\eta(t)$ arbitrary, $\lambda(t)$ and $(\hat{x}(t), \hat{y}(t))$ need to be such that

$$f_y\left(t, \hat{x}(t), \hat{y}(t)\right) + \lambda(t) g_y\left(t, \hat{x}(t), \hat{y}(t)\right) \equiv 0 \text{ for all } t \in [0, t_1]. \tag{7.10}$$

This condition is then sufficient for

$$\int_0^{t_1} \left[f_y\left(t, \hat{x}(t), \hat{y}(t)\right) + \lambda(t) g_y\left(t, \hat{x}(t), \hat{y}(t)\right)\right] \eta(t)\, dt = 0 \text{ for all } \eta(t).$$

Conversely, if (7.10) were not satisfied, there would exist some variation $\eta(t)$ that would make the previous integral either positive or negative and thus the necessary condition (7.8) could not be satisfied. This establishes the necessary conditions for an interior continuous solution to the problem of maximizing (7.1) subject to (7.2) and (7.3)—that there should exist a continuously differentiable function $\lambda(\cdot)$ that satisfies (7.9), (7.10) and $\lambda(t_1) = 0$. Exercise 7.1 outlines a slightly different argument to reach the same conclusion.

The condition that $\lambda(t_1) = 0$ is the *transversality condition* of continuous-time optimization problems. This is naturally related to the transversality condition we encountered in the previous chapter. Intuitively, this condition captures the fact that after the planning horizon, there is no value to having more x.

This derivation, which builds on the standard arguments of calculus of variations, has therefore established the following theorem.[2]

THEOREM 7.1. ***(Necessary Conditions)*** *Consider the problem of maximizing (7.1) subject to (7.2) and (7.3), with f and g continuously differentiable. Suppose that this problem has an interior continuous solution $\hat{y}(t) \in Int\mathcal{Y}(t)$ and $\hat{x}(t) \in Int\mathcal{X}(t)$. Then, there exists a continuously differentiable costate function $\lambda(\cdot)$ defined over $t \in [0, t_1]$ such that (7.2), (7.9) and (7.10) hold, and moreover $\lambda(t_1) = 0$.*

The hypothesis that "an interior continuous solution exists" is important in this theorem. The conditions are *necessary* conditions only for such a solution. Solutions that are not interior need not satisfy (7.2), (7.9) and (7.10). Our strategy will be to use these necessary conditions to locate a candidate optimal path and then use sufficiency conditions, developed below, to verify that they indeed correspond to optimal solutions.

[2]See below for the alternative proof of Theorem 7.9, which generalizes the results in Theorem 7.2 in a number of dimensions.

Let us next consider a slightly different version of Theorem 7.1, where the terminal value of the state variable, x_1, is fixed, so that the maximization problem is

$$\max_{x(t),y(t)} W\left(x\left(t\right),y\left(t\right)\right) \equiv \int_0^{t_1} f\left(t,x\left(t\right),y\left(t\right)\right)dt, \tag{7.11}$$

subject to (7.2) and (7.3). The only difference is that there is no longer a choice over the terminal value of the state variable, x_1. In this case, we have:

THEOREM 7.2. ***(Necessary Conditions II)*** *Consider the problem of maximizing (7.11) subject to (7.2) and (7.3), with f and g continuously differentiable. Suppose that this problem has an interior continuous solution $\hat{y}\left(t\right) \in Int\mathcal{Y}\left(t\right)$ and $\hat{x}\left(t\right) \in Int\mathcal{X}\left(t\right)$. Then, there exists a continuously differentiable costate function $\lambda\left(\cdot\right)$ defined over $t \in [0,t_1]$ such that (7.2), (7.9) and (7.10) hold.*

PROOF. The proof is similar to the argument used in Exercise 7.1 for establishing Theorem 7.1, with the main change that now $x\left(t_1,\varepsilon\right)$ must equal x_1 for feasibility, so $x_\varepsilon\left(t_1,0\right)=0$ and $\lambda\left(t_1\right)$ is unrestricted. Exercise 7.5 asks you to complete the details. □

The new feature in this theorem is that the transversality condition $\lambda\left(t_1\right)=0$ is no longer present, but we need to know what the terminal value of the state variable x should be.[3] Let us start with an application of the necessary conditions in Theorem 7.2 to a simple economic problem. More interesting economic examples are provided later in the chapter and in the exercises.

EXAMPLE 7.1. Consider a relatively common application of the techniques developed so far, which is the problem of utility-maximizing choice of consumption plan by an individual that lives between dates 0 and 1 (perhaps the most common application of these techniques is a physical one, that of finding the shortest curve between two points in the plane, see Exercise 7.4). The individual has an instantaneous utility function $u\left(c\right)$ and discounts the future exponentially at the rate $\rho > 0$. Suppose that $u : [0,1] \to \mathbb{R}$ is a strictly increasing, continuously differentiable and strictly concave function. The individual starts with a level of assets equal to $a\left(0\right) > 0$, earns an interest rate r on his asset holdings and also has a constant flow of labor earnings equal to w. Let us also suppose that the individual can never have negative asset position, so that $a\left(t\right) \geq 0$ for all t. Therefore, the problem of the individual can be written as

$$\max_{[c(t),a(t)]_{t=0}^{1}} \int_0^1 \exp\left(-\rho t\right) u\left(c\left(t\right)\right) dt$$

subject to

$$\dot{a}\left(t\right) = r\left[a\left(t\right) + w - c\left(t\right)\right]$$

[3]It is also worth noting that the hypothesis that there exists an interior solution is more restrictive in this case than in Theorem 7.1. This is because the set of admissible control and state variables, Ω, may be empty or have an empty interior, making it impossible that an interior solution exists. See Exercise 7.23 for an example and Section 7.6 for a formal definition of the set Ω.

and $a(t) \geq 0$, with an initial value of $a(0) > 0$. In this problem, consumption is the control variable, while the asset holdings of the individual are the state variable.

To be able to apply Theorem 7.2, we need a terminal condition for $a(t)$, that is, some value a_1 such that $a(1) = a_1$. The economics of the problem makes it clear that the individual would not like to have any positive level of assets at the end of his planning horizon (since he could consume all of these at date $t = 1$ or slightly before, and $u(\cdot)$ is strictly increasing). Therefore, we must have $a(1) = 0$.

With this observation, Theorem 7.2 provides the following the necessary conditions for an interior continuous solution: there exists a continuously differentiable costate variable $\lambda(t)$ such that the optimal path of consumption and asset holdings, $(\hat{c}(t), \hat{a}(t))$, satisfy a consumption Euler equation similar to eq. (6.35) in Example 6.5 in the previous chapter:

$$\exp(-\rho t) u'(\hat{c}(t)) = \lambda(t) r.$$

In particular, this equation can be rewritten as $u'(\hat{c}(t)) = \beta r \lambda(t)$, with $\beta = \exp(-\rho t)$, and would be almost identical to eq. (6.35), except for the presence of $\lambda(t)$ instead of the derivative of the value function. But as shown below, $\lambda(t)$ is exactly the derivative of the value function, so that the consumption Euler equations in discrete and continuous time are identical. This is of course not surprising, since they capture the same economic phenomenon, in slightly different mathematical formulations.

The next necessary condition determines the behavior of $\lambda(t)$ as

$$\dot{\lambda}(t) = -\lambda(t) r.$$

Now using this condition and differentiating $u'(\hat{c}(t)) = \beta r \lambda(t)$ yields a differential equation in consumption. This differential equation, derived in the next chapter in a somewhat more general context, will be the key consumption Euler equation in continuous time. Leaving the derivation of this equation to the next chapter, we can make progress here by simply integrating this condition to obtain

$$\lambda(t) = \lambda(0) \exp(-rt).$$

Combining this with the first-order condition for consumption yields a straightforward expression for the optimal consumption level at time t:

$$\hat{c}(t) = u'^{-1}[R\lambda(0) \exp((\rho - r) t)],$$

where $u'^{-1}[\cdot]$ is the inverse function of the marginal utility u'. It exists and is strictly decreasing in view of the fact that u is strictly concave. This equation therefore implies that when $\rho = r$, so that the discount factor and the rate of return on assets are equal, the individual will have a constant consumption profile. When $\rho > r$, the argument of u'^{-1} is increasing over time, so consumption must be declining. This reflects the fact that the individual discounts the future more heavily than the rate of return, thus wishes to have a front-loaded consumption profile. In contrast, when $\rho < r$, the opposite reasoning applies and the individual chooses a back-loaded consumption profile. These are of course identical

to the conclusions reached in the discrete-time intertemporal consumer optimization problem in Example 6.5, in particular, eq. (6.37).

The only variable to determine in order to completely characterize the consumption profile is the initial value of the costate variable. This comes from the budget constraint of the individual together with the observation that the individual will run down all his assets by the end of his planning horizon, thus $a(1) = 0$. Now using the consumption rule, we have

$$\dot{a}(t) = R\left\{a(t) + w - u'^{-1}\left[R\lambda(0)\exp\left((\rho - R)t\right)\right]\right\}.$$

The initial value of the costate variable, $\lambda(0)$, then has to be chosen such that $a(1) = 0$. You are asked to complete the details of this step in Exercise 7.6.

Example 7.1 applied the results of Theorem 7.2. It may at first appear that Theorem 7.1 is more convenient to use than Theorem 7.2, since it would enable us to directly formulate the problem as one of dynamic optimization rather than first argue about what the terminal value of the state variable, $a(1)$, should be (based on economic reasoning as we did in Example 7.1). However, as the continuation of the previous example illustrates, this is not necessarily the case:

EXAMPLE 7.1 (CONTINUED). Let us try to apply Theorem 7.1 to the economic environment in Example 7.1. The first-order necessary conditions still give

$$\lambda(t) = \lambda(0)\exp(-Rt).$$

However, since $\lambda(1) = 0$, this is only possible if $\lambda(t) = 0$ for all $t \in [0, 1]$. But the necessary conditions still imply the Euler equation

$$\exp(-\rho t)\, u'(\hat{c}(t)) = \lambda(t) R,$$

which cannot be satisfied in view of the fact that $u' > 0$. This implies that when the terminal value of the assets, $a(1)$, is a choice variable, there exists no solution (at least no solution with an interior continuous control). How is this possible?

The answer is that Theorem 7.1 cannot be applied to this problem, because there is an additional constraint that $a(t) \geq 0$. We would need to consider a version of Theorem 7.1 with inequality constraints. The necessary conditions with inequality constraints are messier and more difficult to work with. Using a little bit of economic reasoning to observe that the terminal value of the assets must be equal to zero and then applying Theorem 7.2 simplifies the analysis considerably.

This discussion highlights that it may also be useful to have a version of Theorem 7.2, where the terminal condition is specified as an inequality—as $x(t_1) \geq x_1$ rather than as $x(t_1) = x_1$.

THEOREM 7.3. ***(Necessary Conditions III)*** *Consider the problem of maximizing (7.1) subject to (7.2) and to* $(x(t), y(t)) \in \mathcal{X}(t) \times \mathcal{Y}(t)$ *for all* t, $x(0) = x_0$ *and* $x(t_1) \geq x_1$, *with* f *and* g *continuously differentiable. Suppose that this problem has an interior continuous*

solution $\hat{y}(t) \in Int\mathcal{Y}(t)$ *and* $\hat{x}(t) \in Int\mathcal{X}(t)$*. Then, there exists a continuously differentiable costate function* $\lambda(\cdot)$ *defined over* $t \in [0, t_1]$ *such that (7.2), (7.9) and (7.10) hold, and moreover* $\lambda(t_1)(x(t_1) - x_1) = 0$*.*

PROOF. See Exercise 7.9. □

7.2. The Maximum Principle: A First Look

7.2.1. The Hamiltonian and the Maximum Principle. By analogy with the Lagrangian, a much more economical way of expressing Theorem 7.2 is to construct the *Hamiltonian:*[4]

$$H(t, x(t), y(t), \lambda(t)) \equiv f(t, x(t), y(t)) + \lambda(t) g(t, x(t), y(t)). \tag{7.12}$$

Since f and g are continuously differentiable, so is H. Denote the partial derivatives of the Hamiltonian with respect to $x(t)$, $y(t)$ and $\lambda(t)$, by H_x, H_y and H_λ. Theorem 7.2 then immediately leads to the following result:

THEOREM 7.4. ***(Simplified Maximum Principle)*** *Consider the problem of maximizing (7.1) subject to (7.2) and (7.3), with* f *and* g *continuously differentiable. Suppose that this problem has an interior continuous solution* $\hat{y}(t) \in Int\mathcal{Y}(t)$ *and* $\hat{x}(t) \in Int\mathcal{X}(t)$*. Then, there exists a continuously differentiable function* $\lambda(t)$ *such that the optimal control* $\hat{y}(t)$ *and the corresponding path of the state variable* $\hat{x}(t)$ *satisfy the following necessary conditions:* $x(0) = x_0$*,*

$$H_y(t, \hat{x}(t), \hat{y}(t), \lambda(t)) = 0 \text{ for all } t \in [0, t_1], \tag{7.13}$$

$$\dot{\lambda}(t) = -H_x(t, \hat{x}(t), \hat{y}(t), \lambda(t)) \text{ for all } t \in [0, t_1], \tag{7.14}$$

$$\dot{x}(t) = H_\lambda(t, \hat{x}(t), \hat{y}(t), \lambda(t)) \text{ for all } t \in [0, t_1], \tag{7.15}$$

and $\lambda(t_1) = 0$*, with the Hamiltonian* $H(t, x, y, \lambda)$ *given by (7.12). Moreover, the Hamiltonian* $H(t, x, y, \lambda)$ *also satisfies the* Maximum Principle *that*

$$H(t, \hat{x}(t), \hat{y}(t), \lambda(t)) \geq H(t, \hat{x}(t), y, \lambda(t)) \text{ for all } y \in \mathcal{Y}(t),$$

for all $t \in [0, t_1]$*.*

For notational simplicity, in eq. (7.15), I wrote $\dot{x}(t)$ instead of $\dot{\hat{x}}(t)$ $(= \partial\hat{x}(t)/\partial t)$. The latter notation is rather cumbersome, and I refrain from using it as long as the context makes it clear that $\dot{x}(t)$ stands for this expression.

[4]More generally, the Hamiltonian should be written as

$$H(t, x, y, \lambda) \equiv \lambda_0 f(t, x(t), y(t)) + \lambda(t) g(t, x(t), y(t)).$$

for some $\lambda_0 \geq 0$. In some pathological cases λ_0 may be equal to 0. However, in all economic applications this will not be the case, and we will have $\lambda_0 > 0$. When $\lambda_0 > 0$, it can be normalized to 1 without loss of any generality. Thus the definition of the Hamiltonian in (7.12) is appropriate for economic applications.

Theorem 7.4 is a simplified version of the celebrated *Maximum Principle* of Pontryagin. The more general version of this Maximum Principle will be given below. For now, a couple of features are worth noting:

(1) As in the usual constrained maximization problems, the optimal solution is characterized jointly with a set of "multipliers," here the *costate* variable $\lambda(t)$, and the optimal path of the control and state variables, $\hat{y}(t)$ and $\hat{x}(t)$.
(2) Again as with the Lagrange multipliers in the usual constrained maximization problems, the costate variable $\lambda(t)$ is informative about the value of relaxing the constraint (at time t). In particular, $\lambda(t)$ is the value of an infinitesimal increase in $x(t)$ at time t (see below).
(3) With this interpretation, it makes sense that $\lambda(t_1) = 0$ is part of the necessary conditions. After the planning horizon, there is no value to having more x. This is therefore the finite-horizon equivalent of the *transversality condition* in the previous chapter.

As emphasized above, Theorem 7.4 gives necessary conditions for an interior continuous solution. However, we do not know that such a solution exists. Moreover, even when such a solution exists, it may correspond to a "stationary point" rather than a maximum, or it may be a local rather than a global maximum of the optimization problem. Therefore, a sufficiency result is even more important in this context than in finite-dimensional optimization problems. Sufficiency is again guaranteed by imposing concavity. The following theorem, first proved by Mangasarian, shows that concavity of the Hamiltonian ensures that conditions (7.13)-(7.15) are not only necessary but also sufficient for a maximum.

Theorem 7.5. ***(Mangasarian's Sufficiency Conditions)*** *Consider the problem of maximizing (7.1) subject to (7.2) and (7.3), with f and g continuously differentiable. Define $H(t,x,y,\lambda)$ as in (7.12), and suppose that an interior continuous solution $\hat{y}(t) \in Int\mathcal{Y}(t)$ and $\hat{x}(t) \in Int\mathcal{X}(t)$ exists and satisfies (7.13)-(7.15). Suppose also that given the resulting costate variable $\lambda(t)$, $H(t,x,y,\lambda)$ is jointly concave in (x,y) for all $t \in [0,t_1]$, then the $\hat{y}(t)$ and the corresponding $\hat{x}(t)$ achieve a global maximum of (7.1). Moreover, if $H(t,x,y,\lambda)$ is strictly jointly concave in (x,y) for all $t \in [0,t_1]$, then the pair $(\hat{x}(t),\hat{y}(t))$ is the unique solution to (7.1).*

The proof of Theorem 7.5 follows from the proof of Theorem 7.6 and is discussed after the next theorem.[5] The next result, Theorem 7.6, which was first derived by Arrow, weakens the condition that $H(t,x,y,\lambda)$ is jointly concave in (x,y). Before stating this result, let us

[5]The statements that $(\hat{x}(t),\hat{y}(t))$ is unique in the last part of the theorem requires some comment. When the problem is formulated in its most general form, $\hat{y}(t)$ needs to be Lebesgue measurable, and thus $(\hat{x}(t),\hat{y}(t))$ will be unique with the qualification "almost everywhere," though $\hat{x}(t)$ would continue to be uniquely defined. When $\hat{y}(t)$ is piecewise continuous, there is no need for this qualification. I ignore this qualification throughout.

define the *maximized Hamiltonian* as

$$M\left(t,x\left(t\right),\lambda\left(t\right)\right)\equiv\max_{y\in\mathcal{Y}(t)}H\left(t,x\left(t\right),y,\lambda\left(t\right)\right), \tag{7.16}$$

with $H\left(t,x\left(t\right),y\left(t\right),\lambda\left(t\right)\right)$ itself defined as in (7.12). Clearly, the necessary conditions for an interior maximum in (7.16) is (7.13). Therefore, an interior pair of state and control variables $\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)$ satisfies (7.13)-(7.15), then $M\left(t,\hat{x}\left(t\right),\lambda\left(t\right)\right)\equiv H\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right),\lambda\left(t\right)\right)$.

THEOREM 7.6. ***(Arrow's Sufficiency Conditions)*** *Consider the problem of maximizing (7.1) subject to (7.2) and (7.3), with f and g continuously differentiable. Define $H\left(t,x,y,\lambda\right)$ as in (7.12), and suppose that an interior continuous solution $\hat{y}\left(t\right)\in Int\mathcal{Y}\left(t\right)$ and $\hat{x}\left(t\right)\in Int\mathcal{X}\left(t\right)$ exists and satisfies (7.13)-(7.15). Given the resulting costate variable $\lambda\left(t\right)$, define the maximized Hamiltonian $M\left(t,\hat{x},\lambda\right)$ as in (7.16). If $\mathcal{X}\left(t\right)$ is a convex set for all t and $M\left(t,\hat{x},\lambda\right)$ is concave in $x\in\mathcal{X}\left(t\right)$ for all $t\in\left[0,t_1\right]$, then $\hat{y}\left(t\right)$ and the corresponding $\hat{x}\left(t\right)$ achieve a global maximum of (7.1). Moreover, if $M\left(t,\hat{x},\lambda\right)$ is strictly concave in x for all $t\in\left[0,t_1\right]$, then the pair $\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)$ achieves the unique global maximum of (7.1) and $\hat{x}\left(t\right)$ is uniquely defined.*

PROOF. Consider an admissible pair $\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)$ satisfying the necessary conditions (7.13)-(7.15) as well as (7.2) and (7.3). Consider also an arbitrary pair $\left(x\left(t\right),y\left(t\right)\right)$ that satisfies (7.2) and (7.3) and define $M\left(t,x,\lambda\right)\equiv\max_y H\left(t,x,y,\lambda\right)$. Since f and g are differentiable, $H\left(t,x\left(t\right),y\left(t\right),\lambda\left(t\right)\right)$ and $M\left(t,x\left(t\right),\lambda\left(t\right)\right)$ are also differentiable in $x\left(t\right)$ at time t. Denote the derivative of M with respect to x by M_x. Since M is concave, we have (recall, for example, Corollary A.4 in Appendix Chapter A):

$$M\left(t,x\left(t\right),\lambda\left(t\right)\right)\leq M\left(t,\hat{x}\left(t\right),\lambda\left(t\right)\right)+M_x\left(t,\hat{x}\left(t\right),\lambda\left(t\right)\right)\left(x\left(t\right)-\hat{x}\left(t\right)\right)\text{ for all }t\in\left[0,t_1\right].$$

Integrating both sides over $\left[0,t_1\right]$ yields

$$\int_0^{t_1}M\left(t,x\left(t\right),\lambda\left(t\right)\right)dt\leq\int_0^{t_1}M\left(t,\hat{x}\left(t\right),\lambda\left(t\right)\right)dt+\int_0^{t_1}M_x\left(t,\hat{x}\left(t\right),\lambda\left(t\right)\right)\left(x\left(t\right)-\hat{x}\left(t\right)\right)dt. \tag{7.17}$$

Moreover,

$$\begin{aligned}M_x\left(t,\hat{x}\left(t\right),\lambda\left(t\right)\right) &= H_x\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right),\lambda\left(t\right)\right)\\ &= -\dot{\lambda}\left(t\right),\end{aligned} \tag{7.18}$$

where the first line follows by an Envelope Theorem type reasoning (since $H_y=0$ from eq. (7.13)), while the second line follows from (7.14). Next, exploiting the definition of the maximized Hamiltonian,

$$\int_0^{t_1}M\left(t,x\left(t\right),\lambda\left(t\right)\right)dt=W\left(x\left(t\right),y\left(t\right)\right)+\int_0^{t_1}\lambda\left(t\right)g\left(t,x\left(t\right),y\left(t\right)\right)dt,$$

and

$$\int_0^{t_1}M\left(t,\hat{x}\left(t\right),\lambda\left(t\right)\right)dt=W\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)+\int_0^{t_1}\lambda\left(t\right)g\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right)\right)dt.$$

Equation (7.17) together with (7.18) then implies

$$\begin{aligned} W\left(x\left(t\right),y\left(t\right)\right) &\leq W\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right) \\ &\quad+\int_{0}^{t_{1}}\lambda\left(t\right)\left[g\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right)\right)-g\left(t,x\left(t\right),y\left(t\right)\right)\right]dt \\ &\quad-\int_{0}^{t_{1}}\dot{\lambda}\left(t\right)\left(x\left(t\right)-\hat{x}\left(t\right)\right)dt. \end{aligned} \tag{7.19}$$

Integrating the last term by parts (Theorem B.3) and using the fact that by feasibility $x\left(0\right)=\hat{x}\left(0\right)=x_{0}$ and by the transversality condition $\lambda\left(t_{1}\right)=0$,

$$\int_{0}^{t_{1}}\dot{\lambda}\left(t\right)\left(x\left(t\right)-\hat{x}\left(t\right)\right)dt=-\int_{0}^{t_{1}}\lambda\left(t\right)\left(\dot{x}\left(t\right)-\dot{\hat{x}}\left(t\right)\right)dt.$$

Substituting this into (7.19),

$$\begin{aligned} W\left(x\left(t\right),y\left(t\right)\right) &\leq W\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right) \\ &\quad+\int_{0}^{t_{1}}\lambda\left(t\right)\left[g\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right)\right)-g\left(t,x\left(t\right),y\left(t\right)\right)\right]dt \\ &\quad+\int_{0}^{t_{1}}\lambda\left(t\right)\left[\dot{x}\left(t\right)-\dot{\hat{x}}\left(t\right)\right]dt. \end{aligned} \tag{7.20}$$

By definition of the admissible pairs $\left(x\left(t\right),y\left(t\right)\right)$ and $\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)$, $\dot{\hat{x}}\left(t\right)=g\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right)\right)$ and $\dot{x}\left(t\right)=g\left(t,x\left(t\right),y\left(t\right)\right)$. Thus (7.20) implies that $W\left(x\left(t\right),y\left(t\right)\right)\leq W\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)$ for any admissible pair $\left(x\left(t\right),y\left(t\right)\right)$, establishing the first part of the theorem.

If M is strictly concave in x, then the inequality in (7.17) is strict, and therefore the same argument establishes $W\left(x\left(t\right),y\left(t\right)\right)<W\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)$, and no other admissible $\left(x\left(t\right),y\left(t\right)\right)$ could achieve the same value, establishing the second part. □

Given Theorem 7.6, the proof of Theorem 7.5 follows as a direct corollary, since if a function $H\left(t,x,y,\lambda\right)$ is jointly [strictly] concave in $\left(x,y\right)$, then $M\left(t,x,\lambda\right)\equiv\max_{y}H\left(t,x,y,\lambda\right)$ is [strictly] concave in x (see Exercise 7.7). Nevertheless, in many applications it may be easier to verify that $H\left(t,x,y,\lambda\right)$ is jointly concave in $\left(x,y\right)$ rather than looking at the maximized Hamiltonian. Moreover, in some problems $M\left(t,x,\lambda\right)$ may be concave in x, while $H\left(t,x,y,\lambda\right)$ is concave in $\left(x,y\right)$ and strictly concave in y, and this information may be useful in establishing uniqueness. However, to economize on space, I focus on sufficiency theorems along the lines of Theorem 7.6.

Sufficiency results as in Theorems 7.5 and 7.6 play an important role in the applications of optimal control. They ensure that an admissible pair $\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)$ that satisfies the necessary conditions specified in Theorem 7.4 and the sufficiency conditions in either Theorem 7.5 or Theorem 7.6 is indeed an optimal solution. This is important, since without the sufficiency results, Theorem 7.4 does not tell us that there exists an interior continuous solution, thus an admissible pair that satisfies the conditions of Theorem 7.4 may not be optimal or an optimal solution may not satisfy these "necessary conditions" (because it is not interior or continuous). Sufficiency results circumvent these problems by establishing that a candidate admissible pair

(satisfying the necessary conditions for an interior continuous solution) corresponds to a global maximum or a unique global maximum.

One difficulty in verifying that the conditions in Theorem 7.5 and Theorem 7.6 are satisfied is that neither the concavity nor convexity of the $g(\cdot)$ function would guarantee the concavity of the Hamiltonian unless we know something about the sign of the costate variable $\lambda(t)$. Nevertheless, in many economically interesting situations, we can ascertain that the costate variable $\lambda(t)$ is everywhere positive. For example, $f_y(t,\hat{x}(t),\hat{y}(t),\lambda(t)) \geq 0$ and $g_y(t,\hat{x}(t),\hat{y}(t),\lambda(t)) \leq 0$ will be sufficient to ensure that $\lambda(t) \geq 0$. Once we know that $\lambda(t)$ is nonnegative, checking the sufficiency conditions is straightforward, especially when f and g are concave functions.

7.2.2. Generalizations. The above theorems can be generalized to the case in which the state variable and the controls are vectors and also to the case in which there are other constraints. The constrained case requires *constraint qualification conditions* as in the standard finite-dimensional optimization case (see, e.g., Theorems A.30 and A.31 in Appendix Chapter A). These are slightly more messy to express. Since I make no use of the constrained maximization problems (except in Exercise 10.7 in Chapter 10), I only discuss constrained problems briefly in Exercise 7.10.

The vector-valued theorems are direct generalizations of the ones presented above and are useful in growth models with multiple capital goods. In particular, let

$$\max_{\mathbf{x}(t),\mathbf{y}(t),\mathbf{x}_1} W(\mathbf{x}(t),\mathbf{y}(t)) \equiv \int_0^{t_1} f(t,\mathbf{x}(t),\mathbf{y}(t))\,dt \tag{7.21}$$

subject to

$$\dot{\mathbf{x}}(t) = G(t,\mathbf{x}(t),\mathbf{y}(t)), \tag{7.22}$$

and

$$\mathbf{y}(t) \in \mathcal{Y}(t),\ \mathbf{x}(t) \in \mathcal{X}(t) \text{ for all } t,\ \mathbf{x}(0) = \mathbf{x}_0 \text{ and } \mathbf{x}(t_1) = \mathbf{x}_1, \tag{7.23}$$

where $\mathcal{Y}(t) \subset \mathbb{R}^{K_y}$ and $\mathcal{X}(t) \subset \mathbb{R}^{K_x}$ (with K_y,$K_x \in \mathbb{N}$).

THEOREM 7.7. ***(Maximum Principle for Multivariate Problems)*** *Consider the problem of maximizing (7.21) subject to (7.22) and (7.23), with f and G continuously differentiable, has an interior continuous solution $\hat{\mathbf{y}}(t) \in Int\mathcal{Y}(t)$ and $\hat{\mathbf{x}}(t) \in Int\mathcal{X}(t)$. Let $H(t,\mathbf{x},\mathbf{y},\boldsymbol{\lambda})$ be given by*

$$H(t,\mathbf{x},\mathbf{y},\boldsymbol{\lambda}) \equiv f(t,\mathbf{x}(t),\mathbf{y}(t)) + \boldsymbol{\lambda}(t)\cdot G(t,\mathbf{x}(t),\mathbf{y}(t)), \tag{7.24}$$

where $\boldsymbol{\lambda}(t) \in \mathbb{R}^{K_x}$ (and $\boldsymbol{\lambda}\cdot G$ denotes the inner product of the vectors $\boldsymbol{\lambda}$ and G). Then, the optimal control $\hat{\mathbf{y}}(t)$ and the corresponding path of the state variable $\mathbf{x}(t)$ satisfy the following necessary conditions:

$$D_{\mathbf{y}}H(t,\hat{\mathbf{x}}(t),\hat{\mathbf{y}}(t),\boldsymbol{\lambda}(t)) = 0 \text{ for all } t \in [0,t_1]. \tag{7.25}$$

$$\dot{\boldsymbol{\lambda}}(t) = -D_{\mathbf{x}}H(t,\hat{\mathbf{x}}(t),\hat{\mathbf{y}}(t),\boldsymbol{\lambda}(t)) \text{ for all } t \in [0,t_1]. \tag{7.26}$$

$$\dot{\mathbf{x}}(t) = D_{\boldsymbol{\lambda}} H\left(t, \hat{\mathbf{x}}(t), \hat{\mathbf{y}}(t), \boldsymbol{\lambda}(t)\right) \text{ for all } t \in [0, t_1], \ \mathbf{x}(0) = \mathbf{x}_0 \text{ and } \boldsymbol{\lambda}(t_1) = 0. \tag{7.27}$$

PROOF. See Exercise 7.11. □

Note also that various conditions in this theorem can be relaxed. For example, the conditions that $\hat{\mathbf{y}}(t) \in \text{Int}\mathcal{Y}(t)$ and $\hat{\mathbf{x}}(t) \in \text{Int}\mathcal{X}(t)$ are not necessary, and when either the state or the control variables take boundary values, there may be jumps in the control variables and the Hamiltonian may not be differentiable everywhere (see below). However, recall that all we require is a piecewise continuous (vector-valued) function $\hat{\mathbf{y}}(t)$, so that this does not create major problems. Since in most economic applications both state and control variables will be interior and the corresponding Hamiltonian will be differentiable everywhere, the form of Theorem 7.7 stated here is sufficient for most problems of interest.

Moreover, the sufficiency conditions for a maximum provided above also have straightforward generalizations, which are presented next.

THEOREM 7.8. ***(Sufficiency Conditions For Multivariate Problems)*** *Consider the problem of maximizing (7.21) subject to (7.22) and (7.23), with f and G continuously differentiable. Define $H(t, \mathbf{x}, \mathbf{y}, \boldsymbol{\lambda})$ as in (7.24), and suppose that an interior continuous solution, $\hat{\mathbf{y}}(t) \in Int\mathcal{Y}(t)$ and $\hat{\mathbf{x}}(t) \in Int\mathcal{X}(t)$, satisfies (7.25)-(7.27). If $\mathcal{X}(t)$ is convex for all t and $M(t, \mathbf{x}, \boldsymbol{\lambda}) \equiv \max_{\mathbf{y}(t) \in \mathcal{Y}(t)} H(t, \mathbf{x}, \mathbf{y}, \boldsymbol{\lambda})$ is concave in $\mathbf{x} \in \mathcal{X}(t)$ for all $t \in [0, t_1]$, then $\hat{\mathbf{y}}(t)$ and the corresponding $\hat{\mathbf{x}}(t)$ achieve a global maximum of (7.21). Moreover, if $M(t, \mathbf{x}, \boldsymbol{\lambda})$ is strictly concave in $\mathbf{x}$, then the pair $(\hat{\mathbf{x}}(t), \hat{\mathbf{y}}(t))$ is the unique solution to (7.21).*

PROOF. See Exercise 7.12. □

7.2.3. Limitations. The limitations of the results presented so far are obvious. First, a continuous and interior solution to the optimal control problem has been assumed. Second, and equally important, the analysis has focused on the finite-horizon case whereas analysis of growth models requires us to solve infinite-horizon problems. To deal with both of these issues, we need to look at the more modern theory of optimal control. This is done in the next section.

7.3. Infinite-Horizon Optimal Control

The results presented so far are most useful in developing an intuition for how dynamic optimization in continuous time works. While a number of problems in economics require finite-horizon optimal control, most economic problems—including almost all growth models—are more naturally formulated as infinite-horizon problems. This is obvious in the context of economic growth, but is also the case in repeated games, political economy or industrial organization, where even if individuals may have finite expected lives, the end date of the game or of their lives may be uncertain. For this reason, the canonical model of optimization in economic problems is the infinite-horizon one. In this section, I provide necessary and sufficient conditions for optimality in infinite-horizon optimal control problems. Since these are the results that are most often used in economic applications, I simplify the exposition

and state these results for the case in which both the state variable and the control variable are one dimensional. The more general, multivariate case is discussed in Section 7.6, when I return to the issues of existence of solutions and properties of the value functions.

7.3.1. The Basic Problem: Necessary and Sufficient Conditions. Let us focus on infinite-horizon control with a single control and a single state variable. For reasons that will be explained below, it is useful to generalize the terminal value constraint on the state variable. For this purpose, throughout this chapter, let $b : \mathbb{R}_+ \to \mathbb{R}_+$, such that $\lim_{t\to\infty} b(t)$ exists and satisfies $\lim_{t\to\infty} b(t) < \infty$. Then, the terminal value condition on the state variable will be specified as $\lim_{t\to\infty} b(t) x(t) \geq x_1$ for some $x_1 \in \mathbb{R}$. The special case where $b(t) = 1$ gives us the terminal value constraint as $\lim_{t\to\infty} x(t) \geq x_1$ and is sufficient in many applications. But for the analysis of competitive equilibrium in continuous time, we will need a terminal value constraint of the form $\lim_{t\to\infty} b(t) x(t) \geq x_1$. Now using the same notation as above, the infinite-horizon optimal control problem is

$$\max_{x(t),y(t)} W(x(t), y(t)) \equiv \int_0^\infty f(t, x(t), y(t))\, dt \tag{7.28}$$

subject to

$$\dot{x}(t) = g(t, x(t), y(t)), \tag{7.29}$$

and

$$y(t) \in \mathcal{Y}(t),\ x(t) \in \mathcal{X}(t) \text{ for all } t,\ x(0) = x_0 \text{ and } \lim_{t\to\infty} b(t) x(t) \geq x_1, \tag{7.30}$$

with again $f : \mathbb{R}\times\mathbb{R}\times\mathbb{R} \to \mathbb{R}$ and $g : \mathbb{R}\times\mathbb{R}\times\mathbb{R} \to \mathbb{R}$. Recall that in (7.30), $b : \mathbb{R}_+ \to \mathbb{R}_+$ with $\lim_{t\to\infty} b(t) < \infty$.

The main difference is that now time runs to infinity. Note also that this problem allows for an implicit choice over the endpoint x_1, since there is no terminal date. The last part of (7.30) imposes a lower bound on this endpoint. In addition, I have further simplified the problem by removing the feasibility requirement that the control $y(t)$ should always belong to the set $\mathcal{Y}(t)$ and that $x(t)$ should belong to the set $\mathcal{X}(t)$, instead simply requiring both to be real-valued functions of time. Because the state variable $x(t)$ does not necessarily lie in a compact set, the results developed here can be applied to endogenous growth models.

The definition of an *admissible pair* $(x(t), y(t))$ is defined in the same way as above (in particular, recall footnote 1). Since $x(t)$ is given by a continuous differential equation, when $y(t)$ is continuous, $x(t)$ will be differentiable and when $y(t)$ is piecewise continuous (meaning that it has a finite number of jumps), $x(t)$ will be differentiable almost everywhere.

There are a number of technical difficulties when dealing with the infinite-horizon case, which are similar to those in the discrete-time analysis. Primary among those is the fact that the value of the functional in (7.28) may not be finite. These issues will be dealt with below.

The main theorem for the infinite-horizon optimal control problem is the following more general version of the *Maximum Principle*. Before stating this theorem, let us recall that the Hamiltonian is defined by (7.12), with the only difference that the horizon is now infinite.

In addition, let us define the *value function*, which is the analog of the value function in discrete-time dynamic programming introduced in the previous chapter:

$$\begin{aligned} V(t_0, x_0) &\equiv \sup_{(x(t),y(t))\in\mathcal{X}(t)\times\mathcal{Y}(t)} \int_{t_0}^{\infty} f(t, x(t), y(t))\, dt \\ \text{subject to } \dot{x}(t) &= g(t, x(t), y(t)),\ x(t_0) = x_0 \text{ and } \lim_{t\to\infty} x(t) \geq x_1. \end{aligned} \tag{7.31}$$

In words, $V(t_0, x_0)$ gives the optimal value of the dynamic maximization problem starting at time t_0 with state variable x_0. Clearly,

$$V(t_0, x_0) \geq \int_{t_0}^{\infty} f(t, x(t), y(t))\, dt \text{ for any admissible pair } (x(t), y(t)). \tag{7.32}$$

Our focus throughout will be on problems where there exist admissible pairs $(\hat{x}(t), \hat{y}(t))$ that reach $V(t_0, x_0) < \infty$ (see Theorem 7.15 below for sufficient conditions). When $(\hat{x}(t), \hat{y}(t))$ is such a pair, then

$$V(t_0, x_0) = \int_{t_0}^{\infty} f(t, (\hat{x}(t), \hat{y}(t)))\, dt. \tag{7.33}$$

Our first result is a weaker version of the *Principle of Optimality*, which we encountered in the context of discrete-time dynamic programming in the previous chapter:

LEMMA 7.1. ***(Principle of Optimality)*** *Suppose that the pair $(\hat{x}(t), \hat{y}(t))$ is an optimal solution to (7.28) subject to (7.29) and (7.30), i.e., it reaches the maximum value $V(t_0, x_0)$. Then,*

$$V(t_0, x_0) = \int_{t_0}^{t_1} f(t, \hat{x}(t), \hat{y}(t))\, dt + V(t_1, \hat{x}(t_1)) \text{ for all } t_1 \geq t_0. \tag{7.34}$$

PROOF. We have

$$\begin{aligned} V(t_0, x_0) &\equiv \int_{t_0}^{\infty} f(t, \hat{x}(t), \hat{y}(t))\, dt \\ &= \int_{t_0}^{t_1} f(t, \hat{x}(t), \hat{y}(t))\, dt + \int_{t_1}^{\infty} f(t, \hat{x}(t), \hat{y}(t))\, dt. \end{aligned}$$

The proof is completed if $V(t_1, \hat{x}(t_1)) = \int_{t_1}^{\infty} f(t, \hat{x}(t), \hat{y}(t))\, dt$. By definition $V(t_1, \hat{x}(t_1)) \geq \int_{t_1}^{\infty} f(t, x(t), y(t))\, dt$ for all admissible $(x(t), y(t))$. Thus this equality can only fail if $V(t_1, \hat{x}(t_1)) > \int_{t_1}^{\infty} f(t, \hat{x}(t), \hat{y}(t))\, dt$. To obtain a contradiction, suppose that this is the case. Then, there must exist an admissible pair from t_1 onwards, $(\tilde{x}(t), \tilde{y}(t))$ with $\tilde{x}(t_1) = \hat{x}(t_1)$ such that $\int_{t_1}^{\infty} f(t, \tilde{x}(t), \tilde{y}(t))\, dt > \int_{t_1}^{\infty} f(t, \hat{x}(t), \hat{y}(t))\, dt$. Then, construct the pair $(\vec{x}(t), \vec{y}(t))$ such that $(\vec{x}(t), \vec{y}(t)) = (\hat{x}(t), \hat{y}(t))$ for all $t \in [t_0, t_1]$ and $(\vec{x}(t), \vec{y}(t)) = (\tilde{x}(t), \tilde{y}(t))$ for all $t \geq t_1$. Since $(\tilde{x}(t), \tilde{y}(t))$ is admissible from t_1 onwards with $\tilde{x}(t_1) = \hat{x}(t_1)$,

$(\vec{x}(t), \vec{y}(t))$ is admissible, and moreover,

$$\begin{aligned}
\int_{t_0}^{\infty} f(t, \vec{x}(t), \vec{y}(t))\, dt &= \int_{t_0}^{t_1} f(t, \vec{x}(t), \vec{y}(t))\, dt + \int_{t_1}^{\infty} f(t, \vec{x}(t), \vec{y}(t))\, dt \\
&= \int_{t_0}^{t_1} f(t, \hat{x}(t), \hat{y}(t))\, dt + \int_{t_1}^{\infty} f(t, \tilde{x}(t), \tilde{y}(t))\, dt \\
&> \int_{t_0}^{t_1} f(t, \hat{x}(t), \hat{y}(t))\, dt + \int_{t_1}^{\infty} f(t, \hat{x}(t), \hat{y}(t))\, dt \\
&= \int_{t_0}^{\infty} f(t, \hat{x}(t), \hat{y}(t))\, dt \\
&= V(t_0, x_0),
\end{aligned}$$

which contradicts (7.32) establishing that $V(t_1, \hat{x}(t_1)) = \int_{t_1}^{\infty} f(t, \hat{x}(t), \hat{y}(t))\, dt$ and thus (7.34). □

Two features in this version of the Principle of Optimality are noteworthy. First, in contrast to the similar equation in the previous chapter, it may appear that there is no discounting in (7.34). This is not the case, since the discounting is embedded in the instantaneous payoff function f, and is thus implicit in $V(t_1, \hat{x}(t_1))$. Second, this lemma may appear to contradict the discussion of "time-consistency" in the previous chapter, since the lemma is stated without additional assumptions that ensure time-consistency. The important point here is that in the time-consistency discussion, the decision-maker considered updating his or her plan, with the payoff function being potentially different after date t_1 (at least because bygones were bygones). In contrast, here the payoff function remains constant. The issue of time-consistency is discussed further in Exercise 7.22. I next state one of the main results on necessary conditions. In this theorem, I also slightly relax the assumption that the optimal control $\hat{y}(t)$ is continuous.

THEOREM 7.9. ***(Infinite-Horizon Maximum Principle)*** *Suppose that problem of maximizing (7.28) subject to (7.29) and (7.30), with f and g continuously differentiable, has a piecewise continuous interior solution, $\hat{y}(t) \in Int\mathcal{Y}(t)$ with a corresponding state variable $\hat{x}(t) \in Int\mathcal{X}(t)$. Let $H(t, x, y, \lambda)$ be defined in (7.12). Then, the optimal control $\hat{y}(t)$ and the corresponding path of the state variable $\hat{x}(t)$ are such that the Hamiltonian $H(t, x, y, \lambda)$ satisfies the Maximum Principle, that*

$$H(t, \hat{x}(t), \hat{y}(t), \lambda(t)) \geq H(t, \hat{x}(t), y(t), \lambda(t))$$

for all $y(t) \in \mathcal{Y}(t)$ and for all $t \in \mathbb{R}$. Moreover, whenever $\hat{y}(t)$ is continuous, the following necessary conditions are satisfied:

$$H_y(t, \hat{x}(t), \hat{y}(t), \lambda(t)) = 0, \tag{7.35}$$

$$\dot{\lambda}(t) = -H_x(t, \hat{x}(t), \hat{y}(t), \lambda(t)), \tag{7.36}$$

$$\dot{x}(t) = H_\lambda(t, \hat{x}(t), \hat{y}(t), \lambda(t)), \text{ with } x(0) = x_0 \text{ and } \lim_{t\to\infty} b(t)\, x(t) \geq x_1, \tag{7.37}$$

for all $t \in \mathbb{R}_+$.

The proof of this theorem is relatively long and will be provided later in this section. Notice that whenever an optimal solution of the specified form exists, it satisfies the Maximum Principle. Thus in some ways Theorem 7.9 can be viewed as stronger than the theorems presented in the previous chapter, especially since it does not impose compactness type conditions. Nevertheless, this theorem only applies when the maximization problem has a piecewise continuous solution $\hat{y}(t)$.[6] In addition, Theorem 7.9 states that if the optimal control, $\hat{y}(t)$, is a continuous function of time, conditions (7.35)-(7.37) are satisfied everywhere. The fact that $\hat{y}(t)$ is a piecewise continuous function implies that the optimal control may include discontinuities, but these will be relatively "rare"—in particular, it will be continuous "most of the time". The added generality of allowing discontinuities is somewhat superfluous in most economic applications, because economic problems often have enough structure to ensure that $\hat{y}(t)$ is indeed a continuous function of time. Consequently, in most economic problems (and in all of the models studied in this book) it will be sufficient to focus on the necessary conditions (7.35)-(7.37).

It is also useful to have a different version of the necessary conditions in Theorem 7.9, which are directly comparable to the necessary conditions generated by dynamic programming in the discrete-time dynamic optimization problems studied in the previous chapter. In particular, the necessary conditions can also be expressed in the form of the so-called *Hamilton-Jacobi-Bellman* (HJB) equation.

THEOREM 7.10. ***(Hamilton-Jacobi-Bellman Equation)*** *Let $V(t,x)$ be as defined in (7.31) and suppose that the hypotheses in Theorem 7.9 hold. Then, whenever $V(t,x)$ is differentiable in (t,x), the optimal pair $(\hat{x}(t),\hat{y}(t))$ satisfies the HJB equation:*

$$f(t,\hat{x}(t),\hat{y}(t))+\frac{\partial V(t,\hat{x}(t))}{\partial t}+\frac{\partial V(t,\hat{x}(t))}{\partial x}g(t,\hat{x}(t),\hat{y}(t))=0 \text{ for all } t\in\mathbb{R}. \tag{7.38}$$

PROOF. From Lemma 7.1, for the optimal pair $(\hat{x}(t),\hat{y}(t))$,

$$V(t_0,x_0)=\int_{t_0}^{t} f(s,\hat{x}(s),\hat{y}(s))\,ds+V(t,\hat{x}(t)) \text{ for all } t.$$

Differentiating this with respect to t and using the differentiability of V and Leibniz's rule (Theorem B.4 in Appendix Chapter B),

$$f(t,\hat{x}(t),\hat{y}(t))+\frac{\partial V(t,\hat{x}(t))}{\partial t}+\frac{\partial V(t,\hat{x}(t))}{\partial x}\dot{x}(t)=0 \text{ for all } t.$$

Setting $\dot{x}(t)=g(t,\hat{x}(t),\hat{y}(t))$ gives (7.38). □

The HJB equation will be useful in providing an intuition for the Maximum Principle. More importantly, it is used directly in many economic models, including in the endogenous technology models studied in Part 4 below.

[6]The reader may also wonder when an optimal piecewise continuous solution will exist as hypothesized in the theorem. Theorem 7.15 below will provide the conditions to ensure that a solution exists. Ensuring the existence of a solution that is piecewise continuous is much harder. Nevertheless, in most economic problems there will be enough structure to ensure the existence of an interior solution and this structure will also often guarantee that the solution is also piecewise continuous and in fact fully continuous.

It is also worth noting a few important features concerning the HJB equation. First, given that the continuous differentiability of f and g, the assumption that $V(t,x)$ is differentiable is not very restrictive, though it is not always satisfied. From the definition (7.31), at all t where $\hat{y}(t)$ is continuous and $g(t,x,y)$ is differentiable in t, $V(t,x)$ will also be differentiable in t. Moreover, an Envelope Theorem type argument also implies that when $\hat{y}(t)$ is continuous, $V(t,x)$ should also be differentiable in x (the differentiability of $V(t,x)$ in x can also be established directly, see Theorem 7.17 below). Second, (7.38) is a partial differential equation, since it features the derivative of V with respect to both time and the state variable x. Third, this partial differential equation also has a similarity to the Euler equation derived in the context of discrete-time dynamic programming. In particular, the simplest Euler equation, (6.28), in the previous chapter required the current gain from increasing the control variable to be equal to the discounted loss of value. The current equation has a similar interpretation, with the first term corresponding to the current gain and the last term to the potential discounted loss of value. The second term results from the fact that the maximized value can also change over time.

7.3.2. Heuristic Derivation of the Stationary HJB Equation. Given its prominent role in dynamic economic analysis, it is useful to provide an alternative heuristic (and more "intuitive") proof of the HJB equation. For this purpose, let us focus on the simpler stationary version of the HJB equation. This stationary version applies to exponentially discounted maximization problems with time-autonomous constraints (see Sections 7.5 and 7.7 below). Briefly, in these problems the payoff function is exponentially discounted, that is,

$$f(t,x(t),y(t)) = \exp(-\rho t) f(x(t),y(t)),$$

and the law of motion of the state variable is given by an autonomous differential equation, that is,

$$g(t,x(t),y(t)) = g(x(t),y(t)).$$

In this case, one can easily verify that if an admissible pair $(\hat{x}(t),\hat{y}(t))_{t\geq 0}$ is optimal starting at $t=0$ with initial condition $x(0)=x_0$, then it is also optimal starting at $s>0$, starting with the same initial condition, that is, $(\hat{x}(t),\hat{y}(t))_{t\geq s}$ is optimal for the problem with initial condition $x(s)=x_0$ (see Exercise 7.16). In view of this, let us define $V(x)\equiv V(0,x)$, that is, the value of pursuing the optimal plan $(\hat{x}(t),\hat{y}(t))$ starting with initial condition x, evaluated at $t=0$. Since $(\hat{x}(t),\hat{y}(t))$ is an optimal plan regardless of the starting date,

$$V(t,x(t)) \equiv \exp(-\rho t) V(x(t)). \tag{7.39}$$

Then, by definition,

$$\frac{\partial V(t,x(t))}{\partial t} = -\exp(-\rho t)\rho V(x(t)).$$

Moreover, let $\dot{V}(x(t)) \equiv (\partial V(t,x(t))/\partial x)\dot{x}(t)$ be change in the function V over time, which results from the change in the state variable $x(t)$ over time. Substituting these expressions

into (7.38) and noting that $\dot{x}(t) = g(\hat{x}(t), \hat{y}(t))$, we obtain the "stationary" form of the Hamilton-Jacobi-Bellman equation as:

$$\rho V(\hat{x}(t)) = f(\hat{x}(t), \hat{y}(t)) + \dot{V}(\hat{x}(t)). \tag{7.40}$$

This stationary HJB equation is widely used in dynamic economic analysis and can be interpreted as a "no-arbitrage asset value equation," as discussed in subsection 7.3.4 below. The following heuristic argument not only shows how this equation is derived, but also provides a further intuition.

Heuristic Derivation of the Stationary HJB Equation (7.40): Consider the discounted infinite-horizon problem described above and suppose that the admissible pair $(\hat{x}(t), \hat{y}(t))$ is optimal starting at $t = 0$ with initial condition $x(0)$. Recall that the value function starting at $x(0)$ is defined as

$$V(x(0)) \equiv \int_0^{\infty} \exp(-\rho t) f((\hat{x}(t), \hat{y}(t))) dt,$$

with $\hat{x}(0) = x(0)$. Now for a time interval $\Delta t > 0$, the integral can be separated into two parts and written as

$$\begin{aligned} V(x(0)) &= \int_0^{\Delta t} \exp(-\rho t) f((\hat{x}(t), \hat{y}(t))) dt + \int_{\Delta t}^{\infty} \exp(-\rho t) f((\hat{x}(t), \hat{y}(t))) dt, \\ &= f((\hat{x}(0), \hat{y}(0))) \Delta t + o(\Delta t) + \int_{\Delta t}^{\infty} \exp(-\rho t) f((\hat{x}(t), \hat{y}(t))) dt, \end{aligned}$$

where the second line approximates $\int_0^{\Delta t} \exp(-\rho t) f((\hat{x}(t), \hat{y}(t))) dt$ with $f((\hat{x}(0), \hat{y}(0))) \Delta t$, and $o(\Delta t)$ is the residual from this approximation (which results from ignoring both the fact that there is discounting between 0 and Δt and that $(\hat{x}(t), \hat{y}(t))$ is not necessarily equal to $(\hat{x}(0), \hat{y}(0))$ for all $t \in [0, \Delta t]$). For Δt small, the $o(\Delta t)$ term will be very small. More formally, it is a "second-order" term satisfying $\lim_{\Delta t \to 0} o(\Delta t)/\Delta t = 0$, so that as $\Delta t \to 0$, this term vanishes faster than the rate at which Δt approaches 0. Next, using the definition (7.39),

$$V(x(0)) = f((\hat{x}(0), \hat{y}(0))) \Delta t + o(\Delta t) + \exp(-\rho \Delta t) V(\hat{x}(\Delta t)).$$

Subtracting $V(\hat{x}(\Delta t))$ from both sides and dividing both sides by Δt,

$$\frac{V(\hat{x}(0)) - V(\hat{x}(\Delta t))}{\Delta t} = f((\hat{x}(0), \hat{y}(0))) + \frac{o(\Delta t)}{\Delta t} + \frac{\exp(-\rho \Delta t) - 1}{\Delta t} V(\hat{x}(\Delta t)). \tag{7.41}$$

Next take the limit as $\Delta t \to 0$. Under the assumption that $V(x)$ is differentiable at $x(0)$, the left-hand side becomes $-\dot{V}(\hat{x}(0))$. The first term on the right-hand side does not depend on Δt, so it remains unaltered. The second term vanishes by definition. The limit of the third term can be written as

$$\lim_{\Delta t \to 0} \frac{\exp(-\rho \Delta t) - 1}{\Delta t} \times \lim_{\Delta t \to 0} V(\hat{x}(\Delta t)) = -\rho V(\hat{x}(0)).$$

This follows because $\exp(-\rho \times 0) = 1$ and thus as $\Delta t \to 0$, the expression $(\exp(-\rho \Delta t) - 1)/\Delta t = (\exp(-\rho \Delta t) - \exp(-\rho \times 0))/\Delta t$ is simply the derivative of

$\exp(-\rho t)$ at $t = 0$. Substituting these terms into (7.41),

$$-\dot{V}(\hat{x}(0)) = f(\hat{x}(0), \hat{y}(0)) - \rho V(\hat{x}(0)).$$

Rearranging, this is identical to (7.40). □

7.3.3. Transversality Condition and Sufficiency. Since there is no terminal value constraint of the form $x(t_1) = x_1$ in Theorem 7.9, we may therefore may expect that there should be a transversality condition similar to the condition that $\lambda(t_1) = 0$ in Theorem 7.1. For example, we may be tempted to impose a transversality condition of the form

$$\lim_{t\to\infty} \lambda(t) = 0, \tag{7.42}$$

which would be generalizing the condition that $\lambda(t_1) = 0$ in Theorem 7.1. But this is not in general the case. A milder transversality condition of the form

$$\lim_{t\to\infty} H(t, \hat{x}(t), \hat{y}(t), \lambda(t)) = 0 \tag{7.43}$$

for an admissible pair $(\hat{x}(t), \hat{y}(t))$ satisfying the necessary conditions applies under relatively mild conditions (see Theorem 7.12 below), but is not easy to check. Stronger transversality conditions apply when we put more structure on the problem (see Section 7.4 below). Before presenting these results, there are immediate generalizations of the sufficiency theorems to this case.

Theorem 7.11. ***(Sufficiency Conditions for Infinite-Horizon Optimal Control)*** *Consider the problem of maximizing (7.28) subject to (7.29) and (7.30), with f and g continuously differentiable. Define $H(t, x, y, \lambda)$ as in (7.12), and suppose that an admissible pair $(\hat{x}(t), \hat{y}(t))$ satisfies (7.35)-(7.37) except at points of discontinuity of $\hat{y}(t)$. Given the resulting costate variable $\lambda(t)$, define $M(t, x, \lambda) \equiv \max_{y(t)\in\mathcal{Y}(t)} H(t, x, y, \lambda)$. If $\mathcal{X}(t)$ is convex and $M(t, x, \lambda)$ is concave in $x \in \mathcal{X}(t)$ for all t, and $\lim_{t\to\infty} \lambda(t)(\hat{x}(t) - \tilde{x}(t)) \leq 0$ for all $\tilde{x}(t)$ implied by an admissible control path $\tilde{y}(t)$, then the pair $(\hat{x}(t), \hat{y}(t))$ achieves the global maximum of (7.28). Moreover, if $M(t, x, \lambda)$ is strictly concave in x, then $(\hat{x}(t), \hat{y}(t))$ is the unique solution to (7.28).*

Proof. See Exercise 7.13. □

This sufficiency result involves the difficult to check condition $\lim_{t\to\infty} \lambda(t)(x(t) - \tilde{x}(t)) \leq 0$ for all $\tilde{x}(t)$ (that is, for all $\tilde{x}(t)$ implied by an admissible control path $\tilde{y}(t)$). When we impose the appropriate transversality condition, this condition will become more straightforward and easier to verify.

7.3.4. Economic Intuition. The Maximum Principle is not only a powerful mathematical tool, but from an economic point of view, it is *the right tool*, because it captures the essential economic intuition of dynamic economic problems. In this subsection, we provide two different and complementary economic intuitions for the Maximum Principle. One of them is based on the original form as stated in Theorem 7.4 or Theorem 7.9, while the other is based on the dynamic programming (HJB) version provided in Theorem 7.10.

To obtain the first intuition consider the problem of maximizing

$$\int_0^{t_1} H\left(t,\hat{x}\left(t\right),y\left(t\right),\lambda\left(t\right)\right)dt=\int_0^{t_1}\left[f\left(t,\hat{x}\left(t\right),y\left(t\right)\right)+\lambda\left(t\right)g\left(t,\hat{x}\left(t\right),y\left(t\right)\right)\right]dt \tag{7.44}$$

with respect to the entire function $y\left(t\right)$ for given $\lambda\left(t\right)$ and $\hat{x}\left(t\right)$, where t_1 can be finite or equal to $+\infty$. The condition $H_y\left(t,\hat{x}\left(t\right),y\left(t\right),\lambda\left(t\right)\right)=0$ would then be a necessary condition for this alternative maximization problem. Therefore, the Maximum Principle is implicitly maximizing the sum the original maximand $\int_0^{t_1} f\left(t,\hat{x}\left(t\right),y\left(t\right)\right)dt$ plus an additional term $\int_0^{t_1}\lambda\left(t\right)g\left(t,\hat{x}\left(t\right),y\left(t\right)\right)dt$. Understanding why this is true provides much of the intuition for the Maximum Principle.

First recall that $V\left(t,\hat{x}\left(t\right)\right)$ is defined in eq. (7.34) as the value of starting at time t with state variable $\hat{x}\left(t\right)$ and pursuing the optimal policy from then on. The next subsection, in particular eq. (7.48), shows that

$$\lambda\left(t\right)=\frac{\partial V\left(t,\hat{x}\left(t\right)\right)}{\partial x}.$$

Therefore, similar to the Lagrange multipliers in the theory of constraint optimization, $\lambda\left(t\right)$ measures the impact of a small increase in x on the optimal value of the program. Consequently, $\lambda\left(t\right)$ is the (shadow) value of relaxing the constraint (7.29) by increasing the value of $x\left(t\right)$ at time t.[7] Moreover, recall that $\dot{x}\left(t\right)=g\left(t,\hat{x}\left(t\right),y\left(t\right)\right)$, so that the second term in the Hamiltonian is equivalent to $\int_0^{t_1}\lambda\left(t\right)\dot{x}\left(t\right)dt$. This is clearly the shadow value of $x\left(t\right)$ at time t and the increase in the stock of $x\left(t\right)$ at this point. Moreover, recall that $x\left(t\right)$ is the state variable, thus we can think of it as a "stock" variable in contrast to the control $y\left(t\right)$, which corresponds to a "flow" variable.

Therefore, maximizing (7.44) is equivalent to maximizing instantaneous returns as given by the function $f\left(t,\hat{x}\left(t\right),y\left(t\right)\right)$, plus the value of stock of $x\left(t\right)$, as given by $\lambda\left(t\right)$, times the increase in the stock, $\dot{x}\left(t\right)$. This implies that the essence of the Maximum Principle is to maximize the flow return plus the value of the current stock of the state variable. This stock-flow type maximization has a clear economic logic.

Let us next turn to the interpreting the costate equation,

$$\begin{aligned}\dot{\lambda}\left(t\right) &= -H_x\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right),\lambda\left(t\right)\right)\\ &= -f_x\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right)\right)-\lambda\left(t\right)g_x\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right)\right).\end{aligned}$$

This equation is also intuitive. Since $\lambda\left(t\right)$ is the value of the stock of the state variable, $x\left(t\right)$, $\dot{\lambda}\left(t\right)$ is the appreciation in this stock variable. A small increase in x will change the current flow return plus the value of the stock by the amount H_x, but it will also affect the value of the stock by the amount $\dot{\lambda}\left(t\right)$. The Maximum Principle states that this gain should be equal to the depreciation in the value of the stock, $-\dot{\lambda}\left(t\right)$, since, otherwise, it would be possible to change $x\left(t\right)$ and increase the value of (7.44).

[7] Here I am using the language of "relaxing the constraint" implicitly presuming that a high value of $x\left(t\right)$ contributes to increasing the value of the objective function. This simplifies terminology, but it is not necessary for any of the arguments, since $\lambda\left(t\right)$ can be negative.

The second and complementary intuition for the Maximum Principle comes from the HJB equation (7.38) in Theorem 7.10. In particular, let us consider the exponentially discounted problem discussed above (and in greater detail in Section 7.5 below). Recall that in this case, the "stationary" form of the Hamilton-Jacobi-Bellman equation takes

$$\rho V\left(\hat{x}\left(t\right)\right)=f\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)+\dot{V}\left(\hat{x}\left(t\right)\right). \tag{7.45}$$

This widely-used equation can be interpreted as a "no-arbitrage asset value equation". Intuitively, we can think of V as the value of an asset traded in the stock market and ρ as the required rate of return for (a large number of) investors. When will investors be happy to hold this asset? Loosely speaking, they will do so when the asset pays out at least the required rate of return. In contrast, if the asset pays out more than the required rate of return, there would be excess demand for it from the investors until its value adjusts so that its rate of return becomes equal to the "required rate of return". Therefore, we can think of the return on this asset in "equilibrium" being equal to the required rate of return, ρ. The return on the assets come from two sources: first, "dividends," that is, current returns paid out to investors. In the current context, this corresponds to the flow payoff $f\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)$. If this dividend were constant and equal to d, and there were no other returns, then the no-arbitrage condition would imply $V=d/\rho$ or

$$\rho V=d.$$

However, in general the returns to the holding an asset come not only from dividends but also from capital gains or losses (appreciation or depreciation of the asset). In the current context, this is equal to $\dot{V}$. Therefore, instead of $\rho V=d$, the no-arbitrage condition becomes

$$\rho V\left(x\right)=d+\dot{V}\left(x\right).$$

Thus, at an intuitive level, the Maximum Principle (for stationary problems) amounts to requiring that the maximized value of dynamic maximization program, $V\left(x\right)$, and its rate of change, $\dot{V}\left(x\right)$, should be consistent with this no-arbitrage condition.

7.3.5. Proof of Theorem 7.9*. In this subsection, I provide a *sketch* of the proof of Theorem 7.9. A fully rigorous proof of Theorem 7.9 is quite long and involved. It can be found in a number of sources mentioned in the references below. The version provided here contains all the basic ideas, but is stated under the assumption that $V\left(t,x\right)$ is twice differentiable in t and x. As discussed above, the assumption that $V\left(t,x\right)$ is differentiable in t and x is not particularly restrictive and Theorem 7.17 below provide sufficient conditions for this to be the case. Nevertheless, the additional assumption that it is twice differentiable is somewhat more stringent.

The main idea of the proof is due to Pontryagin and co-authors. Instead of smooth variations from the optimal pair $\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)$, the method of proof considers "needle-like" variations, that is, piecewise continuous paths for the control variable that can deviate from the optimal control path by an arbitrary amount for a small interval of time.

Sketch Proof of Theorem 7.9: Suppose that the admissible pair $(\hat{x}(t), \hat{y}(t))$, with $\hat{y}(t) \in \text{Int}\mathcal{Y}(t)$ and $\hat{x}(t) \in \text{Int}\mathcal{X}(t)$, is a solution and attains the maximal value $V(0, x_0)$. Take an arbitrary $t_0 \in \mathbb{R}_+$. Construct the following perturbation: $y_\delta(t) = \hat{y}(t)$ for all $t \in [0, t_0)$ and for some sufficiently small Δt and $\delta \in \mathbb{R}$, $y_\delta(t) = \delta$ for $t \in [t_0, t_0 + \Delta t]$ for all $t \in [t_0, t_0 + \Delta t]$. Moreover, let $y_\delta(t)$ for $t \geq t_0 + \Delta t$ be the optimal control for $V(t_0 + \Delta t, x_\delta(t_0 + \Delta t))$, where $x_\delta(t)$ is the value of the state variable resulting from the perturbed control y_δ, with $x_\delta(t_0 + \Delta t)$ being the value at time $t_0 + \Delta t$. Note by construction $x_\delta(t_0) = \hat{x}(t_0)$ (since $y_\delta(t) = \hat{y}(t)$ for all $t \in [0, t_0]$).

Since the pair $(\hat{x}(t), \hat{y}(t))$ is optimal,

$$\begin{aligned} V(t_0, \hat{x}(t_0)) &= \int_{t_0}^{\infty} f(t, \hat{x}(t), \hat{y}(t))\, dt \\ &\geq \int_{t_0}^{\infty} f(t, x_\delta(t), y_\delta(t))\, dt \\ &= \int_{t_0}^{t_0+\Delta t} f(t, x_\delta(t), y_\delta(t))\, dt + V(t_0 + \Delta t, x_\delta(t_0 + \Delta t)), \end{aligned}$$

where the last equality uses the fact that the admissible pair $(x_\delta(t), y_\delta(t))$ is optimal starting with state variable $x_\delta(t_0 + \Delta t)$ at time $t_0 + \Delta t$. Rearranging terms and dividing by Δt yields

$$\frac{V(t_0 + \Delta t, x_\delta(t_0 + \Delta t)) - V(t_0, \hat{x}(t_0))}{\Delta t} \leq -\frac{\int_{t_0}^{t_0+\Delta t} f(t, x_\delta(t), y_\delta(t))\, dt}{\Delta t} \text{ for all } \Delta t \geq 0.$$

Now take limits as $\Delta t \to 0$ and note that $x_\delta(t_0) = \hat{x}(t_0)$ and that

$$\lim_{\Delta t \to 0} \frac{\int_{t_0}^{t_0+\Delta t} f(t, x_\delta(t), y_\delta(t))\, dt}{\Delta t} = f(t_0, x_\delta(t_0), y_\delta(t_0)).$$

Moreover, let $\mathcal{T} \subset \mathbb{R}_+$ be the set of points where the optimal control $\hat{y}(t)$ is a continuous function of time. Note that $\mathcal{T}$ is a dense subset of $\mathbb{R}_+$ since $\hat{y}(t)$ is a piecewise continuous function. Let us now take V to be a differentiable function of time at all $t \in \mathcal{T}$, so that

$$\begin{aligned} \lim_{\Delta t \to 0} \frac{V(t_0 + \Delta t, x_\delta(t_0 + \Delta t)) - V(t_0, \hat{x}(t_0))}{\Delta t} &= \frac{\partial V(t, x_\delta(t))}{\partial t} + \frac{\partial V(t, x_\delta(t))}{\partial x} \dot{x}_\delta(t), \\ &= \frac{\partial V(t, x_\delta(t))}{\partial t} + \frac{\partial V(t, x_\delta(t))}{\partial x} g(t, x_\delta(t), y_\delta(t)), \end{aligned}$$

where $\dot{x}_\delta(t) = g(t, x_\delta(t), y_\delta(t))$ is the law of motion of the state variable given by (7.29) together with the control y_δ. Putting all these together,

$$f(t_0, x_\delta(t_0), y_\delta(t_0)) + \frac{\partial V(t_0, x_\delta(t_0))}{\partial t} + \frac{\partial V(t_0, x_\delta(t_0))}{\partial x} g(t_0, x_\delta(t_0), y_\delta(t_0)) \leq 0$$

for all $t_0 \in \mathcal{T}$ (which correspond to points of continuity of $\hat{y}(t)$) and for all admissible perturbation pairs $(x_\delta(t), y_\delta(t))$. Moreover, from Theorem 7.10, which applies at all $t_0 \in \mathcal{T}$,

$$f(t_0, \hat{x}(t_0), \hat{y}(t_0)) + \frac{\partial V(t_0, \hat{x}(t_0))}{\partial t} + \frac{\partial V(t_0, \hat{x}(t_0))}{\partial x} g(t_0, \hat{x}(t_0), \hat{y}(t_0)) = 0. \tag{7.46}$$

Once more using the fact that $x_\delta(t_0) = \hat{x}(t_0)$, this implies that

$$f(t_0, \hat{x}(t_0), \hat{y}(t_0)) + \frac{\partial V(t_0, \hat{x}(t_0))}{\partial x} g(t_0, \hat{x}(t_0), \hat{y}(t_0)) \geq f(t_0, x_\delta(t_0), y_\delta(t_0)) + \frac{\partial V(t_0, \hat{x}(t_0))}{\partial x} g(t_0, x_\delta(t_0), y_\delta(t_0)) \tag{7.47}$$

for all $t_0 \in \mathcal{T}$ and for all admissible perturbation pairs $(x_\delta(t), y_\delta(t))$. Now defining

$$\lambda(t_0) \equiv \frac{\partial V(t_0, \hat{x}(t_0))}{\partial x}, \tag{7.48}$$

inequality (7.47) can be written as

$$\begin{aligned} f(t_0, \hat{x}(t_0), \hat{y}(t_0)) + \lambda(t_0) g(t_0, \hat{x}(t_0), \hat{y}(t_0)) \quad &\geq \quad f(t_0, x_\delta(t_0), y_\delta(t_0)) \\ &\quad + \lambda(t_0) g(t_0, x_\delta(t_0), y_\delta(t_0)), \end{aligned}$$

or equivalently,

$$\begin{aligned} H(t_0, \hat{x}(t_0), \hat{y}(t_0)) \quad &\geq \quad H(t_0, x_\delta(t_0), y_\delta(t_0)) \\ &\quad \text{for all admissible } (x_\delta(t_0), y_\delta(t_0)). \end{aligned}$$

Therefore,

$$H(t, \hat{x}(t), \hat{y}(t)) \geq \max_y H(t, \hat{x}(t), y).$$

This establishes the Maximum Principle.

The necessary condition (7.35) directly follows from the Maximum Principle together with the fact that H is differentiable in x and y (a consequence of the fact that f and g are differentiable in x and y). Condition (7.37) holds by definition. Finally, (7.36) follows from differentiating (7.46) with respect to x at all points of continuity of $\hat{y}(t)$, which gives

$$\begin{aligned} &\frac{\partial f(t, \hat{x}(t), \hat{y}(t))}{\partial x} + \frac{\partial^2 V(t, \hat{x}(t))}{\partial t \partial x} \\ + \quad &\frac{\partial^2 V(t, \hat{x}(t))}{\partial x^2} g(t, \hat{x}(t), \hat{y}(t)) + \frac{\partial V(t, \hat{x}(t))}{\partial x} \frac{\partial g(t, \hat{x}(t), \hat{y}(t))}{\partial x} = 0, \end{aligned}$$

for all for all $t \in \mathcal{T}$. Using the definition of the Hamiltonian, this implies (7.36) and completes the proof. □

7.4. More on Transversality Conditions

I next turn to a study of the boundary conditions at infinity in infinite-horizon maximization problems. As in the discrete-time optimization problems, these limiting boundary conditions are referred to as "transversality conditions". As mentioned above, a natural conjecture might be that, as in the finite-horizon case, the transversality condition should be similar to that in Theorem 7.1, with t_1 replaced with the limit of $t \to \infty$, that is, $\lim_{t \to \infty} \lambda(t) = 0$. The following example, which is very close to the original model that Frank Ramsey studied in 1928, illustrates that this is not the case; without further assumptions, the valid transversality condition is given by the weaker condition (7.43).

Example 7.2. Consider the following problem:

$$\max \int_0^\infty \left[\log\left(c\left(t\right)\right) - \log c^*\right] dt$$

subject to

$$\dot{k}\left(t\right) = \left[k\left(t\right)\right]^\alpha - c\left(t\right) - \delta k\left(t\right)$$

$$k\left(0\right) = 1$$

and

$$\lim_{t\to\infty} k\left(t\right) \geq 0$$

where $c^* \equiv \left[k^*\right]^\alpha - \delta k^*$ and $k^* \equiv \left(\alpha/\delta\right)^{1/(1-\alpha)}$. In other words, c^* is the maximum level of consumption that can be achieved in the steady state of this model and k^* is the corresponding steady-state level of capital. This way of writing the objective function makes sure that the integral converges and takes a finite value (since $c\left(t\right)$ cannot exceed c^* forever).

The Hamiltonian is straightforward to construct; it does not explicitly depend on time and takes the form

$$H\left(k,c,\lambda\right) = \left[\log c\left(t\right) - \log c^*\right] + \lambda\left[k\left(t\right)^\alpha - c\left(t\right) - \delta k\left(t\right)\right],$$

and implies the following necessary conditions

$$\begin{aligned} H_c\left(k,c,\lambda\right) &= \frac{1}{c\left(t\right)} - \lambda\left(t\right) = 0 \\ H_k\left(k,c,\lambda\right) &= \lambda\left(t\right)\left(\alpha k\left(t\right)^{\alpha-1} - \delta\right) = -\dot{\lambda}\left(t\right). \end{aligned}$$

It can be verified that any optimal path must feature $c\left(t\right) \to c^*$ as $t \to \infty$. This, however, implies that

$$\lim_{t\to\infty} \lambda\left(t\right) = \frac{1}{c^*} > 0 \text{ and } \lim_{t\to\infty} k\left(t\right) = k^*.$$

Now recall from Theorem 7.3 that the finite-horizon transversality condition in this case would have been $\lambda\left(t_1\right)k\left(t_1\right) = 0$, whereas here $\lim_{t\to\infty}\lambda\left(t\right)k\left(t\right) = k^*/c^* > 0$. Therefore, the equivalent of th finite-horizon transversality conditions do not hold. It can be verified, however, that along the optimal path the following condition holds:

$$\lim_{t\to\infty} H\left(k\left(t\right),c\left(t\right),\lambda\left(t\right)\right) = 0.$$

The next theorem shows that this is indeed a useful transversality condition for infinite-horizon optimization problems.

Theorem 7.12. ***(Transversality Condition for Infinite-Horizon Problems)*** *Suppose that problem of maximizing (7.28) subject to (7.29) and (7.30), with f and g continuously differentiable, has a piecewise continuous optimal control $\hat{y}\left(t\right) \in Int\mathcal{Y}\left(t\right)$ with corresponding path of state variable $\hat{x}\left(t\right) \in Int\mathcal{X}\left(t\right)$. Let $V\left(t,x\left(t\right)\right)$ be the value function defined in (7.31). Suppose that $V\left(t,\hat{x}\left(t\right)\right)$ is differentiable in x and t for t sufficiently large and that $\lim_{t\to\infty}\partial V\left(t,\hat{x}\left(t\right)\right)/\partial t = 0$. Let $H\left(t,x,y,\lambda\right)$ be given by (7.12). Then, the optimal control*

$\hat{y}(t)$ and the corresponding path of the state variable $\hat{x}(t)$ satisfy the necessary conditions (7.35)-(7.37) and the transversality condition

$$\lim_{t\to\infty} H\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right),\lambda\left(t\right)\right)=0. \tag{7.49}$$

PROOF. For t sufficiently large, $V(t,x)$ is differentiable in t and x, thus $\partial V\left(t,\hat{x}\left(t\right)\right)/\partial x=\lambda\left(t\right)$. Then, the Hamilton-Jacobi-Bellman equation, (7.38), holds and implies

$$\begin{aligned}\frac{\partial V\left(t,\hat{x}\left(t\right)\right)}{\partial t}+f\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right)\right)+\lambda\left(t\right)g\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right)\right) &= 0 \text{ for } t \text{ large}\\ \frac{\partial V\left(t,\hat{x}\left(t\right)\right)}{\partial t}+H\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right),\lambda\left(t\right)\right) &= 0 \text{ for } t \text{ large.}\end{aligned} \tag{7.50}$$

Now take the limit as $t\to\infty$. Since $\lim_{t\to\infty}\partial V\left(t,\hat{x}\left(t\right)\right)/\partial t=0$ by hypothesis, (7.50) then implies (7.49) (cf. Fact A.5(5) in Appendix Chapter A). □

As discussed in the previous chapter, the condition that $\lim_{t\to\infty}V\left(t,\hat{x}\left(t\right)\right)$ is finite is natural, since economic problems where agents reach infinite value are not relevant (and moreover, if this limit were equal to infinity, there would typically be no pair $\left(x\left(t\right),y\left(t\right)\right)$ reaching this value). The hypothesis here, $\lim_{t\to\infty}\partial V\left(t,\hat{x}\left(t\right)\right)/\partial t=0$, is only slightly stronger than that $\lim_{t\to\infty}V\left(t,\hat{x}\left(t\right)\right)$ exists and is finite. It is again satisfied in almost all economic problems and will be satisfied in all of the models studied in this book. However, the transversality condition (7.49) is not particularly convenient to work with. In the next section, I present a stronger and more useful version of this transversality condition that applies in the context of discounted infinite-horizon problems.

EXAMPLE 7.3. One of the most common examples of this type of dynamic optimization problem is that of the optimal time path of consuming a non-renewable resource. In particular, imagine the problem of an infinitely-lived individual that has access to a non-renewable or exhaustible resource of size 1. The instantaneous utility of consuming a flow of resources y is $u(y)$, where $u:[0,1]\to\mathbb{R}$ is a strictly increasing, continuously differentiable and strictly concave function. The individual discounts the future exponentially with discount rate $\rho>0$, so that his objective function at time $t=0$ is to maximize

$$\int_0^\infty \exp\left(-\rho t\right)u\left(y\left(t\right)\right)dt.$$

The constraint is that the remaining size of the resource at time t, $x(t)$ evolves according to

$$\dot{x}\left(t\right)=-y\left(t\right),$$

which captures the fact that the resource is not renewable and becomes depleted as more of it is consumed. Naturally, we also need that $x(t)\geq 0$.

The current-value Hamiltonian takes the form

$$\hat{H}\left(x\left(t\right),y\left(t\right),\mu\left(t\right)\right)=u\left(y\left(t\right)\right)-\mu\left(t\right)y\left(t\right).$$

Theorem 7.13 implies the following necessary condition for an interior continuously differentiable solution $\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)$ to this problem. There should exist a continuously differentiable

function $\mu(t)$ such that

$$u'(\hat{y}(t)) = \mu(t),$$

and

$$\dot{\mu}(t) = \rho\mu(t).$$

The second condition follows since neither the constraint nor the objective function depend on $x(t)$. This is the famous *Hotelling rule* for the exploitation of exhaustible resources. It charts a path for the shadow value of the exhaustible resource. In particular, integrating both sides of this equation we obtain

$$\mu(t) = \mu(0)\exp(\rho t).$$

Now combining this with the first-order condition for $y(t)$,

$$\hat{y}(t) = u'^{-1}[\mu(0)\exp(\rho t)],$$

where $u'^{-1}[\cdot]$ is the inverse function of u', which exists and is strictly decreasing by virtue of the fact that u is strictly concave. This equation immediately implies that the amount of the resource consumed is monotonically decreasing over time. This is economically intuitive: because of discounting, there is preference for early consumption, whereas delayed consumption has no return (there is no production or interest payments on the stock). Nevertheless, the entire resource is not consumed immediately, since there is also a preference for smooth consumption arising from the fact that $u(\cdot)$ is strictly concave.

Combining the previous equation with the resource constraint gives

$$\dot{x}(t) = -u'^{-1}[\mu(0)\exp(\rho t)].$$

Integrating this equation and using the terminal value constraint that $x(0) = 1$,

$$\hat{x}(t) = 1 - \int_0^t u'^{-1}[\mu(0)\exp(\rho s)]\,ds.$$

Since along any optimal path $\lim_{t\to\infty}\hat{x}(t) = 0$, we also have

$$\int_0^\infty u'^{-1}[\mu(0)\exp(\rho s)]\,ds = 1.$$

Therefore, the initial value of the costate variable $\mu(0)$ must be chosen so as to satisfy this equation. You are asked to verify that the transversality condition (7.49) is satisfied in Exercise 7.19.

7.5. Discounted Infinite-Horizon Optimal Control

Part of the difficulty, especially regarding the absence of a transversality condition, comes from the fact that we did not impose enough structure on the functions f and g. As discussed above, our interest is with the growth models where the utility is discounted exponentially. Consequently, economically interesting problems often take the following more specific form:

$$\max_{x(t),y(t)} W(x(t), y(t)) \equiv \int_0^\infty \exp(-\rho t) f(x(t), y(t))\,dt \text{ with } \rho > 0, \tag{7.51}$$

subject to

$$\dot{x}(t) = g(t, x(t), y(t)), \tag{7.52}$$

and

$$y(t) \in \text{Int}\mathcal{Y}(t),\ x(t) \in \text{Int}\mathcal{X}(t) \text{ for all } t,\ x(0) = x_0 \text{ and } \lim_{t\to\infty} b(t)\, x(t) \geq x_1, \tag{7.53}$$

where recall again that $b : \mathbb{R}_+ \to \mathbb{R}_+$ and $\lim_{t\to\infty} b(t) < \infty$. Throughout the assumption that there is in fact *discounting*, that is, $\rho > 0$, is implicit.

The special feature of this problem is that the objective function, f, depends on time only through exponential discounting. The Hamiltonian in this case would be:

$$\begin{aligned} H(t, x(t), y(t), \lambda(t)) &= \exp(-\rho t) f(x(t), y(t)) + \lambda(t) g(t, x(t), y(t)) \\ &= \exp(-\rho t) [f(x(t), y(t)) + \mu(t) g(t, x(t), y(t))], \end{aligned}$$

where the second line defines

$$\mu(t) \equiv \exp(\rho t) \lambda(t). \tag{7.54}$$

In fact, in this case, rather than working with the standard Hamiltonian, we can work with *the current-value Hamiltonian*, defined as

$$\hat{H}(t, x(t), y(t), \mu(t)) \equiv f(x(t), y(t)) + \mu(t) g(t, x(t), y(t)). \tag{7.55}$$

When $g(t, x(t), y(t))$ is also an autonomous differential equation of the form $g(x(t), y(t))$, I simplify the notation by writing $\hat{H}(x(t), y(t), \mu(t))$ instead of $\hat{H}(t, x(t), y(t), \mu(t))$.

The next result establishes the necessity of a stronger transversality condition under some additional assumptions. These assumptions can be relaxed, though the result is simpler to understand and prove under these assumptions, and they are typically met in economic applications. Throughout, f and g are continuously differentiable for all admissible $(x(t), y(t))$ (with derivatives f_x, f_y, g_x and g_y).

Assumption 7.1: In the maximization of (7.51) subject to (7.52) and (7.53):

(i) f is weakly monotone in (x, y) and g is weakly monotone in (t, x, y) (for example, f could be nondecreasing in x and nonincreasing in y, and so on);

(ii) there exists $m > 0$ such that $|g_y(t, x(t), y(t))| \geq m$ for all t and for all admissible pairs $(x(t), y(t))$;

(iii) there exists $M < \infty$ such that $|f_y(x, y)| \leq M$ for all x and y.

Theorem 7.13. ***(Maximum Principle for Discounted Infinite-Horizon Problems)*** *Suppose that the problem of maximizing (7.51) subject to (7.52) and (7.53), with f and g continuously differentiable, has an interior piecewise continuous optimal control $\hat{y}(t) \in Int\mathcal{Y}(t)$ with corresponding state variable $\hat{x}(t) \in Int\mathcal{X}(t)$. Let $V(t, x(t))$ be the value function defined in (7.31). Suppose that $V(t, \hat{x}(t))$ is differentiable in x and t for t sufficiently large, that $\lim_{t\to\infty} V(t, \hat{x}(t))$ exists and is finite, and that $\lim_{t\to\infty} \partial V(t, \hat{x}(t))/\partial t = 0$. Let $\hat{H}(t, x, y, \mu)$ be the current-value Hamiltonian given by (7.55). Then, except at points of*

discontinuity of $\hat{y}(t)$, the optimal control pair $(\hat{x}(t), \hat{y}(t))$ satisfies the following necessary conditions:

$$\hat{H}_y(t, \hat{x}(t), \hat{y}(t), \mu(t)) = 0 \text{ for all } t \in \mathbb{R}_+, \tag{7.56}$$

$$\rho\mu(t) - \dot{\mu}(t) = \hat{H}_x(t, \hat{x}(t), \hat{y}(t), \mu(t)) \text{ for all } t \in \mathbb{R}_+, \tag{7.57}$$

$$\dot{x}(t) = \hat{H}_\mu(t, \hat{x}(t), \hat{y}(t), \mu(t)) \text{ for all } t \in \mathbb{R}_+, \; x(0) = x_0 \text{ and } \lim_{t\to\infty} b(t)\, x(t) \geq x_1, \tag{7.58}$$

and the transversality condition

$$\lim_{t\to\infty} \left[\exp(-\rho t)\, \hat{H}(t, \hat{x}(t), \hat{y}(t), \mu(t))\right] = 0. \tag{7.59}$$

Moreover, if Assumption 7.1 is satisfied, then the transversality condition can be strengthened to:

$$\lim_{t\to\infty} [\exp(-\rho t)\, \mu(t)\, \hat{x}(t)] = 0. \tag{7.60}$$

PROOF. The derivation of the necessary conditions (7.56)-(7.58) and the transversality condition (7.59) follow by using the definition of the current-value Hamiltonian and from Theorem 7.12. They are left for as an exercise (see Exercise 7.14).

I only give the proof for the stronger transversality condition (7.60). The weaker transversality condition (7.59) can be written as

$$\lim_{t\to\infty} [\exp(-\rho t)\, f(\hat{x}(t), \hat{y}(t)) + \exp(-\rho t)\, \mu(t)\, g(t, \hat{x}(t), \hat{y}(t))] = 0. \tag{7.61}$$

Recall from (7.33) that $V(t, \hat{x}(t)) = \int_t^\infty \exp(-\rho s)\, f(\hat{x}(s), \hat{y}(s))\, ds$. Since $\lim_{t\to\infty} V(t, \hat{x}(t))$ exists and is finite, the first term in (7.61) must be equal to zero. Therefore,

$$\lim_{t\to\infty} [\exp(-\rho t)\, \mu(t)\, g(t, \hat{x}(t), \hat{y}(t))] = \lim_{t\to\infty} [\exp(-\rho t)\, \mu(t)\, \dot{x}(t)] = 0. \tag{7.62}$$

In the rest of the proof, let lim stand for $\lim_{t\to\infty}$. Also throughout the proof, the statement that "$\lim z(t)$ exists for some process $z(t)$" means that the limit exists in extended real numbers and thus allows for $\lim z(t) = \infty$ or $-\infty$; $\lim z(t) = z^* \in \mathbb{R}$ implies that the limit exists and is finite. Recall also from (7.54) that

$$\lambda(t) \equiv \exp(-\rho t)\, \mu(t).$$

As a first observation, note that from Fact A.5(4) in Appendix Chapter A, (7.62) implies

$$\lim \lambda(t)\dot{x}(t) = \lim \lambda(t)\, |\dot{x}(t)| = \lim |\lambda(t)|\, |\dot{x}(t)| = 0. \tag{7.63}$$

I next present a result that will be used a number of times in the rest of the proof. Recall the definition of a *net* from Definition A.7 in Appendix Chapter A. In view of this, the process $[\mu(t)]_{t=0}^\infty$ can be viewed as a net and denoted by $\{\mu(t)\}_{t\in\mathbb{R}_+}$ or simply by $\{\mu(t)\}$. Then, $\{\mu(t)\}_{t\in\mathcal{T}}$ for $\mathcal{T} \subset \mathbb{R}_+$ is a subnet of $\{\mu(t)\}$.

CLAIM 1. *The process (net) $[\mu(t)]_{t=0}^\infty$ is bounded in the sense that there exists $B < \infty$ such that $|\mu(t)| < B$ for all t*

Proof. Suppose, to obtain a contradiction, that $\{\mu(t)\}$ is not bounded. This implies that there exists a subnet $\{\mu(t)\}_{t\in\mathcal{T}}$ limiting to $+\infty$ or $-\infty$ (for $\mathcal{T}\subset\mathbb{R}_+$). Since $g_y \geq m > 0$ or $g_y \leq m < 0$ (cf. Assumption 7.1(i) and (ii)), $\lim_{t\in\mathcal{T}} \mu(t) g_y(t, \hat{x}(t), \hat{y}(t)) = \pm\infty$. But in this case, (7.64) can only hold if $\lim_{t\in\mathcal{T}} f_y(\hat{x}(t), \hat{y}(t)) = \mp\infty$, which contradicts Assumption 7.1(iii) and establishes that $\{\mu(t)\}$ is bounded. □

Now, let us consider two cases, which will together make up the proof of the result.

Case I: $\lim \hat{x}(t)$ exists and is finite. This immediately implies that we must also have $\lim \dot{x}(t) = 0$. Let $\lim \hat{x}(t) = \hat{x}^* \in \mathbb{R}$. Taking limits in (7.56),

$$\lim \left[f_y(\hat{x}(t), \hat{y}(t)) + \mu(t) g_y(t, \hat{x}(t), \hat{y}(t))\right] = 0. \tag{7.64}$$

Since $\lim \hat{x}(t) = \hat{x}^*$ (and thus $\lim |\hat{x}(t)| = |\hat{x}^*|$), $\lim \exp(-\rho t) = 0$ and $|\mu(t)| < B$, we have

$$\begin{aligned}\lim \left[\exp(-\rho t)\mu(t)|\hat{x}(t)|\right] &\leq B \lim \left[\exp(-\rho t)|\hat{x}(t)|\right] \\ &= B|\hat{x}^*| \lim \exp(-\rho t) \\ &= 0.\end{aligned}$$

But also,

$$\begin{aligned}\lim \left[\exp(-\rho t)\mu(t)|\hat{x}(t)|\right] &\geq -B \lim \left[\exp(-\rho t)|\hat{x}(t)|\right] \\ &= -B|\hat{x}^*| \lim \exp(-\rho t) \\ &= 0.\end{aligned}$$

Therefore, $\lim \left[\exp(-\rho t)\mu(t)|\hat{x}(t)|\right] = \lim \left[\exp(-\rho t)\mu(t)\hat{x}(t)\right] = 0$, establishing (7.60).

Case II: suppose that $\lim \hat{x}(t)$ does not exist or is not finite. Recall from (7.56), that $\exp(-\rho t) f_y(\hat{x}(t), \hat{y}(t)) + \lambda(t) g_y(t, \hat{x}(t), \hat{y}(t)) = 0$, so that from Assumption 7.1(i), we must have that for all t, either $f_y \geq 0$ and $g_y \leq 0$, or $f_y \leq 0$ and $g_y \geq 0$, or $f_y \geq 0$ and $g_y \geq 0$, or finally $f_y \leq 0$ and $g_y \leq 0$. This implies that $\lambda(t) \geq 0$ for all t (in the first two cases) or $\lambda(t) \leq 0$ for all t (in the last two cases).

Suppose $\lambda(t) \geq 0$. From (7.62), we still have $\lim \lambda(t)\dot{x}(t) = 0$ and from Fact A.5(4),

$$\lim \lambda(t)\dot{x}(t) = \lim \lambda(t)|\dot{x}(t)| = 0. \tag{7.65}$$

Next, define

$$z(t) = \sup_{s\in[0,t]} |\dot{x}(s)|.$$

Clearly $\lim z(t)$ exists, though it may be infinite (cf. Fact A.5(4)). Observe also that (7.65) implies

$$\lim \lambda(t)|z(t)| = \lim |\lambda(t)||z(t)| = 0. \tag{7.66}$$

To see this, first consider the case in which $\dot{x}(t)$ is bounded. Then $\lim |z(t)|$ must exist and be finite, say $z^* \in \mathbb{R}$. Claim 1 then again implies that

$$0 = -Bz^* \lim \exp(-\rho t) \leq \lim \left[\exp(-\rho t)\mu(t)|z(t)|\right] \leq Bz^* \lim \exp(-\rho t) = 0,$$

so that $\lim \lambda(t)|z(t)| = 0$. Next, suppose that $\dot{x}(t)$ is unbounded. Then $\{|z(t)|\}$ is identical to a subnet $\{|\dot{x}(t)|\}_{t\in\mathcal{T}}$ of $\{|\dot{x}(t)|\}$. Since $\lim \lambda(t)|\dot{x}(t)| = 0$, we must also have $\lim_{t\in\mathcal{T}} \lambda(t)|\dot{x}(t)| = 0$ and therefore $\lim \lambda(t)|z(t)| = 0$, again implying (7.66).

Now define

$$\tilde{x}(t) = \int_0^t |z(s)|\,ds + x_0 \geq \int_0^t \dot{x}(s)ds + x_0 = \hat{x}(t).$$

Since $|\dot{x}(t)| \geq 0$ for all t, $\tilde{x}(t)$ is nondecreasing and thus $\lim \tilde{x}(t)$ exists (Fact A.4). Consider two subcases:

(i) $\lim z(t)$ is equal to infinity, which implies that $|\tilde{x}(t)|$ grows *at an exponential rate or faster*. Therefore, there exists $\gamma > 0$ such that for every $\varepsilon \in (0,\gamma)$ and some $T < \infty$, we have that for all $t \geq T$, $z(t) \geq (\gamma - \varepsilon)\tilde{x}(t)$. Substituting this in (7.66),

$$\lim |\lambda(t)|\,|z(t)| \geq (\gamma - \varepsilon) \lim |\lambda(t)|\,|\tilde{x}(t)| \geq 0, \tag{7.67}$$

Since $\lim |\lambda(t)|\,|z(t)| = 0$, both inequalities in (7.67) must hold as equality and therefore $\lim [\exp(-\rho t)\mu(t)\tilde{x}(t)] = 0$.

(ii) $\lim z(t) = z^* \in \mathbb{R}$, which implies that $|\tilde{x}(t)|$ grows at *less than an exponential rate*. In this case, for any $\gamma > 0$ and some T, we have that $|z(t)/\tilde{x}(t)| < \gamma$ for all $t \geq T < \infty$. Since $\rho > 0$, this implies that $\lim |\exp(-\rho t)|\,|\tilde{x}(t)| = 0$. Claim 1 again implies that

$$0 = -B \lim [\exp(-\rho t)\,\mu(t)\,|\tilde{x}(t)|] \leq \lim [\exp(-\rho t)\,\mu(t)\,|\tilde{x}(t)|] \leq B \lim [\exp(-\rho t)\,\mu(t)\,|\tilde{x}(t)|] = 0,$$

and thus $\lim [\exp(-\rho t)\mu(t)\tilde{x}(t)] = 0$. These two subcases together imply that

$$\lim [\lambda(t)\,\tilde{x}(t)] = 0. \tag{7.68}$$

Since $\lambda(t) \geq 0$ and $\tilde{x}(t) \geq \hat{x}(t)$, we then have

$$0 = \lim \lambda(t)\,\tilde{x}(t) \geq \lim \lambda(t)\,\hat{x}(t). \tag{7.69}$$

Next, repeat the same argument with

$$\vec{x}(t) = -\int_0^t |z(s)|\,ds + x_0 \leq \hat{x}(t).$$

Since $\vec{x}(t)$ is nonincreasing for all t, the same steps in subcases (i) and (ii) now yield $\lim \lambda(t)\,\vec{x}(t) = 0$. Since $\lambda(t) \geq 0$ and $\vec{x}(t) \leq \hat{x}(t)$,

$$0 = \lim \lambda(t)\,\vec{x}(t) \leq \lim \lambda(t)\,\hat{x}(t). \tag{7.70}$$

Combining (7.69) with (7.70) implies that $\lim \lambda(t)\,\hat{x}(t) = 0$ for the case in which $\lambda(t) \geq 0$.

The argument for the case where $\lambda(t) \leq 0$ is identical and completes the proof of the theorem. □

The proof of Theorem 7.13 also clarifies the importance of discounting. For example, neither the key equation, (7.64) nor the argument in subcase (ii) in the proof apply without discounting. Exercise 7.17 discusses how a similar result can be established without discounting under stronger assumptions. Exercise 7.20, on the other hand, shows why Assumption 7.1(iii) is necessary. This condition will be the harder one to verify in some economic models,

though in all models studied in the rest of the book we can make sure that it is satisfied and thus the current version of Theorem 7.13 can be applied directly in all of these cases.

Notice that compared to the transversality condition in the finite-horizon case (e.g., Theorem 7.1), there is the additional term $\exp(-\rho t)$ in (7.60). This is because the transversality condition applies to the original costate variable $\lambda(t)$, i.e., $\lim_{t\to\infty}[x(t)\lambda(t)]=0$, and as shown above, the current-value costate variable $\mu(t)$ is given by $\mu(t)=\exp(\rho t)\lambda(t)$. Note also that the stronger transversality condition takes the form $\lim_{t\to\infty}[\exp(-\rho t)\mu(t)\hat{x}(t)]=0$, not simply $\lim_{t\to\infty}[\exp(-\rho t)\mu(t)]=0$. Exercise 7.18 illustrates why this is.

It is important to emphasize that Theorem 7.13 only provides necessary conditions for interior continuous solutions. In this light, (7.60) should be interpreted as a *necessary transversality condition* for such a solution. This implies that (7.60) is neither necessary (the solution may not be interior or continuous) nor sufficient. However, the next theorem shows that under the appropriate concavity conditions (7.60) is also a *sufficient transversality condition*. More remarkably, it also shows that for concave problems, Assumption 7.1 is no longer necessary, thus, in concave problems, we can use (7.60) without having to check Assumption 7.1. Therefore, the next theorem will be the most important result of this chapter and will be used repeatedly in applications.

Theorem 7.14. ***(Sufficiency Conditions for Discounted Infinite-Horizon Problems)*** *Consider the problem of maximizing (7.51) subject to (7.52) and (7.53), with f and g continuously differentiable. Define $\hat{H}(t,x,y,\mu)$ as the current-value Hamiltonian as in (7.55), and suppose that a solution $\hat{y}(t)$ and the corresponding path of state variable $\hat{x}(t)$ satisfy (7.56)-(7.59). Given the resulting current-value costate variable $\mu(t)$, define $M(t,x,\mu)\equiv\max_{y(t)\in\mathcal{Y}(t)}\hat{H}(t,x,y,\mu)$. Suppose that $\lim_{t\to\infty}V(t,\hat{x}(t))$ exists and is finite (where $V(t,x(t))$ is defined in (7.34)), that for any admissible pair $(x(t),y(t))$, $\lim_{t\to\infty}[\exp(-\rho t)\mu(t)x(t)]\geq 0$, and that$\mathcal{X}(t)$ is convex and $M(t,x,\mu)$ is concave in $x\in\mathcal{X}(t)$ for all t. Then, the pair $(\hat{x}(t),\hat{y}(t))$ achieves the unique global maximum of (7.51). Moreover, if $M(t,x,\mu)$ is strictly concave in x, $(\hat{x}(t),\hat{y}(t))$ is the unique solution to (7.51).*

Proof. Consider $(\hat{x}(t),\hat{y}(t))$ that satisfy (7.56)-(7.58) and (7.59) and an arbitrary admissible pair $(x(t),y(t))$. By the same argument as in the proof of Theorem 7.6

$$M(t,x(t),\mu(t))\leq M(t,\hat{x}(t),\mu(t))+M_x(t,\hat{x}(t),\mu(t))(x(t)-\hat{x}(t)) \text{ for all } t\geq 0.$$

Integrating both sides over $[0,\infty)$ yields

$$\int_0^\infty M(t,x(t),\mu(t))\,dt\leq\int_0^\infty M(t,\hat{x}(t),\mu(t))\,dt+\int_0^\infty M_x(t,\hat{x}(t),\mu(t))(x(t)-\hat{x}(t))\,dt.$$

Moreover,

$$\begin{aligned}M_x(t,\hat{x}(t),\mu(t)) &= \hat{H}_x(t,\hat{x}(t),\hat{y}(t),\mu(t))\\ &= -\dot{\mu}(t).\end{aligned}$$

Next, exploiting the definition of the maximized Hamiltonian,

$$\int_0^\infty M\left(t, x\left(t\right), \mu\left(t\right)\right) dt = W\left(x\left(t\right), y\left(t\right)\right) + \int_0^\infty \lambda\left(t\right) g\left(t, x\left(t\right), y\left(t\right)\right) dt,$$

where $\lambda\left(t\right) = \exp\left(-\rho t\right) \mu\left(t\right)$, and

$$\int_0^\infty M\left(t, \hat{x}\left(t\right), \mu\left(t\right)\right) dt = W\left(\hat{x}\left(t\right), \hat{y}\left(t\right)\right) + \int_0^\infty \lambda\left(t\right) g\left(t, \hat{x}\left(t\right), \hat{y}\left(t\right)\right) dt.$$

Combining these equations, we have

$$\begin{aligned} W\left(x\left(t\right), y\left(t\right)\right) \leq\ & W\left(\hat{x}\left(t\right), \hat{y}\left(t\right)\right) \\ & + \int_0^\infty \lambda\left(t\right) \left[g\left(t, \hat{x}\left(t\right), \hat{y}\left(t\right)\right) - g\left(t, x\left(t\right), y\left(t\right)\right)\right] dt \\ & - \int_0^\infty \dot{\lambda}\left(t\right) \left(x\left(t\right) - \hat{x}\left(t\right)\right) dt. \end{aligned} \tag{7.71}$$

Integrating the last term by parts and recalling that by $x\left(0\right) = \hat{x}\left(0\right) = x_0$, we obtain

$$\int_0^\infty \dot{\lambda}\left(t\right) \left(x\left(t\right) - \hat{x}\left(t\right)\right) dt = \lim_{t \to \infty} \left[\lambda\left(t\right) \left(x\left(t\right) - \hat{x}\left(t\right)\right)\right] - \int_0^{t_1} \lambda\left(t\right) \left(\dot{x}\left(t\right) - \dot{\hat{x}}\left(t\right)\right) dt.$$

In addition, from (7.59), $\lim_{t \to \infty} \left[\exp\left(-\rho t\right) \mu\left(t\right) \hat{x}\left(t\right)\right] = \lim_{t \to \infty} \left[\lambda\left(t\right) \hat{x}\left(t\right)\right] = 0$. Moreover, by hypothesis $\lim_{t \to \infty} \lambda\left(t\right) x\left(t\right) \geq 0$, and thus

$$\int_0^\infty \dot{\lambda}\left(t\right) \left(x\left(t\right) - \hat{x}\left(t\right)\right) dt \geq - \int_0^{t_1} \lambda\left(t\right) \left(\dot{x}\left(t\right) - \dot{\hat{x}}\left(t\right)\right) dt.$$

Combining this with (7.71),

$$\begin{aligned} W\left(x\left(t\right), y\left(t\right)\right) \leq\ & W\left(\hat{x}\left(t\right), \hat{y}\left(t\right)\right) \\ & + \int_0^{t_1} \lambda\left(t\right) \left[g\left(t, \hat{x}\left(t\right), \hat{y}\left(t\right)\right) - g\left(t, x\left(t\right), y\left(t\right)\right)\right] dt \\ & + \int_0^{t_1} \lambda\left(t\right) \left[\dot{x}\left(t\right) - \dot{\hat{x}}\left(t\right)\right] dt. \end{aligned} \tag{7.72}$$

By definition of the admissible pairs $\left(x\left(t\right), y\left(t\right)\right)$ and $\left(\hat{x}\left(t\right), \hat{y}\left(t\right)\right)$, $\dot{\hat{x}}\left(t\right) = g\left(t, \hat{x}\left(t\right), \hat{y}\left(t\right)\right)$ and $\dot{x}\left(t\right) = g\left(t, x\left(t\right), y\left(t\right)\right)$. Thus (7.72) implies that $W\left(x\left(t\right), y\left(t\right)\right) \leq W\left(\hat{x}\left(t\right), \hat{y}\left(t\right)\right)$ for any admissible pair $\left(x\left(t\right), y\left(t\right)\right)$, establishing the first part of the theorem.

If M is strictly concave in x, then the inequality in (7.72) is strict, and therefore the same argument establishes $W\left(x\left(t\right), y\left(t\right)\right) < W\left(\hat{x}\left(t\right), \hat{y}\left(t\right)\right)$, and no other $\hat{x}\left(t\right)$ could achieve the same value, establishing the second part. □

Theorem 7.14 is very useful and powerful. It states that, for a concave problem, an admissible pair $\left(\hat{x}\left(t\right), \hat{y}\left(t\right)\right)$ satisfying (7.56)-(7.58) as well as the transversality condition (7.59) will necessarily dominate any other admissible pair $\left(x\left(t\right), y\left(t\right)\right)$ for which $\lim_{t \to \infty} \left[\exp\left(-\rho t\right) \mu\left(t\right) x\left(t\right)\right] \geq 0$. For this sufficiency result, we do not need to check Assumption 7.1. Given this result, the following strategy will be used in most problems.

(1) We will use the conditions in Theorem 7.13 to locate *a candidate interior solution* $\left(\hat{x}\left(t\right), \hat{y}\left(t\right)\right)$ satisfying (7.56)-(7.58) and (7.59).

(2) We will then verify the concavity conditions of Theorem 7.14 and then simply check that $\lim_{t\to\infty}\left[\exp\left(-\rho t\right)\mu\left(t\right)x\left(t\right)\right]\geq 0$ for other admissible pairs. If these conditions are satisfied, we will have characterized a global maximum.

An important feature of this theorem and the strategy outlined here is that they can be directly applied to unbounded problems (for example, to endogenous growth models). Therefore, as long as the conditions of Theorem 7.14 are satisfied, there is no need to make boundedness assumptions to characterize household behavior or optimal growth solutions.

7.6. Existence of Solutions, Concavity and Differentiability*

The theorems presented so far characterize the properties of a solution to a continuous-time maximization problem. The natural question of when a solution exists has not been posed or answered so far. This might appear curious, since both in finite-dimensional and in discrete-time infinite-horizon optimization problems studied in the previous chapter the analysis starts with existence theorems. There is a good reason for this omission, however. In continuous-time optimization problems, establishing existence of solutions is considerably more difficult than the characterization of solutions. I now present the general theorem on existence of solutions to continuous-time optimization problems and two additional results providing conditions under which the value function $V\left(t,x\right)$, defined in eq. (7.31) and Lemma 7.1, is concave and differentiable.

The reader may have already wondered how valid the approach of using the necessary conditions provided so far, which did not verify the existence of a solution, would be in practice. This is an important concern, and ordinarily such an approach would open the door for potential mistakes. One line of defense, however, is provided by the sufficiency theorems, for example, Theorems 7.11 or 7.14 for infinite-horizon problems. If, given a continuous-time optimization problem, we find an admissible pair $\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)$ that satisfies the necessary conditions (e.g., those in Theorem 7.9) and we can then verify that the optimization problem satisfies the conditions in any one of Theorems 7.11 or 7.14, then we must have characterized an optimal solution and can dispense with an existence theorem. Therefore, the sufficiency results contained in these theorems enable us to bypass the step of checking for the existence of a solution. Nevertheless, this is only valid when the problem possesses sufficient concavity to satisfy the conditions of Theorems 7.11 or 7.14. For problems that do not fall into this category, the justification for using the necessary conditions is much weaker. For this reason, and to provide a more complete treatment of continuous-time optimization problems, in this section I present an existence theorem for optimal control problems. Unfortunately, however, existence theorems for this class of problems are both somewhat complicated to state and difficult to prove. In particular, they require both measure-theoretic ideas and more advanced tools (some of them stated in the starred section, Section A.5, of Appendix Chapter A). I therefore provide a sketch of the proof of one of the most useful and powerful theorems, but I only give a sketch of some of the measure-theoretic details.

For the purposes of this existence theorem, I consider the most general problem studied so far, which is

$$\max_{\mathbf{x}(t),\mathbf{y}(t)} W(\mathbf{x}(t),\mathbf{y}(t)) \equiv \int_0^\infty f(t,\mathbf{x}(t),\mathbf{y}(t))\,dt \tag{7.73}$$

subject to

$$\dot{\mathbf{x}}(t) = G(t,\mathbf{x}(t),\mathbf{y}(t)), \tag{7.74}$$

and

$$\mathbf{y}(t) \in \mathcal{Y}(t),\ \mathbf{x}(t) \in \mathcal{X}(t) \text{ for all } t,\ \mathbf{x}(0) = \mathbf{x}_0 \text{ and } \lim_{t\to\infty} \mathbf{x}(t) \geq \mathbf{x}_1, \tag{7.75}$$

where $\mathcal{Y}(t) \subset \mathbb{R}^{K_y}$ and $\mathcal{X}(t) \subset \mathbb{R}^{K_x}$, with $K_y, K_x \in \mathbb{N}$, and $f : \mathbb{R}\times\mathbb{R}^{K_x}\times\mathbb{R}^{K_y} \to \mathbb{R}$ and $G : \mathbb{R}\times\mathbb{R}^{K_x}\times\mathbb{R}^{K_y} \to \mathbb{R}^{K_x}$.

Clearly, finite-horizon problems are special cases of this more general problem, obtained by setting $f(t,\mathbf{x},\mathbf{y}) \equiv 0$ for all $t \geq t_1$ for some $t_1 > 0$. Also the terminal value constraint is taken to be of the specific form $\lim_{t\to\infty}\mathbf{x}(t) \geq \mathbf{x}_1$ for concreteness. Different forms of the terminal value constraints, which may involve some of the state variables being entirely free and others having to satisfy equality constraints, can be handled using exactly the same strategy of proof as in Theorem 7.15.

To state the result on existence of solutions, let us define

$$M \equiv \{(t,\mathbf{x},\mathbf{y}) : \mathbf{y} \in \mathcal{Y}(t),\ \mathbf{x} \in \mathcal{X}(t) \text{ and } t \in \mathbb{R}_+\},$$

and

$$M' \equiv \{(t,\mathbf{x}) : \mathbf{x}(t) \in \mathcal{X}(t) \text{ and } t \in \mathbb{R}_+\}.$$

Recall also $(\mathbf{x}(t),\mathbf{y}(t))$ is an admissible pair if $\mathbf{x}(t) \in \mathcal{X}(t)$ is *absolutely continuous* and $\mathbf{y}(t) \in \mathcal{Y}(t)$ is (*Lebesgue*) *measurable* for all t, and $\mathbf{x}(t)$ and $\mathbf{y}(t)$ satisfy (7.74) and $\mathbf{x}(t)$ satisfies the initial and terminal value constraints $\mathbf{x}(0) = \mathbf{x}_0$ and $\lim_{t\to\infty}\mathbf{x}(t) \geq \mathbf{x}_1$ (recall footnote 1 and see Definition A.26 in Appendix Chapter A). For a given initial condition $\mathbf{x}(0) = \mathbf{x}_0$, I denote the set of admissible pairs by

$$\Omega(0,\mathbf{x}_0) \equiv \{[\mathbf{x}(t),\mathbf{y}(t)]_{t=0}^\infty : (\mathbf{x}(t),\mathbf{y}(t)) \text{ admissible}\},$$

and a generic element of $\Omega(0,\mathbf{x}_0)$ by $(\mathbf{x}(t),\mathbf{y}(t))$. The set $\Omega(0,\mathbf{x}_0)$ can be more explicitly written as

$$\begin{aligned}\Omega(0,\mathbf{x}_0) \equiv\ & \{[\mathbf{x}(t),\mathbf{y}(t)]_{t=0}^\infty : \mathbf{y}(t) \in \mathcal{Y}(t),. \\ & \mathbf{x}(t) = \int_0^t G(s,\mathbf{x}(s),\mathbf{y}(s))\,ds + \mathbf{x}_0 \in \mathcal{X}(t) \text{ for all } t \in \mathbb{R}_+, \text{ and } \lim_{t\to\infty}\mathbf{x}(t) \geq \mathbf{x}_1\}\end{aligned}$$

THEOREM 7.15. ***(Existence of Solutions)*** *Consider the maximization of (7.73) subject to (7.74) and (7.75). Suppose that*

(1) $\mathcal{Y}(t) : \mathbb{R}_+ \rightrightarrows \mathbb{R}^{K_y}$ *and* $\mathcal{X}(t) : \mathbb{R}_+ \rightrightarrows \mathbb{R}^{K_x}$ *are compact-valued, closed-valued and upper hemi-continuous correspondences.*

(2) f *and* G *are continuous on* M.

(3) *The sets* $\Omega(0,\mathbf{x}_0)$ *and*

$$Q(t,\mathbf{x}) \equiv \left\{(p,\mathbf{z}) \in \mathbb{R}\times\mathbb{R}^{K_x}\text{: } p \leq f(t,\mathbf{x},\mathbf{y}),\ \mathbf{z} = G(t,\mathbf{x},\mathbf{y}),\ \text{and } \mathbf{y} \in \mathcal{Y}(t)\right\}$$

are nonempty for $\mathbf{x}_0 \in \mathcal{X}(t)$ *and for all* $(t,\mathbf{x}) \in M'$*. Moreover, the correspondence* $Q(t,\mathbf{x})$ *is closed and convex-valued and upper hemi-continuous for all all* $(t,\mathbf{x}) \in M'$*.*

(4) *For any interval* $[t_1, t_1+\delta]$ *and any real number* $\varepsilon > 0$*, there exists a continuous function* $\Phi_{t_1\delta\varepsilon}(t)$ *such that for any* $T \in [0,\infty]$*,* $\int_0^T \Phi_{t_1\delta\varepsilon}(t)\,dt$ *exists and is less than or equal to* $\Phi < \infty$ *and*

$$\varepsilon f(t,\mathbf{x},\mathbf{y}) \leq \Phi_{t_1\delta\varepsilon}(t) - \|G(t,\mathbf{x},\mathbf{y})\|$$

for all $t \in [t_1, t_1+\delta]$ *and for all* $(t,\mathbf{x},\mathbf{y}) \in M$*.*

(5) *There exists a positive function* $\phi(t)$ *(i.e.,* $\phi(t) \geq 0$ *for all* t*) such that* $\int_0^\infty \phi(t)\,dt$ *exists and is less than or equal to* $\phi < \infty$ *and*

$$f(t,\mathbf{x},\mathbf{y}) \leq \phi(t)$$

for all $(t,\mathbf{x},\mathbf{y}) \in M$*.*

Under hypotheses 1-5, there exists an admissible pair $(\mathbf{x}(t),\mathbf{y}(t)) \in \Omega(0,\mathbf{x}_0)$ *that is a solution to the maximization problem (7.73) subject to (7.74) and (7.75). That is,* $W(\mathbf{x}(t),\mathbf{y}(t)) = \bar{W} \geq W(\mathbf{x}'(t),\mathbf{y}'(t))$ *for any* $(\mathbf{x}'(t),\mathbf{y}'(t)) \in \Omega(0,\mathbf{x}_0)$*.*

Proof. **(Sketch)** From hypothesis (5),

$$\int_0^\infty f(t,\mathbf{x}(t),\mathbf{y}(t))\,dt \leq \int_0^\infty \phi(t)\,dt \leq \phi < \infty$$

for all $(t,\mathbf{x}(t),\mathbf{y}(t)) \in M$. Since $\Omega(0,\mathbf{x}_0)$ is nonempty,

$$\bar{W} = \sup_{(\mathbf{x}(t),\mathbf{y}(t))\in\Omega} \int_0^\infty f(t,\mathbf{x}(t),\mathbf{y}(t))\,dt \leq \phi < \infty.$$

For an admissible $(\mathbf{x}^n(t),\mathbf{y}^n(t))$, define

$$\begin{aligned} w^n(t) &\equiv \int_0^t f(s,\mathbf{x}^n(s),\mathbf{y}^n(s))\,ds \\ w_+(t) &\equiv \int_0^t \phi(s)\,ds \in [0,\phi] \\ w_-^n(t) &\equiv -\int_0^t [\phi(s) - f(s,\mathbf{x}^n(s),\mathbf{y}^n(s))]\,ds, \end{aligned}$$

where the inequality that $0 \leq w_+(t) \leq \phi$ follows from hypothesis (5). Moreover, again from hypothesis (5), $\phi(t) - f(t,\mathbf{x}^n(t),\mathbf{y}^n(t)) \geq 0$, so $w_-^n(t)$ is negative and nonincreasing. Also, evidently,

$$w_-^n(t) + w_+(t) = w^n(t).$$

Since $w_-^n(t)$ is nonincreasing, it converges on the extended real line, so that

$$\lim_{t\to\infty} w_-^n(t) = -W_-^n,$$

where W_{-}^{n} is bounded above (by $\bar{W} < +\infty$, since $\phi(t) - f(t, \mathbf{x}^{n}(t), \mathbf{y}^{n}(t)) \geq 0$ and $\int_{0}^{t} \phi(t)\, dt \leq \phi < \infty$) and is also negative, since $w_{-}^{n}(t) \leq 0$ for all t. Then, we can write

$$\begin{aligned}\lim_{t\to\infty} w_{-}^{n}(t) &= \lim_{t\to\infty} \left[w_{-}^{n}(t) + w_{+}(t) - w_{+}(t)\right] \\ &= \lim_{t\to\infty} \left[w^{n}(t) - w_{+}(t)\right] \\ &= -W_{-}^{n}\end{aligned}$$

The last equality ensures that $\lim_{t\to\infty} w^{n}(t)$ converges (on the extended real line) to W^{n}, which is less than $\phi - W_{-}^{n} < \infty$. Therefore, by the definition of the supremum, the sequence $\{(\mathbf{x}^{n}(t), \mathbf{y}^{n}(t))\}_{n=1}^{\infty}$ in Ω can be chosen such that

$$\bar{W} - \frac{K}{n} \leq \bar{W}^{n} \leq \bar{W}$$

for some real number $K > 0$. Consequently, $\lim_{n\to\infty} \bar{W}^{n}$ must exist and must be equal to $\bar{W}$.

To complete the proof, we must show that $\{(\mathbf{x}^{n}(t), \mathbf{y}^{n}(t))\}_{n=1}^{\infty}$ converge to an admissible pair. This step has four parts. We must show: (i) $\mathbf{x}^{n}(t) \to \mathbf{x}(t)$ such that $\mathbf{x}(t) \in \mathcal{X}(t)$; (ii) $\mathbf{x}(t)$ is absolutely continuous; (iii) whenever it exists, $\dot{\mathbf{x}}(t)$ satisfies (7.74); and (iv) $\mathbf{y}^{n}(t) \to \mathbf{y}(t) \in \mathcal{Y}(t)$, where $\mathbf{y}(t)$ is (Lebesgue) measurable.

Since $\mathcal{Y}(t)$ and $\mathcal{X}(t)$ are closed-valued (hypothesis (1)), M is closed. This also implies that M' is closed. Since any interval $[t_1, t_1 + \delta]$ is closed and compact and$\mathcal{X}(t)$ is compact for each t, $M_{t_1\delta} \equiv M' \cap \left([t_1, t_1 + \delta] \times \mathbb{R}^{K_x}\right)$ is compact (Lemma A.2 in Appendix Chapter A) and thus any sequence of (vector) functions $\mathbf{x}^{n}(t) \in M_{t_1\delta}$ is uniformly bounded. We next show that any sequence of continuous functions $\{\mathbf{x}^{n}(t)\}_{n=1}^{\infty}$ in $M_{t_1\delta}$ is equicontinuous (Definition A.30). Let $\varepsilon > 0$. Since $\int_{0}^{T} \Phi_{t_1\delta\varepsilon}(t)\, dt$ exists for any T and $\Phi_{t_1\delta\varepsilon}(t)$ is continuous (hypothesis (4)), there exists $\chi > 0$ such that

$$\int_{t_1}^{t_1+\delta} \Phi_{t_1\delta\varepsilon}(t)\, dt < \frac{\chi}{2}.$$

Using hypothesis (4) again, and setting

$$\varepsilon \equiv \frac{\chi}{2\left(\Phi - \bar{W}\right)} > 0$$

for each $n = 1, 2, ...$, we have that

$$\begin{aligned}\int_{t_1}^{t_1+\delta} \|\dot{\mathbf{x}}^{n}(t)\|\, dt &= \int_{t_1}^{t_1+\delta} \|\mathbf{G}(t, \mathbf{x}^{n}(t), \mathbf{y}^{n}(t))\|\, dt \\ &\leq \int_{t_1}^{t_1+\delta} \left[\Phi_{t_1\delta\varepsilon}(t) - \varepsilon f(t, \mathbf{x}^{n}(t), \mathbf{y}^{n}(t))\right] dt \\ &\leq \int_{t_1}^{t_1+\delta} \Phi_{t_1\delta\varepsilon}(t)\, d + \varepsilon \int_{0}^{\infty} \left[\phi(t) - f(t, \mathbf{x}^{n}(t), \mathbf{y}^{n}(t))\right] dt \\ &\leq \frac{\chi}{2} + \varepsilon\left(\Phi - \bar{W}\right) \\ &\leq \chi.\end{aligned}$$

Since this is true for each $n = 1, 2, ...$ and each interval $[t_1, t_1 + \delta]$, the vector functions $\mathbf{x}^n(t)$ are equicontinuous on each $[t_1, t_1 + \delta]$. Now the corollary to the Arzela-Ascoli Theorem (Corollary A.2 in Appendix Chapter A) implies that there exists a subsequence $\{\mathbf{x}^{n_k}(t)\}$ of $\{\mathbf{x}^n(t)\}$ such that $\mathbf{x}^{n_k}(t) \to \mathbf{x}(t)$ and that $\mathbf{x}(t)$ is absolutely continuous on $[t_1, t_1 + \delta]$. Since $[0, \infty)$ is a countable union of intervals of the form $[t_1, t_1 + \delta]$, repeated application of the preceding argument for countably many times implies that there exists a subsequence $\{\mathbf{x}^{n_{k'}}(t)\}$ of $\{\mathbf{x}^n(t)\}$ such that $\mathbf{x}^{n_{k'}}(t) \to \mathbf{x}(t)$ on $[0, \infty)$. Definition A.26 then implies that $\mathbf{x}(t)$ is absolutely continuous on $[0, \infty)$, completing the proof of parts (i) and (ii) above.

For (iii), we need to show that whenever $\dot{\mathbf{x}}(t)$ exists, there is $a(t)$ such that $(a(t), \dot{\mathbf{x}}(t)) \in Q(t, \mathbf{x}(t))$ for all t. For an arbitrary real number $\varepsilon > 0$, consider $Q_\varepsilon(t, \mathbf{x})$ be the ε neighborhood of the set $Q(t, \mathbf{x})$, that is, $Q_\varepsilon(t, \mathbf{x}) = \{(p, \mathbf{z}) : d((p, \mathbf{z}), Q(t, \mathbf{x})) < \varepsilon\}$, where d is the Euclidean metric. Since $Q(t, \mathbf{x})$ is convex-valued (hypothesis (2)), so is $Q_\varepsilon(t, \mathbf{x})$. By definition of admissibility, for the sequence $\{(\mathbf{x}^n(t), \mathbf{y}^n(t))\}_{n=1}^{\infty}$ considered above, $(p^n(t), \dot{\mathbf{x}}^n(t)) \in Q(t, \mathbf{x}^n(t))$ for all t (where $p^n(t) = f(t, \mathbf{x}^n(t), \mathbf{y}^n(t))$) And since $Q(t, \mathbf{x}^n(t)) \subset Q_\varepsilon(t, \mathbf{x}^n(t))$ for each t, $(p^n(t), \dot{\mathbf{x}}^n(t)) \in Q_\varepsilon(t, \mathbf{x}^n(t))$. Moreover, since $\varepsilon > 0$, there exists $h > 0$ such that $(p^n(t + \eta), \dot{\mathbf{x}}^n(t + \eta)) \in Q_\varepsilon(t, \mathbf{x}^n(t))$ for all $\eta \leq h$ and for each $n = 1, 2....$ Integrating (7.74) between t and $t + h$ for an admissible pair $(\mathbf{x}^n(t), \mathbf{y}^n(t))$, we obtain

$$\mathbf{x}^n(t + h) = \int_t^{t+h} G(s, \mathbf{x}^n(s), \mathbf{y}^n(s))\, ds + \mathbf{x}^n(t),$$

and therefore

$$\frac{\mathbf{x}^n(t + h) - \mathbf{x}^n(t)}{h} = \frac{\int_t^{t+h} \dot{\mathbf{x}}^n(s)\, ds}{h}$$

for each $n = 1, 2....$ Since $Q_\varepsilon(t, \mathbf{x}^n(t))$ is convex-valued and $\dot{\mathbf{x}}^n(s) \in Q_\varepsilon(t, \mathbf{x}^n(t))$ for all $s \in [t, t + h]$, $\left(\int_t^{t+h} \dot{\mathbf{x}}^n(s)\, ds\right)/h \in \overline{Q_\varepsilon(t, \mathbf{x}^n(t))}$ (recall that $\overline{Q_\varepsilon(t, \mathbf{x}^n(t))}$ is the closure of $Q_\varepsilon(t, \mathbf{x}^n(t))$). By hypothesis (3), $Q_\varepsilon(t, \mathbf{x}^n(t))$ is closed-valued and upper hemi-continuous and thus $\overline{Q_\varepsilon(t, \mathbf{x}^n(t))}$ is also closed-valued and upper hemi-continuous. In view of this, taking the limit $n \to \infty$ yields

$$\frac{\mathbf{x}(t + h) - \mathbf{x}(t)}{h} \in \overline{Q_\varepsilon(t, \mathbf{x}(t))}$$

Now taking the limit as $h \to 0$ and using the hypothesis that $\dot{\mathbf{x}}(t)$ exists, we have

$$\dot{\mathbf{x}}(t) \in \overline{Q_\varepsilon(t, \mathbf{x}(t))}.$$

Since $\varepsilon > 0$ and t were arbitrary, we conclude that

$$\dot{\mathbf{x}}(t) \in \overline{Q(t, \mathbf{x}(t))} = Q(t, \mathbf{x}(t)),$$

for all t, completing the proof of part (iii).

For part (iv), note that for each $n = 1, 2, ...$, $(\mathbf{x}^n(t), \mathbf{y}^n(t))$ is admissible, thus $\mathbf{y}^n(t)$ must be measurable in t. Moreover, since $\mathcal{Y}(t)$ is compact, so by Helly's Selection Theorem (e.g., Kolmogorov and Fomin, 1970, Theorem 5, p. 372), there exists a subsequence of $\{\mathbf{y}^n(t)\}, \{\mathbf{y}^{n_k}(t)\}$ such that $\mathbf{y}^{n_k}(t) \to \mathbf{y}(t)$. Since $\mathbf{x}(t)$ is absolutely continuous, it is differentiable almost everywhere (Fact A.16 in Appendix Chapter A). Therefore, $\dot{\mathbf{x}}(t)$ exists almost everywhere and is thus measurable. Moreover, $\mathbf{x}(t)$ is absolutely continuous and thus

measurable. Next let $D = \{(t, \mathbf{y}(t)) \in \mathbb{R}_+ \times \mathcal{Y}(t) : \dot{\mathbf{x}}(t) = G(t, \mathbf{x}(t), \mathbf{y}(t))\}$. Since $\mathbb{R}_+$ is a countable union of compact intervals and $\mathcal{Y}(t)$ is closed, D can be expressed as a countable union of compact sets. Since G is continuous in all of its arguments on D (hypothesis (2)) and $\dot{\mathbf{x}}(t) = G(t, \mathbf{x}(t), \mathbf{y}(t))$ is measurable for all $t \in \mathbb{R}_+$, $\mathbf{y}(t)$ is also measurable for all $t \in \mathbb{R}_+$. □

Theorem 7.15 establishes the existence of solutions to continuous-time maximization problems under relatively mild assumptions (at least from an economic point of view). Hypotheses (1) and (2) are the usual compactness and continuity assumption, necessary in finite-dimensional and discrete-time optimization problems as well (recall, for example, Theorem A.9 in Appendix Chapter A). Hypotheses (4) and (5) are somewhat stronger versions of the boundedness assumptions that are also necessary in finite-dimensional and discrete-time optimization problems (these are sometimes referred to as "growth conditions" since they restrict the rate at which the objective function can grow). Hypothesis (3) is rather unusual. Such a convexity assumption is not necessary in finite-dimensional and discrete-time optimization problems. However, in continuous-time problems this assumption cannot be dispensed with (see Exercise 7.24).

While Theorem 7.15 establishes the existence of solutions, providing sufficient conditions for the solution (the control $\mathbf{y}(t)$) to be continuous turns out to be much harder. Without knowing that the solution is continuous, the necessary conditions in Theorems 7.9 and/or 7.13 may appear incomplete. Although there is some truth to this, recall that both Theorems 7.9 and 7.13 provided necessary conditions only for all t except those at which $\mathbf{y}(t)$ is discontinuous. Therefore, these theorems remain valid as stated. Moreover, most economic problems possess enough structure to ensure that $\mathbf{y}(t)$ is continuous and $\mathbf{x}(t)$ is smooth. For example, in many economic problems Inada-type conditions ensure that optimal controls remain within the interior of the feasible set and concavity rules out discontinuous controls. However, instead of using these conditions to prove existence of continuous solutions, the indirect approach outlined above turns out to be more useful. In particular, in most economic problems the following approach works: (i) characterize a solution that satisfies the necessary conditions in Theorem 7.9 or in Theorem 7.13 (provided that such a solution exists); and (ii) verify the sufficiency conditions in one of Theorems 7.11 or 7.14. This approach, when it works, guarantees both the existence of a solution and continuity of the control (and smoothness of the state variable). Since all the problems in this book satisfy these conditions, this approach will be the one I will adopt in the remainder.

Let us next discuss how the conditions of Theorem 7.15 can be verified in the context of the optimal growth example, which will be discussed in greater detail in the next section.

Example 7.4. Consider the problem

$$\max_{[k(t),c(t)]_{t=0}^{\infty}} \int_0^{\infty} \exp(-\rho t)\, u(c(t))\, dt,$$

subject to

$$\dot{k}(t) = f(k(t)) - \delta k(t) - c(t),$$

$k(t) \geq 0$ for all t, and $k(0) > 0$. Recall that $u : R_+ \to R$ is strictly increasing, continuously differentiable and strictly concave, while $f(\cdot)$ satisfies our basic assumptions, Assumptions 1 and 2 (this also implies that $c(t) \geq 0$ for all t). Let us now check that hypotheses (1)-(5) are satisfied. First, note that in view of the Inada conditions in Assumption 2, there exists $\bar{k} > 0$ such that $f(\bar{k}) = \delta \bar{k}$, so that for any $k > \bar{k}$, $\dot{k}$ would necessarily be strictly negative (even if $c = 0$). Thus we can restrict attention to $k(t) \in \left[0, \vec{k}\right]$, where $\vec{k} \equiv \max\left\{\bar{k}, k(0)\right\}$. Thus, $k(t) \in \mathcal{X}(t) \equiv \left[0, \vec{k}\right]$ with $\mathcal{X}(t)$ compact and thus closed-valued, and also continuous in t. Similarly, $c(t) \in \left[0, f\left(\vec{k}\right)\right]$, thus $c(t) \in \mathcal{Y}(t) \equiv \left[0, f\left(\vec{k}\right)\right]$, so that $\mathcal{Y}(t)$ is compact-valued and continuous in t. Therefore, hypothesis (1) is verified. Both u and f are continuous, thus hypothesis (2) is verified as well. The correspondence $Q(t, \mathbf{x})$ takes the form

$$Q(t,k) = \left\{(p,z) \in \mathbb{R} \times \mathbb{R}\text{: } p \leq \exp(-\rho t)\, u(c),\ z = f(k) - \delta k - c \text{ and } c \in \left[0, f\left(\vec{k}\right)\right]\right\},$$

which is clearly closed- and convex-valued and continuous, and both $Q(t, \mathbf{x})$ and $\Omega(\mathbf{x}_0)$ are nonempty, thus hypothesis (3) is also satisfied.

In view of the fact that $c(t) \in \left[0, f\left(\vec{k}\right)\right]$, we have

$$\exp(-\rho t)\, u(t) \leq \exp(-\rho t) \max\left\{0, u\left(f\left(\vec{k}\right)\right)\right\},$$

so that hypothesis (5) is satisfied as well. Finally, using once more the fact that $c(t) \in \left[0, f\left(\vec{k}\right)\right]$ and $k(t) \in \left[0, \vec{k}\right]$, we have that for any interval $[t_1, t_1 + \delta]$ and any $\varepsilon > 0$, let

$$\Phi(t) \equiv \exp(-\rho t)\left[u\left(\vec{k}\right) + \max\left\{\vec{k}, \left[f\left(\vec{k}\right) - \delta \vec{k}\right]\right\}\right].$$

Then, $\varepsilon u(c) + |f(k) - \delta k - c| \leq \Phi(t)$ for any $t \in [t_1, t_1 + \delta]$ and

$$\int_0^\infty \exp(-\rho t)\, \Phi(t) \leq \frac{u\left(\vec{k}\right) + \max\left\{\vec{k}, \left[f\left(\vec{k}\right) - \delta \vec{k}\right]\right\}}{\rho} < \infty,$$

thus verifying hypothesis (4). Theorem 7.15 therefore guarantees the existence of solutions in this problem. In the next section, we will see that Theorem 7.14 will in fact ensure the existence of a continuous solution to this problem.

I next briefly present conditions under which the value function $V(t, x)$ is concave and differentiable. Recall the definition of $V(t, x)$ in (7.31), which was done for the case of a one-dimensional state variable. Let us consider the more general problem of maximizing (7.73) subject to (7.74) and (7.75). Presuming that the solution exists, the value function is defined as

$$V(t_0, \mathbf{x}_0) = \max_{[\mathbf{x}(t), \mathbf{y}(t)]_{t=t_0}^{\infty} \in \Omega(t_0, \mathbf{x}_0)} \int_{t_0}^{\infty} f(t, \mathbf{x}(t), \mathbf{y}(t))\, dt.$$

Notice that all of the constraints are already incorporated into the set $\Omega(t_0, \mathbf{x}_0)$.

The next theorem provides sufficient conditions for the value function $V(t,\mathbf{x})$ to be concave in $\mathbf{x}$. The theorem is simple, though in many problems ensuring that the constraint set, $\Omega(t,\mathbf{x})$, is convex may be difficult. In particular, $\Omega(t,\mathbf{x})$ is convex if, whenever $[\mathbf{x}(t),\mathbf{y}(t)]_{t=t_0}^{\infty} \in \Omega(t,\mathbf{x})$ and $[\mathbf{x}'(t),\mathbf{y}'(t)]_{t=t_0}^{\infty} \in \Omega(t,\mathbf{x}')$, then $[\mathbf{x}_\alpha(t),\mathbf{y}_\alpha(t)]_{t=t_0}^{\infty} \in \Omega(t,\mathbf{x}_\alpha)$ for any $\alpha \in [0,1]$, where $\mathbf{x}_\alpha(t) \equiv \alpha\mathbf{x}(t) + (1-\alpha)\mathbf{x}'(t)$, $\mathbf{y}_\alpha(t) \equiv \alpha\mathbf{y}(t) + (1-\alpha)\mathbf{y}'(t)$, and $\mathbf{x}_\alpha \equiv \alpha\mathbf{x} + (1-\alpha)\mathbf{x}'$. For example, linear systems of differential equations of the form $\dot{\mathbf{x}}(t) = A(t) + B(t)\mathbf{x}(t) + C(t)\mathbf{y}(t)$ satisfy this convexity condition.

Theorem 7.16. ***(Concavity of the Value Function)*** *Consider the problem of maximizing (7.73) subject to (7.74) and (7.75), and suppose that the hypotheses of Theorem 7.15 are satisfied so that a solution to this maximization problem exists and let the value function starting with initial value $\mathbf{x}_0$ at time t_0 be $V(t_0,\mathbf{x}_0)$. Suppose also that $f(t,\mathbf{x}(t),\mathbf{y}(t))$ is jointly concave in $\mathbf{x}(t)$ and $\mathbf{y}(t)$ for all $t \in [t_0,\infty)$ and $\Omega(t_0,\mathbf{x})$ is convex for all $\mathbf{x} \in \mathcal{X}(t_0)$. Then, $V(t_0,\mathbf{x})$ is concave in $\mathbf{x}$.*

Proof. Suppose $[\mathbf{x}(t),\mathbf{y}(t)]_{t=t_0}^{\infty} \in \Omega(t,\mathbf{x})$ and $[\mathbf{x}'(t),\mathbf{y}'(t)]_{t=t_0}^{\infty} \in \Omega(t,\mathbf{x}')$ are optimal solutions. Then

$$V(t_0,\mathbf{x}) = \int_{t_0}^{\infty} f(t,\mathbf{x}(t),\mathbf{y}(t))\,dt \text{ and } V(t_0,\mathbf{x}') = \int_{t_0}^{\infty} f\left(t,\mathbf{x}'(t),\mathbf{y}'(t)\right)dt. \tag{7.76}$$

Now consider $V(t_0,\mathbf{x}_\alpha)$ with $\mathbf{x}_\alpha$ defined as above. By the hypothesis that $\Omega(t_0,\mathbf{x})$ is convex, $[\mathbf{x}_\alpha(t),\mathbf{y}_\alpha(t)]_{t=t_0}^{\infty} \in \Omega(t,\mathbf{x}_\alpha)$ with $\mathbf{x}_\alpha(t)$ and $\mathbf{y}_\alpha(t)$ defined as above. Therefore,

$$\begin{aligned}
V(t_0,\mathbf{x}_\alpha) &\geq \int_{t_0}^{\infty} f(t,\mathbf{x}_\alpha(t),\mathbf{y}_\alpha(t))\,dt \\
&= \int_{t_0}^{\infty} f\left(t,\alpha\mathbf{x}(t)+(1-\alpha)\mathbf{x}'(t),\alpha\mathbf{y}(t)+(1-\alpha)\mathbf{y}'(t)\right)dt \\
&\geq \alpha\int_{t_0}^{\infty} f(t,\mathbf{x}(t),\mathbf{y}(t))\,dt + (1-\alpha)\int_{t_0}^{\infty} f\left(t,\mathbf{x}'(t),\mathbf{y}'(t)\right)dt \\
&= \alpha V(t_0,\mathbf{x}) + (1-\alpha)V\left(t_0,\mathbf{x}'\right),
\end{aligned}$$

where the first line uses the fact that $[\mathbf{x}_\alpha(t),\mathbf{y}_\alpha(t)]_{t=t_0}^{\infty}$ is feasible, thus $V(t_0,\mathbf{x}_\alpha)$ must be at least as large as the returns from $[\mathbf{x}_\alpha(t),\mathbf{y}_\alpha(t)]_{t=t_0}^{\infty}$. The second line writes out $\mathbf{x}_\alpha(t)$ and $\mathbf{y}_\alpha(t)$ explicitly, the third line uses the joint concavity of f, and the fourth line uses (7.76). This establishes the concavity of $V(t_0,\mathbf{x})$. □

Theorem 7.17. ***(Differentiability of the Value Function)*** *Consider the problem of maximizing (7.73) subject to (7.74) and (7.75), and suppose that the hypotheses of Theorems 7.15 and 7.16 are satisfied so that a solution exists and $V(t_0,\mathbf{x})$ is concave. Suppose also that f and G are differentiable in $(\mathbf{x},\mathbf{y})$ for all t and that the optimal pair $(\hat{\mathbf{x}}(t),\hat{\mathbf{y}}(t))$ is such that there exists $\Delta t > 0$, with $\hat{\mathbf{x}}(t) \in Int\mathcal{X}(t)$ and $\hat{\mathbf{y}}(t) \in Int\mathcal{Y}(t)$ for $t \in [t_0,t_0+\Delta t]$. Then, $V(t_0,\mathbf{x})$ is differentiable in $\mathbf{x}$ at $\mathbf{x}_0$.*

Proof. Use the equivalent of Lemma 7.1 for multivariate problems to write

$$V(t_0,\mathbf{x}_0) = \int_{t_0}^{t_0+\Delta t} f(t,\hat{\mathbf{x}}(t),\hat{\mathbf{y}}(t))\,dt + V(t_0+\Delta t,\hat{\mathbf{x}}(t_0+\Delta t)).$$

Since $\hat{\mathbf{x}}(t) \in \text{Int}\mathcal{X}(t)$ and $\hat{\mathbf{y}}(t) \in \text{Int}\mathcal{Y}(t)$ for $t \in [t_0, t_0 + \Delta t]$, for $\varepsilon > 0$ and sufficiently small, there exists a neighborhood $\mathcal{N}_\varepsilon(\mathbf{x}_0)$ of $\mathbf{x}_0$ such that for all $\mathbf{x} \in \mathcal{N}_\varepsilon(\mathbf{x}_0)$, there exists $[\mathbf{x}'(t), \mathbf{y}'(t)]_{t=t_0}^{\infty} \in \Omega(t, \mathbf{x})$ such that $\mathbf{x}'(t + \Delta t) = \hat{\mathbf{x}}(t + \Delta t)$ and $[\mathbf{x}'(t), \mathbf{y}'(t)]_{t=t_0+\Delta t}^{\infty} = [\hat{\mathbf{x}}(t), \hat{\mathbf{y}}(t)]_{t=t_{0+\Delta t}}^{\infty}$. Let the value from this potentially suboptimal admissible pair be

$$\bar{V}(t_0, \mathbf{x}) = \int_{t_0}^{t_0+\Delta t} f\left(t, \mathbf{x}'(t), \mathbf{y}'(t)\right) dt + V\left(t_0, \hat{\mathbf{x}}(t_0 + \Delta t)\right).$$

By definition, $\bar{V}(t_0, \mathbf{x})$ exists and satisfies $\bar{V}(t_0, \mathbf{x}) \leq V(t_0, \mathbf{x})$ for all $\mathbf{x} \in \mathcal{N}_\varepsilon(\mathbf{x}_0)$. Moreover, $\bar{V}(t_0, \mathbf{x})$ is differentiable in $\mathbf{x}$ by the fact that f is differentiable and $[\mathbf{x}'(t), \mathbf{y}'(t)]_{t=t_0}^{t_0+\Delta t}$ is a differentiable trajectory by Theorem B.14 in Appendix Chapter B. Moreover, by Theorem 7.16, $V(t_0, \mathbf{x}_0)$ is concave and appealing to the same result from convex analysis as in the proof of Theorem 6.6 in the previous chapter, the convex function $-V(t_0, \mathbf{x})$ is subdifferentiable and has a nonempty, closed and convex set of subgradients, such that for any subgradient $\mathbf{p}$ and for all $\mathbf{x} \in \mathcal{N}_\varepsilon(\mathbf{x}_0)$,

$$\mathbf{p} \cdot (\mathbf{x} - \mathbf{x}_0) \geq V(t_0, \mathbf{x}) - V(t_0, \mathbf{x}_0).$$

Moreover, since $\bar{V}(t_0, \mathbf{x}) \leq V(t_0, \mathbf{x})$ and $\bar{V}(t_0, \mathbf{x}_0) = V(t_0, \mathbf{x}_0)$,

$$\mathbf{p} \cdot (\mathbf{x} - \mathbf{x}_0) \geq V(t_0, \mathbf{x}) - V(t_0, \mathbf{x}_0) \geq \bar{V}(t_0, \mathbf{x}) - \bar{V}(t_0, \mathbf{x}_0)$$

for all $\mathbf{x} \in \mathcal{N}_\varepsilon(\mathbf{x}_0)$. Since $\bar{V}(t_0, \mathbf{x})$ is differentiable for all $\mathbf{x} \in \mathcal{N}_\varepsilon(\mathbf{x}_0)$, $\mathbf{p}$ must be uniquely defined, which implies that $V(t_0, \mathbf{x})$ is also differentiable. □

7.7. A First Look at Optimal Growth in Continuous Time

In this section, I briefly show how the main results developed so far, Theorems 7.13 and 7.14, can be applied to the problem of optimal growth. Recall that the optimal growth problem involves the maximization of the utility of the representative household. I do not provide a full treatment of this model here, since this is the topic of the next chapter. My objective is to illustrate how Theorems 7.13 and 7.14 can be applied in economic growth problems by using this canonical problem. I will show that checking the conditions for Theorem 7.13 requires some work, while verifying the conditions for Theorem 7.14 is much more straightforward.

Consider the neoclassical economy without any population growth and without any technological progress. In this case, the optimal growth problem in continuous time can be written as:

$$\max_{[k(t),c(t)]_{t=0}^{\infty}} \int_0^{\infty} \exp(-\rho t) u(c(t)) dt,$$

subject to

$$\dot{k}(t) = f(k(t)) - \delta k(t) - c(t),$$

$k(t) \geq 0$ for all t, and $k(0) > 0$. Let us suppose that the utility function $u(c)$ is defined as $u : \mathbb{R}_+ \setminus \{0\} \to \mathbb{R}$ or as $u : \mathbb{R}_+ \to \mathbb{R}$ and is strictly increasing, continuously differentiable and strictly concave (for example, when $u(c) = \log c$, the domain is $\mathbb{R}_+ \setminus \{0\}$). Suppose also that $f(\cdot)$ satisfies Assumptions 1 and 2 from Chapter 2.

Let us first set up the current-value Hamiltonian, which, in this case, does not directly depend on time and can be written as

$$\hat{H}(k,c,\mu) = u(c(t)) + \mu(t)[f(k(t)) - \delta k(t) - c(t)], \tag{7.77}$$

with state variable k, control variable c and current-value costate variable μ.

From Theorem 7.13, let us look for a path of capital stock and consumption per capita that satisfies the necessary conditions,

$$\hat{H}_c(k,c,\mu) = u'(c(t)) - \mu(t) = 0 \tag{7.78}$$

$$\hat{H}_k(k,c,\mu) = \mu(t)\left(f'(k(t)) - \delta\right) = \rho\mu(t) - \dot{\mu}(t). \tag{7.79}$$

In addition, we would like to use the stronger transversality condition (7.60), which here takes the form

$$\lim_{t\to\infty}\left[\exp(-\rho t)\,\mu(t)\,k(t)\right] = 0. \tag{7.80}$$

If we wish to show that this transversality condition is necessary (for an interior solution), we need to verify that Assumption 7.1 is satisfied. First, the objective function $u(c)$ is increasing in c and independent of k, so it is weakly monotone. The constraint function, $f(k) - \delta k - c$, is decreasing in c, but may be nonmonotone in k. However, without loss of any generality, attention can be restricted to $k(t) \in \left[0, \bar{k}\right]$, where $\bar{k}$ is defined such that $f'\left(\bar{k}\right) = \delta$. Since increasing the capital stock above this level would reduce output and thus consumption both today and in the future, and optimal control will never involve $k(t) \notin \left[0, \bar{k}\right]$ for any $t > 0$. When $k(t) \in \left[0, \bar{k}\right]$, the constraint function is also weakly monotone in k. This ensures that Assumption 7.1(i) is satisfied. Assumption 7.1(ii) is immediate, since the derivative of $f(k) - \delta k - c$ with respect to c is equal to -1 and is thus uniformly bounded away from 0. Assumption 7.1(iii) is not necessarily satisfied, however (for example, with $u(c) = \log c$, $\lim_{c\to 0} u'(c) = \infty$). Nevertheless, we could modify the problem by taking the utility function as $u : [\varepsilon, +\infty) \to \mathbb{R}$ for some $\varepsilon > 0$ sufficiently small. In this case, Assumption 7.1(iii) would also be satisfied. We can then verify that any solution to this modified problem involves $c(t) > \varepsilon$ for all t is also a solution to the original problem where the domain of the utility function is $\mathbb{R}_+ \setminus \{0\}$ (see as Exercise 7.25). Moreover, it can also be shown that for any optimal growth problem as formulated here, there exists $\varepsilon > 0$ such that $c(t) > \varepsilon$ for all t (see again Exercise 7.25). Thus there is no loss of generality in choosing the domain of the utility function as $[\varepsilon, +\infty)$. This then establishes that the necessary conditions in Theorem 7.13, including the strong version of the transversality condition, (7.60), can be used in the optimal growth problem. These give us (7.78)-(7.80).

However, a more direct line of attack would be to use Theorem 7.14. Recall that this theorem does *not* require Assumption 7.1. Therefore, there is no need to look for a modified problem with the restriction that $c(t) \in [\varepsilon, +\infty)$. We simply look for a candidate path that satisfies conditions (7.78)-(7.79) as well as the transversality condition (7.80). We then verify that (7.78) implies $\mu(t) > 0$ along this candidate path (since $u' > 0$ everywhere). Consequently, the current-value Hamiltonian given in (7.77) consists of the sum of two strictly

concave functions and is itself strictly concave (and since $\hat{H}(k,c,\mu)$ is strictly concave, so is $M(k,\mu) \equiv \max_c \hat{H}(k,c,\mu)$). Second, since $\mu(t) > 0$ and $k(t) > 0$ (by feasibility), any alternative path must satisfy $\lim_{t\to\infty} [\exp(-\rho t)\mu(t)k(t)] \geq 0$. Therefore, the conditions in Theorem 7.14 are satisfied. We can therefore conclude that the candidate solution that satisfies (7.78)-(7.80) gives a global maximum. Notice that in applying Theorem 7.14, we are using the equivalent of (7.60), (7.80), as a *sufficient* transversality condition under concavity (and thus we do not need to check Assumption 7.1).

Since an analysis of optimal growth in the neoclassical model is more relevant in the context of the next chapter, I do not provide further details here.

7.8. The q-Theory of Investment and Saddle-Path Stability

As another application of the methods developed in this chapter, I now consider the canonical model of investment under adjustment costs, also known as the q-theory of investment. This problem is not only useful as an application of optimal control techniques, but it is one of the basic models of standard macroeconomic theory. Moreover, I will use this model to illustrate the notion of *saddle-path stability*, which will play a crucial role in the analysis of optimal and equilibrium growth.

The economic problem is that of a price-taking firm trying to maximize the present discounted value of its profits. The only twist relative to the problems studied so far is that this firm is subject to "adjustment" costs when it changes its capital stock. In particular, let the capital stock of the firm be $k(t) \geq 0$ and suppose that the firm has access to a production function $f(k)$ that satisfies Assumptions 1 and 2. For simplicity, let us normalize the price of the output of the firm to 1 in terms of the final good at all dates. The firm is subject to adjustment costs captured by the function $\phi(i)$, which is strictly increasing, continuously differentiable and strictly convex, and satisfies $\phi(0) = \phi'(0) = 0$. This implies that in addition to the cost of purchasing investment goods (which given the normalization of price is equal to i for an amount of investment i), the firm incurs a cost of adjusting its production structure given by the convex function $\phi(i)$. In some models, the adjustment cost is taken to be a function of investment relative to capital—$\phi(i/k)$ instead of $\phi(i)$—but this makes no difference for our main focus. I also assume that installed capital depreciates at an exponential rate δ and that the firm maximizes its net present discounted earnings with a discount rate equal to the interest rate r, which is assumed to be constant.

The firm's problem can be written as

$$\max_{[k(t),i(t)]_{t=0}^{\infty}} \int_0^{\infty} \exp(-rt)\left[f(k(t)) - i(t) - \phi(i(t))\right] dt$$

subject to

$$\dot{k}(t) = i(t) - \delta k(t) \tag{7.81}$$

and $k(t) \geq 0$, with $k(0) > 0$ given. Notice that $\phi(i)$ does not contribute to capital accumulation; it is simply a cost. Moreover, since ϕ is strictly convex, it implies that it is not

optimal for the firm to make "large" adjustments. Therefore it will act as a force towards a smoother time path of investment.

To characterize the optimal investment plan of the firm, let us use the same strategy as in the previous section. In particular, let us write the current-value Hamiltonian:

$$\hat{H}(k,i,q) \equiv [f(k(t)) - i(t) - \phi(i(t))] + q(t)[i(t) - \delta k(t)],$$

where I used $q(t)$ instead of the familiar $\mu(t)$ for the costate variable, for reasons that will be apparent soon.

The necessary conditions for an interior solution to this problem, including the transversality condition, can be written as

$$\begin{aligned} \hat{H}_i(k,i,q) &= -1 - \phi'(i(t)) + q(t) = 0 \\ \hat{H}_k(k,i,q) &= f'(k(t)) - \delta q(t) = rq(t) - \dot{q}(t) \\ \lim_{t\to\infty} \exp(-rt)\, q(t)\, k(t) &= 0, \end{aligned} \tag{7.82}$$

where I suppressed the "^" to reduce notation. The first necessary condition implies that

$$q(t) = 1 + \phi'(i(t)) \text{ for all } t. \tag{7.83}$$

Let us next check sufficiency. Since $q(t) > 0$ for all t, $\hat{H}$ is strictly concave. Moreover, since $k(t) \geq 0$ by feasibility, we also have $\lim_{t\to\infty} \exp(-rt)\, q(t)\, k(t) \geq 0$ for any feasible investment and capital stock process. Consequently, Theorem 7.14 applies and shows that a solution to (7.82) characterizes the unique profit-maximizing investment and capital stock process. Once again, it turns out that using Theorem 7.14 is both straightforward and powerful. In particular, the transversality condition corresponding to (7.60) is being used as a sufficient condition and there is no need to check Assumption 7.1.

Next, differentiating (7.83) with respect to time,

$$\dot{q}(t) = \phi''(i(t))\, \dot{i}(t). \tag{7.84}$$

Substituting this into the second necessary condition, we obtain the following law of motion for investment:

$$\dot{i}(t) = \frac{1}{\phi''(i(t))} \left[(r+\delta)\left(1 + \phi'(i(t))\right) - f'(k(t)) \right]. \tag{7.85}$$

A number of interesting economic features emerge from this equation. First, as $\phi''(i)$ tends to zero, it can be verified that $\dot{i}(t)$ diverges, meaning that investment jumps to a particular value. In other words, it can be shown that this value is such that the capital stock immediately reaches its steady-state value (see Exercise 7.28). This is intuitive. As $\phi''(i)$ tends to zero, $\phi(i)$ becomes linear. In this case, adjustment costs simply increase the cost of investment linearly and do not create any need for smoothing. In contrast, when $\phi''(i(t)) > 0$, there will be a motive for smoothing, $\dot{i}(t)$ will take a finite value, and investment will adjust slowly. Therefore, as claimed above, adjustment costs lead to a smoother path of investment.

The behavior of investment and capital stock can now be analyzed using the differential equations (7.81) and (7.85). First, it can be verified easily that there exists a unique steady-state solution with $k > 0$. This solution involves a level of capital stock k^* for the firm and investment just enough to replenish the depreciated capital, $i^* = \delta k^*$. This steady-state level of capital satisfies the first-order condition (corresponding to the right-hand side of (7.85) being equal to zero):

$$f'(k^*) = (r + \delta)\left(1 + \phi'(\delta k^*)\right).$$

This first-order condition differs from the standard "modified golden rule" condition, which requires the marginal product of capital to be equal to the interest rate plus the depreciation rate, because an additional cost of having a higher capital stock is that there will have to be more investment to replenish depreciated capital. This is captured by the term $\phi'(\delta k^*)$. Since ϕ is strictly convex and f is strictly concave and satisfies the Inada conditions (from Assumption 2), there exists a unique value of k^* that satisfies this condition.

The analysis of dynamics in this case requires somewhat different ideas than those used in the basic Solow growth model (cf., Theorems 2.4 and 2.5). In particular, instead of global stability in the k-i space, the correct concept is one of *saddle-path stability*. The reason for this is that instead of an initial value constraint, $i(0)$ is pinned down by a boundary condition at "infinity," that is, to satisfy the transversality condition,

$$\lim_{t \to \infty} \left[\exp(-rt)\, q(t)\, k(t)\right] = 0.$$

This implies that in the context of the current theory, with one state and one control variable, we should have a one-dimensional manifold (a curve) along which capital-investment pairs tend towards the steady state. This manifold is also referred to as the "stable arm". The initial value of investment, $i(0)$, will then be determined so that the economy starts along this curve. In fact, if any capital-investment pair (rather than only pairs along this curve) were to lead to the steady state, we would not know how to determine $i(0)$; in other words, there would be an "indeterminacy" of equilibria. Mathematically, rather than requiring all eigenvalues of the linearized system to be negative, what we require now is saddle-path stability, which involves the number of (strictly) negative eigenvalues to be the same as the number of state variables.

This notion of saddle path stability is central in most growth models. Let us now make these notions more precise by considering the following generalizations of Theorems 2.4 and 2.5 (see Appendix Chapter B):

Theorem 7.18. ***(Saddle-Path Stability in Linear Systems)*** *Consider the following linear differential equation system*

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{b} \tag{7.86}$$

with initial value $\mathbf{x}(0)$*, where* $\mathbf{x}(t) \in \mathbb{R}^n$ *for all* t *and* $\mathbf{A}$ *is an* $n \times n$ *matrix. Let* $\mathbf{x}^*$ *be the steady state of the system given by* $\mathbf{A}\mathbf{x}^* + \mathbf{b} = 0$*. Suppose that* $m \leq n$ *of the eigenvalues of* $\mathbf{A}$ *have negative real parts. Then, there exists an* m*-dimensional subspace* M *of* $\mathbb{R}^n$ *such*

that starting from any $\mathbf{x}(0) \in M$*, the differential equation (7.86) has a unique solution with* $\mathbf{x}(t) \to \mathbf{x}^*$*.*

Theorem 7.19. ***(Saddle-Path Stability in Nonlinear Systems)*** *Consider the following nonlinear autonomous differential equation*

$$\dot{\mathbf{x}}(t) = \mathbf{G}[\mathbf{x}(t)] \tag{7.87}$$

where $\mathbf{G} : \mathbb{R}^n \to \mathbb{R}^n$ *and suppose that* $\mathbf{G}$ *is continuously differentiable, with initial value* $\mathbf{x}(0)$*. Let* $\mathbf{x}^*$ *be a steady-state of this system, given by* $\mathbf{F}(\mathbf{x}^*) = 0$*. Define*

$$\mathbf{A} = D\mathbf{G}(\mathbf{x}^*),$$

and suppose that $m \leq n$ *of the eigenvalues of* $\mathbf{A}$ *have (strictly) negative real parts and the rest have (strictly) positive real parts. Then, there exists an open neighborhood of* $\mathbf{x}^*$*,* $\mathbf{B}(\mathbf{x}^*) \subset \mathbb{R}^n$ *and an* m*-dimensional manifold* $M \subset \mathbf{B}(\mathbf{x}^*)$ *such that starting from any* $\mathbf{x}(0) \in M$*, the differential equation (7.87) has a unique solution with* $\mathbf{x}(t) \to \mathbf{x}^*$*.*

Put differently, these two theorems state that when only a subset of the eigenvalues have negative real parts, a lower-dimensional subset of the original space leads to stable solutions. Fortunately, in this context this is exactly what we require, since $i(0)$ should adjust in order to place us on exactly such a lower-dimensional subset of the original space.

Armed with these theorems, we can now investigate the transitional dynamics in the q-theory of investment. To see that the equilibrium will tend to this steady-state level of capital stock, let us to plot (7.81) and (7.85) in the k-i space. This is done in Figure 7.1. The curve corresponding to $\dot{k} = 0$, (7.81), is upward sloping, since a greater level of capital stock requires more investment to replenish the depreciated capital. When we are above this curve, there is more investment than necessary for replenishment, so that $\dot{k} > 0$. When we are below this curve, then $\dot{k} < 0$. On the other hand, the curve corresponding to $\dot{i} = 0$, (7.85), can be nonmonotonic. Nevertheless, it is straightforward to verify that in the neighborhood of the steady-state it is downward sloping (see Exercise 7.28). When we are to the right of this curve, $f'(k)$ is lower, thus $\dot{i} > 0$. When we are to its left, $\dot{i} < 0$. The resulting phase diagram, together with the one-dimensional stable manifold, is shown in Figure 7.1 (see again Exercise 7.28 for a different proof).

We will next see that starting with an arbitrary level of capital stock, $k(0) > 0$, the unique optimal solution involves an initial level of investment $i(0) > 0$, followed by a smooth and monotonic approach to the steady-state investment level of δk^*. In particular, it can be shown easily that when $k(0) < k^*$, $i(0) > i^*$ and it monotonically decreases towards i^* (see Exercise 7.28). This is intuitive. Adjustment costs discourage large values of investment, thus the firm cannot adjust its capital stock to its steady-state level immediately. However, because of diminishing returns, the benefit of increasing the capital stock is greater when the level of capital stock is low. Therefore, at the beginning the firm is willing to incur greater adjustment costs in order to increase its capital stock and $i(0)$ is high. As capital

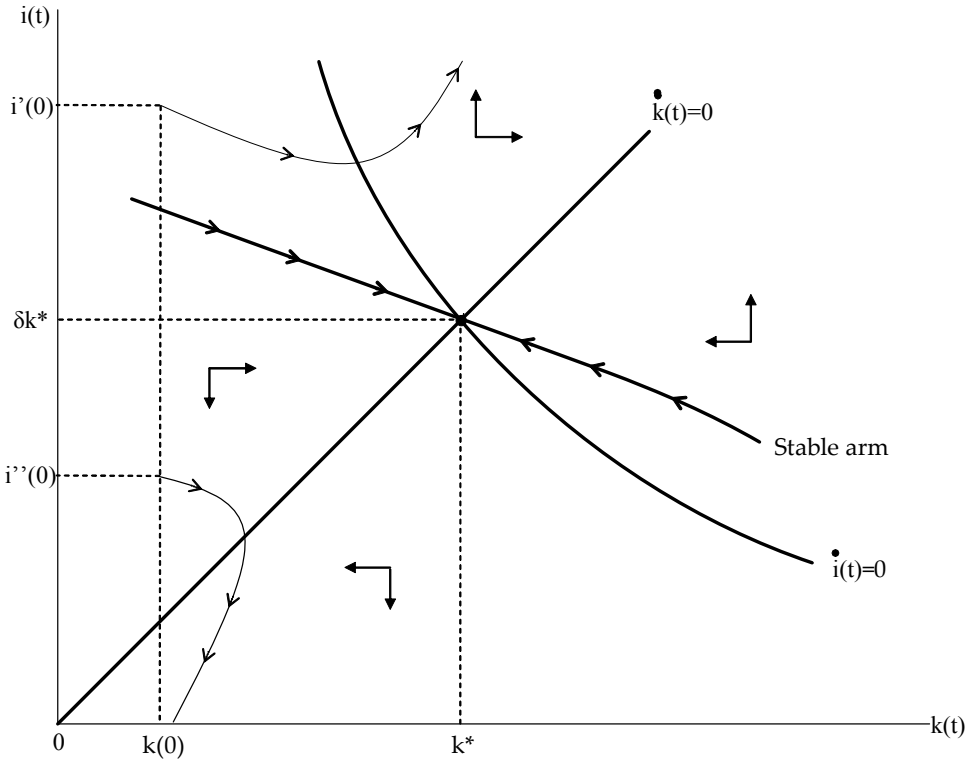

Figure 7.1. Dynamics of capital and investment in the q-theory.

accumulates and $k(t) > k(0)$, the benefit of boosting the capital stock declines and the firm also reduces investment towards the steady-state investment level.

There are two ways to see why the solution starting with $(k(0), i(0))$ and converging to (k^*, i^*) as shown in Figure 7.1 is the unique optimal solution. The first one, which is more rigorous and straightforward, is to use Theorem 7.14. As noted above, the conditions of this theorem hold in this problem. Thus we know that a path of capital and investment that satisfies the necessary conditions (that is, a path starting with $(k(0), i(0))$ and converging to (k^*, i^*)) is the unique optimal path. By implication, other paths, for example, those that start in $i'(0)$ or $i''(0)$ in Figure 7.1, cannot be optimal.

The second argument is common in the literature, though it is somewhat less complete than the argument just given. It involves showing that initial investment levels other than $i(0)$ would violate either the transversality condition or the first-order necessary conditions. Consider, for example, $i'(0) > i(0)$ as the initial level of investment. The phase diagram in Figure 7.1 makes it clear that starting from such a level of investment, the necessary conditions imply that $i(t)$ and $k(t)$ would tend to infinity. It can be verified that in this

case $q(t)k(t)$ would tend to infinity at a rate faster than r, thus violating the transversality condition, $\lim_{t\to\infty}\exp(-rt)q(t)k(t)=0$. To see this more explicitly, note that along a trajectory starting at $i'(0)$, $\dot{k}(t)/k(t)>0$, and thus

$$\begin{aligned}\frac{d(q(t)k(t))/dt}{q(t)k(t)} &\geq \frac{\dot{q}(t)}{q(t)}\\ &= \frac{\dot{i}(t)\phi''(i(t))}{1+\phi'(i(t))}\\ &= r+\delta-f'(k(t))/(1+\phi'(i(t))),\end{aligned}$$

where the second line uses (7.83) and (7.84), while the third line substitutes from (7.85). As $k(t)\to\infty$, $f'(k(t))\to 0$, implying that

$$\lim_{t\to\infty}\exp(-rt)q(t)k(t)\geq\lim_{t\to\infty}\exp(-rt)\exp((r+\delta)t)=\lim_{t\to\infty}\exp(\delta t)>0,$$

violating the transversality condition. In contrast, if we start with $i''(0)<i(0)$ as the initial level, $i(t)$ would tend to 0 in finite time (as shown by the fact that the trajectories hit the horizontal axis) and $k(t)$ would also tend towards zero (though not reaching it in finite time). After the time where $i(t)=0$, we also have $q(t)=1$ and thus $\dot{q}(t)=0$ (from (7.83)). Moreover, by the Inada conditions, as $k(t)\to 0$, $f'(k(t))\to\infty$. Consequently, after $i(t)$ reaches 0, the necessary condition $\dot{q}(t)=(r+\delta)q(t)-f'(k(t))$ is violated. This establishes that the unique optimal path involves investment starting at $i(0)$. The reason why this argument is not fully rigorous is that it is exploiting the necessary conditions for an interior continuous solution. However, when $i(0)=0$, we are no longer in the interior of the feasibility set for the control variable (here $\mathbb{R}_+$) and we cannot be sure that the necessary conditions still apply. Despite the fact that this argument is not complete, it is often used in many different contexts and especially in the analysis of competitive growth problems. Therefore, I will outline how a similar argument can be developed for the neoclassical growth model in the next chapter, though the main tool of analysis will again be the sufficiency result in Theorem 7.14.

Let us next turn to the "q-theory" aspects. James Tobin argued that the value of an extra unit of capital to the firm divided by its replacement cost is a measure of the "value of investment to the firm". In particular, when this ratio is high, the firm would like to invest more. In steady state, the firm will settle where this ratio is 1 or close to 1. In our formulation, the costate variable $q(t)$ measures Tobin's q. To see this, let us denote the current (maximized) value of the firm when it starts with a capital stock of $k(t)$ by $V(k(t))$. The same arguments as above imply that

$$V'(k(t))=q(t), \tag{7.88}$$

so that $q(t)$ measures exactly by how much one dollar increase in capital will raise the value of the firm.

In steady state, $\dot{q}(t)=0$, so that $q^*=f'(k^*)/(r+\delta)$, which is approximately equal to 1 when $\phi'(\delta k^*)$ is small. Nevertheless, out of steady state, $q(t)$ can be significantly greater than

this amount, signaling that there is need for greater investments. Therefore, in this model Tobin's q, or alternatively the costate variable $q(t)$, will signal when investment demand is high.

The q-theory of investment is one of the workhorse models of macroeconomics and finance, since proxies for Tobin's q can be constructed using stock market prices and book values of firms. When stock market prices are greater than book values, this corresponds to periods in which the firm in question has a high Tobin's q—meaning that the value of installed capital is greater than its replacement cost, which appears on the books. Nevertheless, whether this is a good approach in practice is intensely debated, in part because Tobin's q does not contain all the relevant information when there are irreversibilities or fixed costs of investment, and also perhaps more importantly, what is relevant in theory (and in practice) is the "marginal q," which corresponds to the marginal increase in value (as suggested by eq. (7.88)). However, in the data most measures compute "average q". The discrepancy between these two concepts can be large.

7.9. Taking Stock

This chapter has reviewed the basic tools of dynamic optimization in continuous time. By its nature, this has been a technical (and unfortunately somewhat dry) chapter. The material covered here may have been less familiar than the discrete-time optimization methods presented in the previous chapter. Part of the difficulty arises from the fact that optimization in continuous time is with respect to functions, even when the horizon is finite (rather than with respect to vectors or infinite sequences as in the discrete-time case). This introduces a range of complications and some technical difficulties, which are not of great interest in the context of economic applications. As a result, this chapter has provided an overview of the main results, with an emphasis on those that are most useful in economic applications, together with some of the proofs. These proofs are included to provide the reader with a sense of where the results come from and to enable them to develop a better feel for their intuition.

It is useful to recap the main approach developed in this chapter. Most of the problems in economic growth and macroeconomics require the use of discounted infinite-horizon optimal control methods. Theorem 7.13 provides necessary conditions for an interior continuous solution to such problems. Theorem 7.14 provides sufficient conditions, related to the concavity of the maximized Hamiltonian, for such a solution to be a global or unique global maximum (these conditions require the existence of a candidate solution, since they use information on the costate variable of this solution). Importantly, the conditions in Theorem 7.14 are more straightforward to verify than those in Theorem 7.13 (in particular, than Assumption 7.1). Therefore, the following strategy will be used in all cases in this book.

(1) We start with the necessary conditions in Theorem 7.13 to construct a *candidate solution*. This can be done even when Assumption 7.1 is not satisfied.

(2) Once a candidate path has been located, we then verify that the concavity conditions in Theorem 7.14 are satisfied. If they are, then we have located a path that is optimal. If, in addition, the maximized Hamiltonian is strictly concave, then this is the unique optimal solution.

It is also worth noting that while the basic ideas of optimal control may be a little less familiar than those of discrete-time dynamic programming, these methods are used in much of growth theory and in other areas of macroeconomics. Moreover, while some problems naturally lend themselves to analysis in discrete time, other problems become easier in continuous time. Some argue that this is indeed the case for growth theory. Regardless of whether one agrees with this assessment, it is important to have a good command of both discrete-time and continuous-time models in macroeconomics, since it should be the context and economic questions that dictate which type of model one should write down, not the force of habit. This motivated my choice of giving roughly equal weight to the two sets of techniques.

There is another reason for studying optimal control. The most powerful theorem in optimal control, Pontryagin's Maximum Principle, is as much an economic result as a mathematical result. As discussed above, the Maximum Principle has a very natural interpretation both in terms of maximizing flow returns plus the value of the stock, and also in terms of an asset value equation for the value of the maximization problem. These economic intuitions are not only useful in illustrating the essence of this mathematical technique, but they also provide a useful perspective on a large set of questions that involve the use of dynamic optimization techniques in macroeconomics, labor economics, finance and other fields.

This chapter also concludes our exposition of the "foundations" of growth theory, which comprised general equilibrium foundations of aggregative models as well as an introduction to mathematical tools necessary for dynamic economic analysis. I next turn to economically more substantive issues.

7.10. References and Literature

The main material covered in this chapter is the topic of many excellent applied mathematics. The purpose here has been to provide a review of the results that are most relevant for economists, together with simplified versions of the most important proofs. The first part of the chapter is closer to the calculus of variations theory, because it makes use of variational arguments combined with continuity properties. Nevertheless, most economists do not need to study the calculus of variations in detail, both because it has been superseded by optimal control theory and also because most of the natural applications of the calculus of variations are in physics and other natural sciences. The interested reader can look at Gelfand and Fomin (2000). Chiang (1992) provides a readable and simple introduction to the calculus of variations with economic examples.

The theory of optimal control was originally developed by Pontryagin et al. (1962). For this reason, the main necessary condition is also referred to as the Pontryagin's (Maximum) Principle. The type of problem considered here (and in economics more generally) is referred

to as the *Lagrange problem* in optimal control theory. The Maximum Principle is generally stated either for the somewhat simpler Meyer problem or the more general Bolza problem, though all of these problems are essentially equivalent, and when the problem is formulated in vector form, one can easily go back and forth between these different problems by simple transformations. A more modern approach, which underlies the necessary conditions used in infinite-horizon problems, is developed in Rockefeller (1971).

There are several books with varying levels of difficulty dealing with optimal control. Many of these books are not easy to read, but are also not entirely rigorous in their proofs. An excellent source that provides an advanced and complete treatment is Fleming and Rishel (1975). The first part of this book provides a complete (but rather different) proof of Pontryagin's Maximum Principle and various applications. This book also provides a number of theorems on existence and continuity of optimal controls, though for more specialized problems than those covered in Theorem 7.15 (and than those necessary for economic applications). The proof of existence of solutions in Section 7.6 above combines certain ideas from Baum's (1976) proof, which in turn extends Cesari's (1966) classic proof of existence of solutions to infinite-horizon problems, with part of the proof in Fleming and Rishel (1975, Chapter 3). In particular, the last part of Theorem 7.15, which established the measurability of control $\mathbf{y}(t)$, can be shown in greater detail using a similar line of argument to that in Fleming and Rishel (which involves the use of Lusin's Theorem). In the economics literature, existence theorems are provided in Magill (1981) and Romer (1986b), but under somewhat more restrictive conditions, and using a different method of proof.

A deeper understanding of sufficient conditions for existence of solution and the structure of necessary conditions can be gained from the excellent (but abstract and difficult) book by Luenberger (1969). The results in this book are general enough to cover both discrete-time and continuous-time dynamic optimization. This book also gives a very good sense of why maximization in function spaces is different from finite-dimensional maximization, and when such infinite-dimensional maximization problems may fail to have solutions.

The main theorems in the infinite-horizon case (Theorems 7.9, 7.13, 7.11 and 7.14) have been presented with the terminal value constraint $\lim_{t\to\infty} b(t)\, x(t) \geq x_1$. This is important, since the constraint on household assets in the competitive equilibrium of the neoclassical growth model (the no-Ponzi condition) takes this form. The standard results with terminal value constraints of the form $\lim_{t\to\infty} x(t) \geq x_1$ cannot be applied directly. Many authors seem to use the following reasoning: ignore the terminal value constraint, apply the Maximum Principle, and then if necessary, use the terminal value constraint at the end. While this procedure typically gives the "right" answer, it is not mathematically correct. The Maximum Principle cannot be applied in economic problems without a terminal value constraint, since in that case a solution typically fails to exist (see, for example, Exercise 8.2 in the next chapter). Therefore the application of the Maximum Principle to these problems is vacuous. A slight strengthening of the terminal value constraints of the Maximum Principle circumvents this problem.

Note also that in contrast to the standard practice in economic growth and macroeconomics, the emphasis here has been mostly on the sufficiency results for concave problems, in particular, on Theorem 7.14. This is because the standard form of the Maximum Principle only gives necessary conditions for interior continuous solutions. But it is not easy to verify that such a solution exists. Since most problems in economics are concavey, Theorem 7.14 or some of the other sufficiency results (for example, Theorems 7.5, 7.6, 7.8 or 7.11) are easy to apply and enable us to verify that a candidate solution that satisfies the Maximum Principle is indeed a solution and achieves the global optimum. It should also be noted that all of the sufficiency results here have been stated and proved assuming that the control function $y(t)$ (or $\mathbf{y}(t)$) is continuous. The logic of the proof is very similar when these functions are piecewise continuous and a formal proof along these lines is provided in Seierstad and Sydsaeter (1977).

Books that develop the theory of optimal control with economic applications may be more accessible for economists. Here the best reference is Seierstad and Sydsaeter (1987). This book is not as rigorous as Fleming and Rishel (1975) and in fact does not contain proofs for most results. Nevertheless, it does provide a number of useful results and is more interesting to read for economists. It also shows how the results can be applied to economic problems. Other references in economics are Kamien and Schwartz (1991) and Leonard and Van Long (1992). Another classic is Arrow and Kurz's (1970) book, which covers the same material and also presents rich economic insight on growth theory and related problems. This book also states and provides a proof of Arrow's sufficiency theorem, which also appears in Arrow (1968).

Two recent books on applications of optimal control in economics, Caputo (2005) and Weitzman (2003), are somewhat more accessible. My treatment of the sufficiency results here is very similar to Caputo (2005). Weitzman (2003) provides a lively discussion of the applications of the Maximum Principle, especially in the context of environmental economics and depletion of natural resources.

There is some confusion in the literature over the role of the transversality condition. The example provided in Section 7.4 shows that the stronger transversality condition, which is very useful in many applications, does not always hold. Halkin (1974) was the first to provide an example to show this. The weaker form of the transversality condition (7.49) was derived in Michel (1982). His results are similar to those in Theorem 7.12, though Theorem 7.12 is stated under assumptions that are more relevant in economic situations. Michel assumes that the objective function is nonnegative, which is violated by many of the common utility functions used in economic growth models, and also imposes an additional technical assumption that is not easy to verify. Michel (1982) also provides another set of sufficient conditions for the stronger transversality condition (7.60). More general (weaker) transversality conditions appropriate for economic models are presented in Benveniste and Scheinkman (1982) and Araujo and Scheinkman (1983). Theorem 7.14 here is stated under somewhat different, and easy to check, assumptions.

The original economic interpretation of the Maximum Principle appeared in Dorfman (1969). The interpretation here builds on the discussion by Dorfman, but also extends this based on the no-arbitrage interpretation of asset values in the Hamilton-Jacobi-Bellman equation. This interpretation of Hamilton-Jacobi-Bellman is well known in many areas of macroeconomics and labor economics, but is not often used to provide a general economic interpretation for the Maximum Principle. Weitzman (2003) also provides an economic interpretation for the Maximum Principle related to the Hamilton-Jacobi-Bellman equation.

The classic reference for exploitation of a non-renewable resource is Hotelling (1931). Weitzman (2003) provides a detailed treatment and a very insightful discussion. Dasgupta and Heal (1979) and Conrad (1999) are also useful references for applications of similar ideas to sustainability and environmental economics. Classic references on investment with costs of adjustment and the q-theory of investment include Eisner and Strotz (1963), Lucas (1967), Tobin (1969) and Hayashi (1982). Detailed treatments of the q-theory of investment can be found in any graduate-level economics textbook, for example, Blanchard and Fisher (1989) or Romer (1996), as well as in Dixit and Pindyck's (1994) book on investment under uncertainty and Caballero's (1999) survey. Caballero (1999) also includes a critique of the q-theory.

7.11. Exercises

Exercise 7.1. Consider the problem of maximizing (7.1) subject to (7.2) and (7.3) as in Section 7.1. Suppose that $\lambda(t)$ is defined by (7.10). Suppose also that given the pair $(\hat{x}(t), \hat{y}(t))$, there exists a time interval (t', t'') with $t' < t''$ such that

$$\dot{\lambda}(t) \neq -\left[f_x(t, \hat{x}(t), \hat{y}(t)) + \lambda(t) g_x(t, \hat{x}(t), \hat{y}(t))\right] \text{ for all } t \in \left(t', t''\right).$$

Prove that the pair $(\hat{x}(t), \hat{y}(t))$ could not attain the optimal value of (7.1). Using this argument establish that (7.10), (7.9) and $\lambda(t_1) = 0$ are necessary conditions for an interior solution $(\hat{x}(t), \hat{y}(t))$.

Exercise 7.2. * Prove that, given in optimal solution $\hat{x}(t), \hat{y}(t)$ to (7.1), the maximized Hamiltonian defined in (7.16) and evaluated at $\hat{x}(t)$, $M(t, \hat{x}(t), \lambda(t))$, is differentiable in x and satisfies $\dot{\lambda}(t) = -M_x(t, \hat{x}(t), \lambda(t))$ for all $t \in [0, t_1]$.

Exercise 7.3. The key equation of the calculus of variations is the Euler-Legrange equation, which characterizes the solution to the following problem (under similar regularity conditions to those of Theorem 7.2):

$$\max_{x(t)} \int_0^{t_1} F(t, x(t), \dot{x}(t))\, dx$$

subject to $x(t) = 0$. Suppose that F is differentiable in all of its arguments and an interior continuously differentiable solution exists. The so-called Euler-Legrange equation, which provides the necessary conditions for an optimal solution, is

$$\frac{\partial F(t, x(t), \dot{x}(t))}{\partial x(t)} - \frac{\partial^2 F(t, x(t), \dot{x}(t))}{\partial \dot{x}(t)\, \partial t} = 0.$$

Derive this equation from Theorem 7.2. [Hint: define $y(t) \equiv \dot{x}(t)$].

EXERCISE 7.4. This exercise asks you to use the Euler-Legrange equation derived in Exercise 7.3 to solve the canonical problem that motivated Euler and Legrange, that of finding the shortest distance between two points in a plane. In particular, consider a two dimensional plane and two points on this plane with coordinates (z_0, u_0) and (z_1, u_1). We would like to find the curve that has the shortest length that connects these two points. Such a curve can be represented by a function $x : \mathbb{R} \to \mathbb{R}$ such that $u = x(z)$, together with initial and terminal conditions $u_0 = x(z_0)$ and $u_1 = x(z_1)$. It is also natural to impose that this curve $u = x(z)$ be smooth, which corresponds to requiring that the solution be continuously differentiable so that $x'(z)$ exists.

To solve this problem, observe that the (arc) length along the curve x can be represented as

$$A[x(z)] \equiv \int_{z_1}^{z_2} \sqrt{1 + [x'(z)]^2} dz.$$

The problem is to minimize this object by choosing $x(z)$.

Now, without loss of any generality let us take $(z_0, u_0) = (0, 0)$ and let $t = z$ to transform the problem into a more familiar form, which becomes that of maximizing

$$-\int_0^{t_1} \sqrt{1 + [x'(t)]^2} dt.$$

Prove that the solution to this problem requires

$$\frac{d\left[x'(t)\left(1 + (x'(t))^2\right)\right]}{dt} = 0.$$

Show that this is only possible if $x''(t) = 0$, so that the shortest path between two points is a straight-line.

EXERCISE 7.5. Prove Theorem 7.2, in particular, paying attention to constructing feasible variations that ensure $x(t_1, \varepsilon) = x_1$ for all ε in some neighborhood of 0. What happens if there are no such feasible variations?

EXERCISE 7.6.
(1) Provide an expression for the initial level of consumption $c(0)$ as a function of $a(0)$, w, r and β in Example 7.1.
(2) What is the effect of an increase in $a(0)$ on the initial level of consumption $c(0)$? What is the effect on the consumption path?
(3) How would the consumption path change if instead of a constant level of labor earnings, w, the individual faced a time-varying labor income profile given by $[w(t)]_{t=0}^{1}$? Explain the reasoning for the answer in detail.

EXERCISE 7.7. Prove that if a function $\phi(x, y)$ is jointly [strictly] concave in (x, y), then $\Phi = \max_y \phi(x, y)$ is [strictly] concave in x.

EXERCISE 7.8. * Prove a version of Theorem 7.6 corresponding to Theorem 7.2. [Hint: instead of $\lambda(t_1) = 0$, the proof should exploit the fact that $x(1) = \hat{x}(1) = x_1$].

EXERCISE 7.9. Prove Theorem 7.3. [Hint: show that if $x(t_1) > x_1$ in the optimal solution, then $\lambda(t_1)$ must be equal to 0, but not necessarily so when $x(t_1) = x_1$].

EXERCISE 7.10. Consider the problem of maximizing (7.1) subject to (7.2) and (7.3), and the additional constraint that $y(t) \geq 0$ for all t (instead of $y(t) \in \text{Int}\mathcal{Y}(t)$). Assume as usual that f and g continuously differentiable. Show that the necessary conditions for optimality are: $x(0) = x_0$,

$$H_y(t, \hat{x}(t), \hat{y}(t), \lambda(t)) \leq 0 \text{ for all } t \in [0, t_1],$$

$$\dot{\lambda}(t) = -H_x(t, \hat{x}(t), \hat{y}(t), \lambda(t)) \text{ for all } t \in [0, t_1],$$

$$\dot{x}(t) = H_\lambda(t, \hat{x}(t), \hat{y}(t), \lambda(t)) \text{ for all } t \in [0, t_1],$$

and $\lambda(t_1) = 0$, with the Hamiltonian $H(t, x, y, \lambda)$ given by (7.12). Moreover, the Hamiltonian $H(t, x, y, \lambda)$ also satisfies the *Maximum Principle* that for all $t \in [0, t_1]$,

$$H(t, \hat{x}(t), \hat{y}(t), \lambda(t)) \geq H(t, \hat{x}(t), y, \lambda(t)) \text{ for all } y \text{ such that } y(t) \geq 0 \text{ for all } t.$$

EXERCISE 7.11. * Prove Theorem 7.7.

EXERCISE 7.12. Prove Theorem 7.8.

EXERCISE 7.13. Prove Theorem 7.11. [Hint: apply the proof of Theorem 7.6, except that when the expression with $\dot{\lambda}$ is integrated by parts, there is an additional term, which you should show increases the right-hand side of the inequality corresponding to (7.20)].

EXERCISE 7.14. Prove that in the discounted infinite-horizon optimal control problem considered in Theorem 7.13 conditions (7.56)-(7.58) are necessary. [Hint: use Theorem 7.9].

EXERCISE 7.15. Consider a finite-horizon continuous-time maximization problem, where the objective function is

$$W(x(t), y(t)) = \int_0^{t_1} f(t, x(t), y(t))\, dt$$

with $x(0) = x_0$ and $t_1 < \infty$, and the constraint equation is

$$\dot{x}(t) = g(t, x(t), y(t)).$$

Imagine that t_1 is also a choice variable.

(1) Show that $W(x(t), y(t))$ can be written as

$$W(x(t), y(t)) = \int_0^{t_1} \left[H(t, x(t), y(t)) + \dot{\lambda}(t) x(t) \right] dt - \lambda(t_1) x(t_1) + \lambda(0) x_0,$$

where $H(t, x, y) \equiv f(t, x(t), y(t)) + \lambda(t) g(t, x(t), y(t))$ is the Hamiltonian and $\lambda(t)$ is the costate variable.

(2) Now suppose that the pair $(\hat{x}(t), \hat{y}(t))$ together with terminal date $\hat{t}_1$ constitutes an optimal solution for this problem. Consider the following class of variations:

$$\begin{aligned} y(t, \varepsilon) &= \hat{y}(t) + \varepsilon \eta(t) \text{ for } t \in \left[0, \hat{t}_1\right] \text{ and } y(t, \varepsilon) = \hat{y}\left(\hat{t}_1\right) + \varepsilon \eta(t) \text{ for } t \in \left[\hat{t}_1, \hat{t}_1 + \varepsilon \Delta t\right], \\ t_1 &= \hat{t}_1 + \varepsilon \Delta t. \end{aligned}$$

Denote the corresponding path of the state variable by $x(t, \varepsilon)$. Evaluate $W(x(t, \varepsilon), y(t, \varepsilon))$ at this variation. Explain why this variation is feasible for ε small enough.

(3) Show that for a feasible variation,

$$\begin{aligned}\left.\frac{dW\left(x\left(t,\varepsilon\right),y\left(t,\varepsilon\right)\right)}{d\varepsilon}\right|_{\varepsilon=0} &= \int_0^{\hat{t}_1}\left[H_x\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right)\right)+\dot{\lambda}\left(t\right)\right]\frac{\partial x\left(t,\varepsilon\right)}{\partial\varepsilon}dt\\ &\quad+\int_0^{\hat{t}_1}H_y\left(t,\hat{x}\left(t\right),\hat{y}\left(t\right)\right)\eta\left(t\right)dt\\ &\quad+H\left(\hat{t}_1,\hat{x}\left(\hat{t}_1\right),\hat{y}\left(\hat{t}_1\right)\right)\Delta t-\lambda\left(\hat{t}_1\right)\frac{\partial x\left(\hat{t}_1,\varepsilon\right)}{\partial\varepsilon}.\end{aligned}$$

(4) Explain why the previous expression has to be equal to 0.

(5) What happens if you take limit $\hat{t}_1\to\infty$?

Exercise 7.16. Consider the discounted infinite-horizon problem, with $f\left(t,x\left(t\right),y\left(t\right)\right)=\exp\left(-\rho t\right)f\left(x\left(t\right),y\left(t\right)\right)$, and $g\left(t,x\left(t\right),y\left(t\right)\right)=g\left(x\left(t\right),y\left(t\right)\right)$. Prove that if an admissible pair $\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)_{t\geq0}$ is optimal starting at $t=0$ with initial condition $x\left(0\right)=x_0$, then $\left(\hat{x}\left(t\right),\hat{y}\left(t\right)\right)_{t\geq s}$ is also admissible and optimal for the problem starting at $t=s$ with initial condition $x\left(s\right)=x_0$.

Exercise 7.17. * Generalize the transversality condition (7.60) in Theorem 7.13, by strengthening Assumption 7.1, so that there exists $M<\infty$ such that $\left|f_y\left(t,x,y\right)\right|\leq M\exp\left(-\kappa t\right)$ for some $\kappa>0$. Verify that each of the step of the proof carefully goes through under this additional assumption.

Exercise 7.18. Consider the following discounted infinite-horizon maximization problem

$$\max\int_0^\infty\exp\left(-\rho t\right)\left[2y\left(t\right)^{1/2}+\frac{1}{2}x\left(t\right)^2\right]dt$$

subject to

$$\dot{x}\left(t\right)=-\rho x\left(t\right)y\left(t\right)$$

and $x\left(0\right)=1$.

(1) Show that this problem satisfies all the assumptions of Theorem 7.14.

(2) Set up at the current-value Hamiltonian and derive the necessary conditions, with the costate variable $\mu\left(t\right)$.

(3) Show that the following is an optimal solution $y\left(t\right)=1$, $x\left(t\right)=\exp\left(-\rho t\right)$, and $\mu\left(t\right)=\exp\left(\rho t\right)$ for all t.

(4) Show that this optimal solution violates the condition that $\lim_{t\to\infty}\exp\left(-\rho t\right)\mu\left(t\right)=0$, but satisfies (7.60).

Exercise 7.19. Consider the problem of consuming a non-renewable resource in Example 7.3. Show that the solution outlined their satisfies the transversality condition (7.49).

Exercise 7.20. Consider a modified version of the Hotelling resource extraction problem,

$$\max\int_0^\infty\exp\left(-\rho t\right)\log y\left(t\right)dt$$

subject to

$$\dot{x}\left(t\right)=-y\left(t\right)$$

and

$$x(0) = x_0 > 0 \text{ and } \lim_{t\to\infty} x(t) \geq x_1 \in (0, x_0),$$

with $\rho > 0$ and $y \in (0, +\infty)$.

(1) Show that this problem does not satisfy Assumption 7.1.

(2) Set up the current-value Hamiltonian $\hat{H}(x, y, \mu)$ with costate variable $\mu(t)$ and show that the Maximum Principle implies

$$\begin{aligned} \mu(t) &= \mu(0)\exp(\rho t) \\ y(t) &= \mu(0)^{-1}\exp(-\rho t) \end{aligned}$$

for some $\mu(0) > 0$.

(3) Show that the weaker transversality condition (7.59) is satisfied. [Hint: $\lim_{t\to\infty}\left[\exp(-\rho t)\hat{H}(x, y, \mu)\right] = \lim_{t\to\infty}[\exp(-\rho t)(-\log\mu(0) - \rho t - 1)] = 0$].

(4) Show that the stronger transversality condition (7.60) is not satisfied. [Hint: $\lim_{t\to\infty}[\exp(-\rho t)\mu(t)x(t)] = x_1 > 0$.]

(5) Explain why the stronger transfers out the condition does not hold here and how it should be modified. [Hint: try $\lim_{t\to\infty}[\exp(-\rho t)\mu(t)(x(t) - x_1)] = 0$].

Exercise 7.21. This exercise generalizes Example 7.2. Consider the following optimal growth model without discounting:

$$\max \int_0^\infty [u(c(t)) - u(c^*)]\, dt$$

subject to

$$\dot{k}(t) = f(k(t)) - c(t) - \delta k(t)$$

with initial condition $k(0) > 0$, and c^* defined as the golden rule consumption level

$$c^* = f(k^*) - \delta k^*$$

where k^* is the golden rule capital-labor ratio given by $f'(k^*) = \delta$.

(1) Set up the Hamiltonian for this problem with costate variable $\lambda(t)$.

(2) Characterize the solution to this optimal growth program.

(3) Show that the transversality condition that $\lim_{t\to\infty}\lambda(t)k(t) = 0$ is not satisfied at the optimal solution. Explain why this is the case.

Exercise 7.22. Consider the infinite-horizon optimal control problem given by the maximization of (7.28) subject to (7.29) and (7.30). Suppose that the problem has a quasi-stationary structure, so that

$$\begin{aligned} f(t, x, y) &\equiv \beta(t) f(x, y) \\ g(t, x, y) &\equiv g(x, y), \end{aligned}$$

where $\beta(t)$ is the discount factor that applies to returns that are an interval of time t away from the present.

(1) Set up Hamiltonian and characterize the necessary conditions for this problem.

(2) Prove that the solution to this problem is time-consistent (meaning that the solution chosen at some date s cannot be improved upon at some future date s' by changing the continuation plans after this date) if and only if $\beta(t) = \exp(-\rho t)$ for some $\rho \geq 0$.
(3) Interpret this result and explain in what way the conclusion is different from that of Lemma 7.1.

Exercise 7.23. Consider the following maximization problem:

$$\max_{x(t),y(t)} \int_0^1 f(x(t), y(t))\, dt$$

subject to

$$\dot{x}(t) = y(t)^2$$

$x(0) = 0$ and $x(1) = 0$, where $y(t) \in \mathbb{R}$ and f is an arbitrary continuously differentiable function. Show that the unique solution to this maximization problem does not satisfy the necessary conditions in Theorem 7.2. Explain why this is and relate your answer to Theorem 7.15.

Exercise 7.24. * Consider the following infinite-horizon utility maximization problem

$$\max \int_0^\infty \exp(-\rho t)\, u(c(t))\, dt$$

subject to initial capital $k(0)$ and law of motion of capital given by

$$\dot{k}(t) = \begin{cases} f(k(t)) - \delta k(t) - c(t) & \text{if } f(k(t)) - c(t) \geq \underline{k} \\ 0 & \text{if } f(k(t)) - c(t) < \underline{k}, \end{cases}$$

where $\underline{k} > \delta k^*$ for $k^* = f^{-1}(\rho + \delta)$ is a minimum investment size requirement. Suppose that $k(0) = k^*$. Show that there does not exist a solution to this optimal control problem. Explain why this is and relate your answer to Theorem 7.15. [Hint: Show that $k(t) = \delta k^*$ for all t, which would have been the optimal policy without the minimum investment size requirement, is suboptimal. Then, show that any policy that induces $k(t) = k'$ can be improved upon by a policy that alternates between $c(t) = f(k(t)) - \underline{k}$ and $c(t) = f(k(t))$, spending an interval of length $\Delta > 0$ in each of the two investment regime. Finally, show that the same alternating policy for Δ arbitrarily small improves over any other admissible pair of $(k(t), c(t))$].

Exercise 7.25. (1) Consider the optimal growth problem in Section 7.7. Show that if $u(c) = \log c$, Assumption 7.1(iii) is violated.
(2) Suppose next that $u : \mathbb{R}_+ \setminus \{0\} \to \mathbb{R}$ is continuous, strictly increasing and convex, and that u is continously differentiable on $[\varepsilon, +\infty)$ for some $\varepsilon > 0$ sufficiently small. Show that if the modified problem where $c(t)$ is restricted to be in $[\varepsilon, +\infty)$ has an optimal control where $c(t) > \varepsilon$ for all t, then this optimal control also constitutes a solution to the original problem where $u : \mathbb{R}_+ \setminus \{0\} \to \mathbb{R}$. [Hint: use Theorem 7.14].
(3) Show that for any optimal growth problem as described in Section 7.7, there exists $\varepsilon > 0$ such that the optimal control will involve $c(t) > \varepsilon$ for all t. [Hint: use the Inada condition on the production function from Assumption 2].

Exercise 7.26. Consider the following continuous-time discounted infinite-horizon problem:

$$\max \int_0^{\infty} \exp(-\rho t)\, u(c(t))\, dt$$

subject to

$$\dot{x}(t) = g(x(t)) - c(t)$$

with initial condition $x(0) > 0$.

Suppose that $u(\cdot)$ is strictly increasing and strictly concave, with $\lim_{c\to\infty} u'(c) = 0$ and $\lim_{c\to 0} u'(c) = \infty$, and $g(\cdot)$ is increasing and strictly concave with $\lim_{x\to\infty} g'(x) = 0$ and $\lim_{x\to 0} g'(x) = \infty$.

(1) Set up the current-value Hamiltonian and derive the Euler equations for an optimal path.

(2) Show that the standard transversality condition and the Euler equations are necessary and sufficient for a solution.

(3) Characterize the optimal path of solutions and their limiting behavior.

Exercise 7.27. Consider the q-theory model of Section 7.8.

(1) Verify that when $k \in [\varepsilon, \bar{k}]$, Assumption 7.1 is satisfied.

(2) Consider the modified problem with $k \in [\varepsilon, \bar{k}]$. Show that if there exists a solution to this modified problem with $k(t) > \varepsilon$ for all t, then this is also a solution to the original problem where $k \in \mathbb{R}_+$. [Hint: use the same line of argument as in Exercise 7.25].

(3) Show that there always exists $\varepsilon > 0$ sufficiently small, such that the solution to the modified problem involves $k(t) > \varepsilon$ for all t.

Exercise 7.28. (1) In the q-theory of investment, prove that when $\phi''(i) = 0$ (for all i), investment jumps so that the capital stock reaches its steady-state value k^* immediately.

(2) Prove that as shown in Figure 7.1, the curve for (7.85) is downward sloping in the neighborhood of the steady state.

(3) As an alternative to the diagrammatic analysis of Figure 7.1, linearize (7.81) and (7.85), and show that in the neighborhood of the steady state this system has one positive and one negative eigenvalue. Explain why this implies that optimal investment plans will tend towards the stationary solution (steady state).

(4) Prove that when $k(0) < k^*$, $i(0) > i^*$ and $i(t) \downarrow i^*$.

(5) Derive the equations for the q-theory of investment when the adjustment cost takes the form $\phi(i/k)$. How does this affect the steady-state marginal product of capital?

(6) Derive the optimal equation path when investment is irreversible, in the sense that there is an additional constraint, $\dot{i} \geq 0$.

Part 3

Neoclassical Growth

This part of the book covers the basic workhorse models of economic growth. I start with the infinite-horizon neoclassical growth model, which was already discussed in the previous three chapters. A closely related model is the baseline overlapping-generations model of Samuelson and Diamond, and this is the topic of Chapter 9. Despite the similarities between the two models, they have quite different normative and positive implications, and each model may be appropriate for different sets of issues. It is therefore important to have a detailed discussion of both.

This part of the book also presents the basic economic growth model with human capital investments. The inclusion of this model is motivated both because of the increasingly important role that human capital plays in economic growth and macroeconomics, and also as a way of linking macroeconomic approaches to microdata on schooling and returns to education.

Finally, Chapter 11 introduces the first models of sustained economic growth. These are contained in this part of the book rather than the next, because they are models of sustained growth *without* technological change. Despite their simplicity, these models lead to a number of important economic insights and provide a good introduction to the set of issues that will be discussed in the next part of the book.

CHAPTER 8

The Neoclassical Growth Model

We are now ready to start our analysis of the standard neoclassical growth model (also known as the Ramsey or Cass-Koopmans model). This model differs from the Solow model only in one crucial respect: it explicitly models the consumer side and endogenizes savings. In other words, it introduces household optimization. Beyond its use as a basic growth model, this model has become a workhorse for many areas of macroeconomics, including the analysis of fiscal policy, taxation, business cycles, and even monetary policy.

Since both the basic equilibrium and optimal growth models in discrete time were already presented as applications of dynamic programming in Chapter 6, this chapter focuses on the continuous-time neoclassical growth model, except in Section 8.6, which sketches the characterization of the competitive equilibrium in discrete time, and in exercises.

8.1. Preferences, Technology and Demographics

8.1.1. Basic Environment. Consider an infinite-horizon economy in continuous time and suppose that the economy admits a representative household with instantaneous utility function

$$u\left(c\left(t\right)\right). \tag{8.1}$$

The following standard assumption on this utility function will be maintained throughout the book unless stated otherwise.

ASSUMPTION 3. ***(Neoclassical Preferences)*** *The utility function $u\left(c\right)$ is defined on $\mathbb{R}_{+}$ or $\mathbb{R}_{+}\backslash\left\{0\right\}$. It is strictly increasing, concave, twice differentiable in the interior of its domain, with derivatives u' and u''.*

More explicitly, the reader may wish to suppose that the economy consists of a set of identical households (with measure normalized to 1). Each household has an instantaneous utility function given by (8.1). Population within each household grows at the rate n, starting with $L\left(0\right)=1$, so that total population in the economy is

$$L\left(t\right)=\exp\left(nt\right). \tag{8.2}$$

All members of the household supply their labor inelastically.

Our baseline assumption is that the household is fully altruistic towards all of its future members, and always makes the allocations of consumption (among household members)

cooperatively. This implies that the objective function of each household at time $t = 0$, $U(0)$, can be written as

$$U(0) \equiv \int_0^{\infty} \exp\left(-(\rho - n)\, t\right) u\left(c\left(t\right)\right) dt, \tag{8.3}$$

where $c(t)$ is consumption per capita at time t, ρ is the subjective discount rate, and the effective discount rate is $\rho - n$, since it is assumed that the household derives utility from the consumption per capita of its additional members in the future as well (see Exercise 8.1).

It is useful to be a little more explicit about where the objective function (8.3) is coming from. First, given the strict concavity of $u(\cdot)$ and the assumption that within-household allocation decisions are cooperative, each household member will have an equal consumption (Exercise 8.1). This implies that each member will consume

$$c\left(t\right) \equiv \frac{C\left(t\right)}{L\left(t\right)}$$

at date t, where $C(t)$ is total consumption and $L(t)$ is the size of the representative household (equal to total population, since the measure of households is normalized to 1). This implies that the household will receive a utility of $u\left(c\left(t\right)\right)$ per household member at time t, or a total utility of $L\left(t\right) u\left(c\left(t\right)\right) = \exp\left(nt\right) u\left(c\left(t\right)\right)$. Since utility at time t is discounted back to time 0 with a discount rate of $\exp\left(-\rho t\right)$, we obtain the expression in (8.3).

Let us also assume:

Assumption 4′. **(Discounting)**

$$\rho > n.$$

This condition ensures that there is in fact discounting of future utility streams. Otherwise, (8.3) would have infinite value, and standard optimization techniques would not be useful in characterizing optimal plans. Assumption 4′ makes sure that in the model without growth, discounted utility is finite. When there is sustained growth, this condition will be strengthened to Assumption 4.

Let us start with an economy without any technological progress. Factor and product markets are competitive, and the production possibilities set of the economy is represented by the aggregate production function

$$Y\left(t\right) = F\left[K\left(t\right), L\left(t\right)\right],$$

which is a simplified version of the production function (2.1) used in the Solow growth model in Chapter 2. In particular, there is now no technology term (labor-augmenting technological change will be introduced below). As in the Solow model, Assumptions 1 and 2 are imposed throughout. The constant returns to scale feature enables us to work with the per capita

production function $f(\cdot)$ such that, output per capita is given by

$$\begin{aligned} y(t) &\equiv \frac{Y(t)}{L(t)} \\ &= F\left[\frac{K(t)}{L(t)}, 1\right] \\ &\equiv f(k(t)), \end{aligned}$$

where, as before,

$$k(t) \equiv \frac{K(t)}{L(t)}. \tag{8.4}$$

Competitive factor markets then imply that the rental rate of capital and the wage rate at time t are given by:

$$R(t) = F_K[K(t), L(t)] = f'(k(t)) \tag{8.5}$$

and

$$w(t) = F_L[K(t), L(t)] = f(k(t)) - k(t) f'(k(t)). \tag{8.6}$$

The demand side is somewhat more complicated, since each household will solve a continuous-time optimization problem in deciding how to use their assets and allocate consumption over time. To prepare for this, let us denote the asset holdings of the representative household at time t by $\mathcal{A}(t)$. Then, the law of motion for the total assets of the household is

$$\dot{\mathcal{A}}(t) = r(t)\mathcal{A}(t) + w(t)L(t) - c(t)L(t), \tag{8.7}$$

where $c(t)$ is consumption per capita of the household, $r(t)$ is the risk-free market rate of return on assets, and $w(t)L(t)$ is the flow of labor income earnings of the household. Defining per capita assets as

$$a(t) \equiv \frac{\mathcal{A}(t)}{L(t)},$$

dividing (8.7) by $L(t)$, substituting for the previous definition and using the fact that $L(t)$ grows at the rate n, the law of motion of per capita assets is obtained as

$$\dot{a}(t) = (r(t) - n)a(t) + w(t) - c(t). \tag{8.8}$$

In practice, household assets can consist of (claims to) capital stock, $K(t)$, which they rent to firms and government bonds, $B(t)$. In models with uncertainty, households would have a portfolio choice between the capital stock of the corporate sector and riskless bonds (typically assumed to be supplied by the government). Bonds play an important role in models with incomplete markets, allowing households to smooth idiosyncratic shocks. Since these bonds are in zero net supply, in the aggregate $B(t) = 0$, and thus market clearing implies that assets per capita must be equal to the capital stock per capita. That is,

$$a(t) = k(t). \tag{8.9}$$

Because there is no uncertainty here, I ignore government bonds (until Chapter 17). Since household assets are the same as the capital stock and capital depreciates at the rate δ, the market rate of return on assets is

$$r(t) = R(t) - \delta. \tag{8.10}$$

8.1.2. The Natural Debt Limit and the No-Ponzi Game Condition. Equation (8.8) is a flow constraint. As already discussed in detail in Example 6.5 in Chapter 6, this flow constraint is not sufficient as a "proper budget constraint" on household behavior (unless we impose a lower bound on assets, such as $a(t) \geq 0$ for all t, which is too strong; see Exercise 8.28). And if we do not ensure that there is a proper budget constraint on household behavior, the analysis of household maximization leads to nonsensical results. In particular, Exercise 8.2 shows that any solution to the maximization of (8.3) with respect to (8.8)—without an additional of the condition—involves $a(t)$ becoming arbitrarily negative for all t, that is, the representative household holds an arbitrarily negative asset position. But then, the market clearing condition, (8.9), implies that $k(t) = a(t)$ is arbitrarily negative, which violates the feasibility constraint that $k(t)$ has to be nonnegative. Clearly, the flow budget constraint, (8.8), is not sufficient to capture the full set of constraints on household behavior.

There are two ways to proceed. The first is to impose the so-called "no-Ponzi condition," which turns out to be the more convenient and flexible approach. The second involves imposing a *natural debt limit* (as in Example 6.5). Recall that the natural debt limit imposes that $a(t)$ should never become so negative that the household cannot repay its debts even if it chooses zero consumption thereafter. Using (8.8) and assuming that the household does not consume from time t onwards, the natural debt limit in this case is found as

$$a(t) \geq -\int_t^\infty w(s) \exp\left(-\int_t^s (r(z) - n)\, dz\right) ds. \tag{8.11}$$

This is the direct analog of the discrete-time natural debt limit (6.41) in Chapter 6. In particular, the right-hand side is the negative of the net present discounted value of labor income for the household. Exercise 8.3 asks you to work through a more detailed derivation of this condition. Any path of consumption and assets for the household that violates (8.11) is not feasible (unless we allow for bankruptcy). In view of this, the problem of the representative household can be expressed as the maximization of (8.3) subject to (8.8) and (8.11). In fact, one could simply impose the limiting version of this constraint,

$$\lim_{t\to\infty} a(t) \geq \hat{a} \equiv -\lim_{t\to\infty} \int_t^\infty w(s) \exp\left(-\int_t^s (r(z) - n)\, dz\right) ds. \tag{8.12}$$

This is because, if the natural debt limit is violated for some $t' < \infty$, then (8.12) cannot be satisfied either (see Exercise 8.4). There is therefore no loss of generality in imposing (8.12) instead of (8.11).

However, we will see below that the natural debt limit is not useful when we look at economies with sustained growth (because in that case $\hat{a} = -\infty$, see Exercise 8.8). For this reason, we must now undertake the investment for understanding the no-Ponzi game

condition. To start with, let us write the lifetime budget constraint of a household as

$$\begin{aligned} &\int_0^T c(t) L(t) \exp\left(\int_t^T r(s)\, ds\right) dt + \mathcal{A}(T) \\ &= \int_0^T w(t) L(t) \exp\left(\int_t^T r(s)\, ds\right) dt + \mathcal{A}(0) \exp\left(\int_0^T r(s)\, ds\right), \end{aligned} \tag{8.13}$$

for some arbitrary $T > 0$. This constraint states that the household's asset position at time T is given by his total income plus initial assets minus expenditures, all carried forward to date T units. Differentiating this expression with respect to T and dividing $L(t)$ gives (8.8) (see Exercise 8.5).

Now imagine that (8.13) applies to a finite-horizon economy ending at date T. At this point, the household cannot have negative asset holdings, thus (8.13) must hold with $\mathcal{A}(T) \geq 0$. However, inspection of the flow budget constraint (8.8) makes it clear that this constraint does not guarantee $\mathcal{A}(T) \geq 0$. Therefore, in the finite horizon, we also need to impose $\mathcal{A}(T) \geq 0$ as an additional terminal value constraint.

In the infinite-horizon case, we need a similar constraint. The appropriate restriction is the *no-Ponzi condition* (or the no-Ponzi game condition). It takes the form

$$\lim_{t\to\infty} a(t) \exp\left(-\int_0^t (r(s) - n)\, ds\right) \geq 0. \tag{8.14}$$

This condition is stated as an inequality, to ensure that the representative household does not asymptotically tend to a negative wealth. Intuitively, without (8.14), there is no proper lifetime budget constraint on the representative household (and by implication on any of the households in the economy) and they all increase their consumption by borrowing to such a level that feasibility is violated. This could clearly not be an equilibrium. At a more fundamental level, the relationship between the representative household and the financial markets must impose a proper lifetime budget constraint; otherwise financial markets would lose money. In the finite-horizon economy this is the constraint $\mathcal{A}(T) \geq 0$. In the infinite-horizon economy, it is (8.14).

The name "no-Ponzi condition" for (8.14) comes from the chain-letter or pyramid schemes, which are sometimes called Ponzi games, where an individual can continuously borrow from a competitive financial market (or more often, from unsuspecting souls that become part of the chain-letter scheme) and pay his or her previous debts using current borrowings. The consequence of this scheme would be that the asset holding of the individual would tend to $-\infty$ as time goes by, violating feasibility at the economy level.

Let us momentarily return to the finite-horizon problem. In this case, financial markets would impose $\mathcal{A}(T) \geq 0$. But the household itself would never choose $\mathcal{A}(T) > 0$, so the budget constraint could be simplified and written with $\mathcal{A}(T) = 0$. In the same way, the no-Ponzi condition will imply

$$\lim_{t\to\infty} \left[a(t) \exp\left(-\int_0^t (r(s) - n)\, ds\right)\right] = 0. \tag{8.15}$$

To further highlight the parallel between the finite-horizon budget constraint and the no-Ponzi condition, multiply both sides of (8.13) by $\exp\left(-\int_0^T r(s)\,ds\right)$ to obtain

$$\int_0^T c(t)\,L(t)\exp\left(-\int_0^t r(s)\,ds\right)dt+\exp\left(-\int_0^T r(s)\,ds\right)\mathcal{A}(T)$$
$$=\int_0^T w(t)\,L(t)\exp\left(-\int_0^t r(s)\,ds\right)dt+\mathcal{A}(0),$$

then divide everything by $L(0)$ and note that $L(t)$ grows at the rate n, to obtain

$$\int_0^T c(t)\exp\left(-\int_0^t (r(s)-n)\,ds\right)dt+\exp\left(-\int_0^T (r(s)-n)\,ds\right)a(T)$$
$$=\int_0^T w(t)\exp\left(-\int_0^t (r(s)-n)\,ds\right)dt+a(0).$$

Now take the limit as $T\to\infty$ and use the stronger (equality) no-Ponzi condition (8.15) to obtain

$$\int_0^\infty c(t)\exp\left(-\int_0^t (r(s)-n)\,ds\right)dt=a(0)+\int_0^\infty w(t)\exp\left(-\int_0^t (r(s)-n)\,ds\right)dt,$$

which requires the discounted sum of expenditures to be equal to initial income plus the discounted sum of labor income. Therefore this equation is a direct extension of (8.13) to infinite horizon. This derivation makes it clear that the no-Ponzi condition (8.14), or (8.15), essentially ensures that the household's lifetime budget constraint holds in infinite horizon. In the analysis that follows, I will impose this constraint on household behavior.

8.2. Characterization of Equilibrium

8.2.1. Definition of Equilibrium. Let us now define an equilibrium in this dynamic economy. I provide two definitions, the first is somewhat more formal, while the second definition will be more useful in characterizing the equilibrium below.

DEFINITION 8.1. *A **competitive equilibrium** of the neoclassical growth model (Ramsey economy) consists of paths of consumption, capital stock, wage rates and rental rates of capital, $[C(t),K(t),w(t),R(t)]_{t=0}^{\infty}$, such that the representative household maximizes its utility given initial asset holdings (capital stock) $K(0)>0$ and the time path of prices $[w(t),R(t)]_{t=0}^{\infty}$, and all markets clear.*

Recall again that an equilibrium corresponds to the entire time path of real quantities and the associated prices. We may sometimes focus on the steady-state equilibrium, but equilibrium always refers to the entire path.

Since everything can be equivalently defined in terms of per capita variables, an alternative and more convenient definition of equilibrium is as follows:

DEFINITION 8.2. *A **competitive equilibrium** of the neoclassical growth model (Ramsey economy) consists of paths of per capita consumption, capital-labor ratio, wage rates and*

rental rates of capital, $[c(t), k(t), w(t), R(t)]_{t=0}^{\infty}$*, such that the representative household maximizes (8.3) subject to (8.8) and (8.14) given initial per capita asset holdings (capital-labor ratio)* $k(0) > 0$*, factor prices* $[w(t), R(t)]_{t=0}^{\infty}$ *as in (8.5) and (8.6), and the rate of return on assets* $r(t)$ *given by (8.10).*

8.2.2. Household Maximization. Let us start with the problem of the representative household. From the definition of equilibrium we know that this is to maximize (8.3) subject to (8.8) and (8.14). This is a special case of discounted infinite-horizon control problems discussed in Theorem 7.13 in the previous chapter. To derive the necessary conditions for an optimal consumption-saving plan for the household, let us set up the current-value Hamiltonian:

$$\hat{H}(t, a, c, \mu) = u(c(t)) + \mu(t)[w(t) + (r(t) - n)a(t) - c(t)],$$

with state variable a, control variable c and current-value costate variable μ. This problem is closely related to the problem of optimal growth already discussed in Section 7.7 in the previous chapter, but with two main differences: first, the rate of return on assets is time varying; second, the terminal value constraint is represented by the no-Ponzi condition (8.14), which is slightly different from the one in Section 7.7, which was $\lim_{t\to\infty} k(t) \geq 0$. An argument similar to that in Section 7.7 makes sure that Theorem 7.13 can be applied to this problem (see Exercise 8.7).

Now applying Theorem 7.13, the necessary conditions for an interior solution are

$$\hat{H}_c(t, a, c, \mu) = u'(c(t)) - \mu(t) = 0 \tag{8.16}$$

$$\hat{H}_a(t, a, c, \mu) = \mu(t)(r(t) - n) = -\dot{\mu}(t) + (\rho - n)\mu(t), \tag{8.17}$$

and the transition equation (8.8). The transversality condition (the equivalent of (7.60) in the previous chapter) is

$$\lim_{t\to\infty} [\exp(-(\rho - n)t)\mu(t)a(t)] = 0. \tag{8.18}$$

The transversality condition is written in terms of the current-value costate variable, which is more convenient given the rest of the necessary conditions. Note also that (8.16), (8.17) and (8.18) are necessary conditions for an interior solution. Are they also sufficient? Notice that, from (8.16), for any $\mu(t) > 0$. This implies that $\hat{H}(t, a, c, \mu)$ is a concave function of (a, c) (though not strictly so). Moreover, we will see that any $(a(t), c(t))$ that satisfies (8.14) must also satisfy $\lim_{t\to\infty} [\exp(-(\rho - n)t)\mu(t)a(t)] \geq 0$, so that the conditions of Theorem 7.14 are met. Therefore, conditions (8.16)-(8.18) are also sufficient and characterize a solution to the household maximization problem.

Let us now use these equations to derive a more explicit characterization. First, rearrange the second condition to obtain:

$$\frac{\dot{\mu}(t)}{\mu(t)} = -(r(t) - \rho), \tag{8.19}$$

which states that the multiplier changes depending on whether the rate of return on assets is currently greater than or less than the discount rate of the household.

Next, the first necessary condition above implies that

$$u'(c(t)) = \mu(t). \tag{8.20}$$

To make more progress, let us differentiate the previous expression with respect to time and divide by $\mu(t)$, which yields

$$\frac{u''(c(t))\,c(t)}{u'(c(t))}\frac{\dot{c}(t)}{c(t)} = \frac{\dot{\mu}(t)}{\mu(t)}.$$

Substituting this into (8.19) yields another form of the consumer Euler equation:

$$\frac{\dot{c}(t)}{c(t)} = \frac{1}{\varepsilon_u(c(t))}(r(t) - \rho) \tag{8.21}$$

where

$$\varepsilon_u(c(t)) \equiv -\frac{u''(c(t))\,c(t)}{u'(c(t))} \tag{8.22}$$

is the elasticity of the marginal utility $u'(c(t))$. This equation is closely related to the consumer Euler equation derived in the context of the discrete-time problem, eq. (6.36), as well as to the consumer Euler equation in continuous time with constant interest rates in Example 7.1 in the previous chapter. As with eq. (6.36), it states that consumption will grow over time when the discount rate is less than the rate of return on assets. It also specifies the speed at which consumption will grow in response to a gap between this rate of return and the discount rate, which is related to the elasticity of marginal utility of consumption, $\varepsilon_u(c(t))$. The interpretation of $\varepsilon_u(c(t))$ and of this equation will be discussed further below.

Next, integrating eq. (8.19),

$$\begin{aligned}\mu(t) &= \mu(0)\exp\left(-\int_0^t (r(s) - \rho)\,ds\right) \\ &= u'(c(0))\exp\left(-\int_0^t (r(s) - \rho)\,ds\right),\end{aligned} \tag{8.23}$$

where the first line uses the form of the solutions to linear nonhomogeneous differential equations (see Section B.4 in Appendix Chapter B), and the second line uses (8.16) evaluated at time $t = 0$. Substituting (8.23) into the transversality condition yields

$$\begin{aligned}\lim_{t\to\infty}\left[\exp(-(\rho - n)t)\,a(t)\,u'(c(0))\exp\left(-\int_0^t (r(s) - \rho)\,ds\right)\right] &= 0 \\ \lim_{t\to\infty}\left[a(t)\exp\left(-\int_0^t (r(s) - n)\,ds\right)\right] &= 0.\end{aligned} \tag{8.24}$$

Next, note also that, since $a(t) = k(t)$, the transversality condition of the representative household can alternatively be written as

$$\lim_{t\to\infty}\left[k(t)\exp\left(-\int_0^t (r(s) - n)\,ds\right)\right] = 0. \tag{8.25}$$

This equation emphasizes that the transversality condition requires the discounted market value of the capital stock in the very far future to be equal to 0. This "market value" version of the transversality condition is both intuitive and often more convenient to work with.

To complete the analysis of equilibrium growth and show that the conditions described above are also sufficient, we need to ensure that the conditions of Theorem 7.14 are satisfied. As noted above, in view of the fact that $\mu(t) > 0$ for all t, the current-value Hamiltonian $\hat{H}(t, a, c, \mu)$ is a concave function of (a, c). Thus it only remains to show that for any feasible $(a(t), c(t))$, $\lim_{t\to\infty} [\exp(-(\rho - n)t)\mu(t)a(t)] \geq 0$. Substituting for $\mu(t)$ from (8.23), this again requires

$$\lim_{t\to\infty} \left[a(t) \exp\left(- \int_0^t (r(s) - n)\, ds \right) \right] \geq 0.$$

But, this is identical to the no-Ponzi condition (8.14). Thus, no feasible $(a(t), c(t))$ can violate this condition. Therefore, Theorem 7.14 implies that any $(\hat{a}(t), \hat{c}(t))$ that satisfies (8.21) and (8.24) is a solution to the household maximization problem. By implication, any $(\hat{k}(t), \hat{c}(t))$ that satisfies (8.21) and (8.25) will be part of a competitive equilibrium. The full competitive equilibrium is such a $(\hat{k}(t), \hat{c}(t))$ process combined with market clearing prices. Recall that equilibrium prices are given by (8.5) and (8.6). From (8.10), this implies that the market rate of return for consumers, $r(t)$, is given by

$$r(t) = f'(k(t)) - \delta. \tag{8.26}$$

Substituting this into the household's maximization problem, we obtain

$$\frac{\dot{c}(t)}{c(t)} = \frac{1}{\varepsilon_u(c(t))} \left(f'(k(t)) - \delta - \rho \right) \tag{8.27}$$

as the equilibrium version of the consumption growth equation, (8.21). Substituting it into (8.25), the transversality condition becomes

$$\lim_{t\to\infty} \left[k(t) \exp\left(- \int_0^t \left(f'(k(s)) - \delta - n \right) ds \right) \right] = 0. \tag{8.28}$$

Conditions (8.27) and (8.28) are expressed only in terms of the path of the capital-labor ratio and fully characterize a competitive equilibrium. Moreover, we will see shortly that this competitive equilibrium is unique and that these conditions are identical to the Euler equation and the transversality condition in the optimal growth problem (see Exercise 8.11).

8.2.3. Consumption Behavior. Let us now return to the dynamics of the representative household's consumption. First, recall that the Euler equation, (8.21), relates the "slope" of the consumption profile of the households to $\varepsilon_u(c(t))$, which was defined as the elasticity of the marginal utility function, $u'(c)$. Notice, however, that $\varepsilon_u(c(t))$ is not only the elasticity of marginal utility, but even more importantly, it is the inverse of the *intertemporal elasticity of substitution*, which plays a crucial role in macro models. The intertemporal elasticity of substitution regulates the willingness of households to substitute consumption (or labor or any other attribute that yields utility) over time. The elasticity of marginal utility of consumption between the dates t and $s > t$ is defined as

$$\sigma_u(t, s) = -\frac{d \log(c(s)/c(t))}{d \log(u'(c(s))/u'(c(t)))}.$$

As $s \downarrow t$,

$$\sigma_u(t,s) \to \sigma_u(t) = -\frac{u'(c(t))}{u''(c(t))\,c(t)} = \frac{1}{\varepsilon_u(c(t))}. \tag{8.29}$$

This is not surprising, since the concavity of the utility function $u(\cdot)$—or equivalently, the elasticity of marginal utility—determines how willing households are to substitute consumption over time.

It is also possible to derive some further results on the consumption behavior of households. Notice that the term $\exp\left(-\int_0^t r(s)\,ds\right)$ is a present-value factor that converts a unit of income at time t to a unit of income at time 0. In the special case where $r(s) = r$, this factor would be exactly equal to $\exp(-rt)$. But more generally, we can define an average interest rate between dates 0 and t as

$$\bar{r}(t) = \frac{1}{t}\int_0^t r(s)\,ds. \tag{8.30}$$

In that case, the conversion factor between dates 0 and t is $\exp(-\bar{r}(t)\,t)$, and the transversality condition can be written as

$$\lim_{t\to\infty}\left[\exp\left(-\left(\bar{r}(t)-n\right)t\right)a(t)\right] = 0. \tag{8.31}$$

Now integrating (8.21), we obtain (recall again Section B.4 in Appendix Chapter B):

$$c(t) = c(0)\exp\left(\int_0^t \frac{r(s)-\rho}{\varepsilon_u(c(s))}ds\right)$$

as the consumption function. Therefore, given the initial consumption level, $c(0)$, the path of consumption is determined. In the special case where $\varepsilon_u(c(s))$ is constant, for example, $\varepsilon_u(c(s)) = \theta$, this equation simplifies to

$$c(t) = c(0)\exp\left(\left(\frac{\bar{r}(t)-\rho}{\theta}\right)t\right),$$

and moreover, the lifetime budget constraint simplifies to

$$\int_0^\infty c(t)\exp\left(-\left(\bar{r}(t)-n\right)t\right)dt = a(0) + \int_0^\infty w(t)\exp\left(-\left(\bar{r}(t)-n\right)t\right)dt,$$

and substituting for $c(t)$ into this lifetime budget constraint in this isoelastic case, we obtain

$$c(0) = \left[\int_0^\infty \exp\left(-\left(\frac{(1-\theta)\,\bar{r}(t)}{\theta} - \frac{\rho}{\theta} + n\right)t\right)dt\right] \times \left[a(0) + \int_0^\infty w(t)\exp\left(-\left(\bar{r}(t)-n\right)t\right)dt\right] \tag{8.32}$$

as the initial value of consumption. Once the initial consumption level is determined, the Euler equation (8.21) fully determines the path of utility-maximizing consumption of the household.

8.2.4. Using the Natural Debt Limit. Let us now return to the alternative approach of using the natural debt limit (8.12) instead of the no-Ponzi condition (8.14). In principle, $\hat{a}$ in (8.12) could be equal to minus infinity. However, exactly the same steps as in Exercise 8.2 show that this would violate the feasibility constraint that $k(t)$ has to be nonnegative. In light of this, a well-defined equilibrium must involve $\hat{a} > -\infty$. We will see below that factor

prices satisfy $\lim_{t\to\infty} r(t) > \rho$ and $\lim_{t\to\infty} w(t) = w \geq 0$, and this is sufficient for $\hat{a} > -\infty$ (see Exercise 8.9). When $\hat{a} > -\infty$, the maximization of (8.3) subject to (8.8) and (8.11) again satisfies the conditions of Theorems 7.13 and 7.14. Therefore, the alternative approach of modeling the household's problem as maximizing (8.3) subject to (8.8) and (8.11) leads to the same characterization as the one using the no-Ponzie condition (see Exercise 8.10). It can then be verified that $\lim_{t\to\infty} r(t) > \rho$ and $\lim_{t\to\infty} w(t) = w \geq 0$, so that $\hat{a} > -\infty$, and the approach is indeed a valid. However, as Exercise 8.8 shows, this approach does not work when there is sustained growth as in Section 8.7.

8.3. Optimal Growth

Before characterizing the competitive equilibrium further, it is useful to look at the optimal growth problem. Recall that this involves characterizing the capital-labor ratio and consumption path that maximizes the utility of the representative household. The problem of maximizing the utility of the representative househol subject to technology and feasibility constraints can be written as:

$$\max_{[k(t),c(t)]_{t=0}^{\infty}} \int_0^{\infty} \exp\left(-(\rho - n)\, t\right) u\left(c\left(t\right)\right) dt,$$

subject to

$$\dot{k}\left(t\right) = f\left(k\left(t\right)\right) - (n+\delta)k\left(t\right) - c\left(t\right),$$

and $k\left(0\right) > 0$.[1] As noted in Chapter 5, versions of the First and Second Welfare Theorems for economies with a continuum of commodities would imply that the solution to this problem should be the same as the equilibrium growth problem of the previous section. However, in the present context there is no need to appeal to these theorems since it is straightforward to characterize both allocations and show their equivalence.

To do this, let us once again set up the current-value Hamiltonian, which in this case takes the form

$$\hat{H}\left(k, c, \mu\right) = u\left(c\left(t\right)\right) + \mu\left(t\right)\left[f\left(k\left(t\right)\right) - (n+\delta)k\left(t\right) - c\left(t\right)\right],$$

with state variable k, control variable c and current-value costate variable μ. With the same argument as in Section 7.7 in the previous chapter, it can be shown that Theorems 7.13 and 7.14 can be applied to this problem and characterize the unique optimal growth path (in particular, recall Exercise 7.25). Consequently, the necessary and sufficient conditions for

[1] In the case where the infinite-horizon problem represents dynastic utilities as discussed in Chapter 5, this specification presumes that the social planner gives the same weights to different generations as the current dynastic decision-maker. Naturally, there also exist Pareto optimal allocations with unequal distribution of consumption across households, though these are less natural and less interesting in the context of representative household economies.

this optimal path are:

$$\begin{aligned} \hat{H}_c(k,c,\mu) &= 0 = u'(c(t)) - \mu(t), \\ \hat{H}_k(k,c,\mu) &= -\dot{\mu}(t) + (\rho - n)\mu(t) = \mu(t)\left(f'(k(t)) - \delta - n\right), \\ \lim_{t\to\infty}\left[\exp\left(-(\rho-n)t\right)\mu(t)k(t)\right] &= 0. \end{aligned} \tag{8.33}$$

Repeating the same steps as before, we can combine the first two optimality conditions in (8.33) and obtain (8.27) for the path of consumption of the representative household. In addition, once again integrating the second first-order condition,

$$\mu(t) = \mu(0)\exp\left(-\int_0^t \left(f'(k(s)) - \delta - \rho\right) ds\right).$$

Combining this with the first condition in (8.33) evaluated at $t = 0$ implies that $\mu(0) = u'(c(0)) > 0$. Substituting this expression into the transversality condition (the third condition in (8.33)), simplifying and canceling out $\mu(0) > 0$, we obtain (8.28).

These steps establish that the competitive equilibrium is a Pareto optimum and that the optimal growth path can be decentralized as a competitive equilibrium. This result is stated in the next proposition:

PROPOSITION 8.1. *In the neoclassical growth model described above, with Assumptions 1, 2, 3 and 4′, the equilibrium is Pareto optimal and coincides with the optimal growth path maximizing the utility of the representative household.*

8.4. Steady-State Equilibrium

A steady-state equilibrium is defined as in Chapter 2 as an equilibrium path in which capital-labor ratio, consumption and output are constant. The steady-state equilibrium (by the equivalence between the two problems, also the stationary solution to the optimal growth problem) is straightforward to characterize. Steady state requires that consumption per capita is constant, thus

$$\dot{c}(t) = 0.$$

From (8.27), this implies that as long as $f'(k^*) > 0$, *regardless* of the exact utility function, we must have a capital-labor ratio k^* such that

$$f'(k^*) = \rho + \delta, \tag{8.34}$$

which is the equivalent of the steady-state relationship in the discrete-time optimal growth model.[2] This equation pins down the steady-state capital-labor ratio only as a function of the production function, the discount rate and the depreciation rate. This corresponds to the *modified golden rule*, rather than the golden rule in the Solow model (see Exercise 8.17). The modified golden rule involves a level of the capital stock that does not maximize steady-state consumption, because earlier consumption is preferred to later consumption. This is because

[2]In addition, there again exists another, economically uninteresting steady state at $k = 0$. As in Chapter 2, I ignore this steady state throughout. Moreover, as in Chapter 2, starting with any $k(0) > 0$ the economy will always tend to the steady-state capital-labor ratio k^* given by (8.34).

of discounting, which means that the objective is not to maximize steady-state consumption, but involves giving a higher weight to earlier consumption.

Note also at this point that in light of Assumption 4′, (8.34) implies that the steady-state interest rate is $r^* = f'(k^*) - \delta > n$ (given that $\rho > n$). Since in steady state the wage rate is $w^* = f(k^*) - k^* f'(k^*) < \infty$, it is then straightforward to see that households have finite wealth at all points in time.

Given k^*, the steady-state consumption level is also straightforward to determine as:

$$c^* = f(k^*) - (n + \delta)k^*, \tag{8.35}$$

which is similar to the consumption level in the basic Solow model. Moreover, given Assumption 4′, a steady state where the capital-labor ratio and thus output are constant necessarily satisfies the transversality condition.

This analysis therefore establishes:

PROPOSITION 8.2. *In the neoclassical growth model described above, with Assumptions 1, 2, 3 and 4′, the steady-state equilibrium capital-labor ratio, k^*, is uniquely determined by (8.34) and is independent of the utility function. The steady-state consumption per capita, c^*, is given by (8.35).*

As with the basic Solow growth model, there are also a number of straightforward comparative static results that show how the steady-state values of capital-labor ratio and consumption per capita change with the underlying parameters. For this reason, let us again parameterize the production function as follows

$$f(k) = A\tilde{f}(k),$$

where $A > 0$, so that A is again a shift parameter, with greater values corresponding to greater productivity of factors. Since $f(k)$ satisfies the regularity conditions imposed above, so does $\tilde{f}(k)$.

PROPOSITION 8.3. *Consider the neoclassical growth model described above, with Assumptions 1, 2, 3 and 4′, and suppose that $f(k) = A\tilde{f}(k)$. Denote the steady-state level of the capital-labor ratio by $k^*(A, \rho, n, \delta)$ and the steady-state level of consumption per capita by $c^*(A, \rho, n, \delta)$ when the underlying parameters are A, ρ, n and δ. Then,*

$$\frac{\partial k^*(A,\rho,n,\delta)}{\partial A} > 0,\ \frac{\partial k^*(A,\rho,n,\delta)}{\partial \rho} < 0,\ \frac{\partial k^*(A,\rho,n,\delta)}{\partial n} = 0 \text{ and } \frac{\partial k^*(A,\rho,n,\delta)}{\partial \delta} < 0$$

$$\frac{\partial c^*(A,\rho,n,\delta)}{\partial A} > 0,\ \frac{\partial c^*(A,\rho,n,\delta)}{\partial \rho} < 0,\ \frac{\partial c^*(A,\rho,n,\delta)}{\partial n} < 0 \text{ and } \frac{\partial c^*(A,\rho,n,\delta)}{\partial \delta} < 0.$$

PROOF. See Exercise 8.14. □

The new results here relative to the basic Solow model concern the comparative statics with respect the discount factor. In particular, instead of the saving rate, it is now the discount factor that affects the rate of capital accumulation. There is a close link between the discount rate in the neoclassical growth model and the saving rate in the Solow model.

Loosely speaking, a lower discount rate implies greater patience and thus greater savings. In the model without technological progress, the steady-state saving rate can be computed as

$$s^* = \frac{\delta k^*}{f(k^*)}. \tag{8.36}$$

Exercise 8.16 looks at the relationship between the discount rate, the saving rate and the steady-state per capita consumption level.

Another interesting result is that the rate of population growth has no impact on the steady state capital-labor ratio, which contrasts with the basic Solow model. Exercise 8.13 shows that this result depends on the way in which intertemporal discounting takes place. Another important result, which is more general, is that k^* and thus c^* do *not* depend on the instantaneous utility function $u(\cdot)$. The form of the utility function only affects the transitional dynamics, but has no impact on steady states. This is because the steady state is determined by the modified golden rule. This result will not be true when there is technological change.

8.5. Transitional Dynamics

Recall that transitional dynamics in the basic Solow model were given by a single differential equation with an initial condition. This is no longer the case, since the equilibrium is determined by two differential equations, repeated here for convenience:

$$\dot{k}(t) = f(k(t)) - (n+\delta)k(t) - c(t) \tag{8.37}$$

and

$$\frac{\dot{c}(t)}{c(t)} = \frac{1}{\varepsilon_u(c(t))}\left(f'(k(t)) - \delta - \rho\right). \tag{8.38}$$

Moreover, we have an initial condition $k(0) > 0$, also a boundary condition at infinity, of the form

$$\lim_{t\to\infty}\left[k(t)\exp\left(-\int_0^t \left(f'(k(s)) - \delta - n\right)ds\right)\right] = 0. \tag{8.39}$$

As already discussed in the context of the q-theory of investment, this combination of boundary conditions consisting of an initial condition and a transversality condition is quite typical for economic optimal control problems involving the behavior of both state and control variables. The appropriate notion of stability is again that of *saddle-path stability* introduced in Theorems 7.18 and 7.19 (instead of the global stability results in Theorems 2.4 and 2.5). In particular, the consumption level (or equivalently the costate variable μ) is the control variable, and its initial value $c(0)$ (or equivalently $\mu(0)$) is free. It has to adjust so as to satisfy the transversality condition (the boundary condition at infinity). Since $c(0)$ or $\mu(0)$ can jump to any value, we again need that there exists a unique one-dimensional manifold (curve) tending to the steady state. As in the q-theory of investment, if there were more than one paths tending to the steady state, the equilibrium would be indeterminate, since there would be multiple values of $c(0)$ that could be consistent with equilibrium.

Fortunately, the economic forces are such that once we use the correct notion (saddle-path stability), stability is guaranteed and also there exists a unique competitive equilibrium path. In particular, in the neoclassical growth model there exists a one-dimensional manifold (curve) of stable solutions, often referred to as the *stable arm*, tending to the unique steady state. This can be seen in Figure 8.1. The vertical line is the locus of points where $\dot{c}=0$. The reason why the $\dot{c}=0$ locus is just a vertical line is that in view of the consumer Euler equation (8.38), only the unique level of k^* given by (8.34) can keep per capita consumption constant. The inverse U-shaped curve is the locus of points where $\dot{k}=0$ in (8.37). The intersection of these two loci defines the steady state (k^*, c^*). The shape of the $\dot{k}=0$ locus can be understood by analogy to the diagram where we discussed the golden rule in Chapter 2. If the capital stock is too low, steady-state consumption is low, and if the capital stock is too high, then the steady-state consumption is again low. There exists a unique level, k_{gold} that maximizes the steady-state consumption per capita. The $\dot{c}=0$ locus intersects the $\dot{k}=0$ locus always to the left of k_{gold} (see Exercise 8.17). Once these two loci are drawn, the rest of the diagram can be completed by looking at the direction of motion according to the differential equations, (8.37) and (8.38). Given this direction of movements, it is clear that there exists a unique stable arm tending to the steady state. This implies that starting with an initial capital-labor ratio $k(0)>0$, there exists a unique $c(0)$ on the stable arm. If the representative household starts with a per capita consumption level of $c(0)$ at date $t=0$ and then follows the consumption path given by the Euler equation (8.38), consumption per capita and the capital-labor ratio will converge to the unique steady state (k^*, c^*).

Is the path starting with $(k(0), c(0))$ and converging to (k^*, c^*) the unique equilibrium? The answer is yes and there are three ways of seeing this. The first one again involves a direct application of Theorem 7.14. In particular, we have already verified that the concavity conditions of this theorem are satisfied. In addition, we have also shown that for any feasible path $\lim_{t\to\infty}\left[k(t)\exp\left(-\int_0^t (f'(k(s))-\delta-n)\,ds\right)\right]\geq 0$. Therefore, the path starting with $(k(0), c(0))$ and converging to (k^*, c^*), which satisfies (8.37), (8.38) and (8.39), gives a global maximum to the representative household's problem and satisfies market clearing. In addition, no other path will provide the representative household with the same utility (see Exercise 8.11). Thus the path starting with $(k(0), c(0))$ is the unique competitive equilibrium. The same argument applies to the optimal growth problem, which is also described by the same equations. The path starting with $(k(0), c(0))$ is therefore the unique optimal growth path as well.

PROPOSITION 8.4. *In the neoclassical growth model described above, with Assumptions 1, 2, 3 and 4′, there exists a unique equilibrium path starting from any $k(0)>0$ and converging monotonically to the unique steady-state (k^*, c^*) with k^* given by (8.34). Moreover, if $k(0)<k^*$, then $k(t)\uparrow k^*$ and $c(t)\uparrow c^*$, whereas if $k(0)>k^*$, then $k(t)\downarrow k^*$ and $c(t)\downarrow c^*$. This equilibrium path is also identical to the unique optimal growth path.*

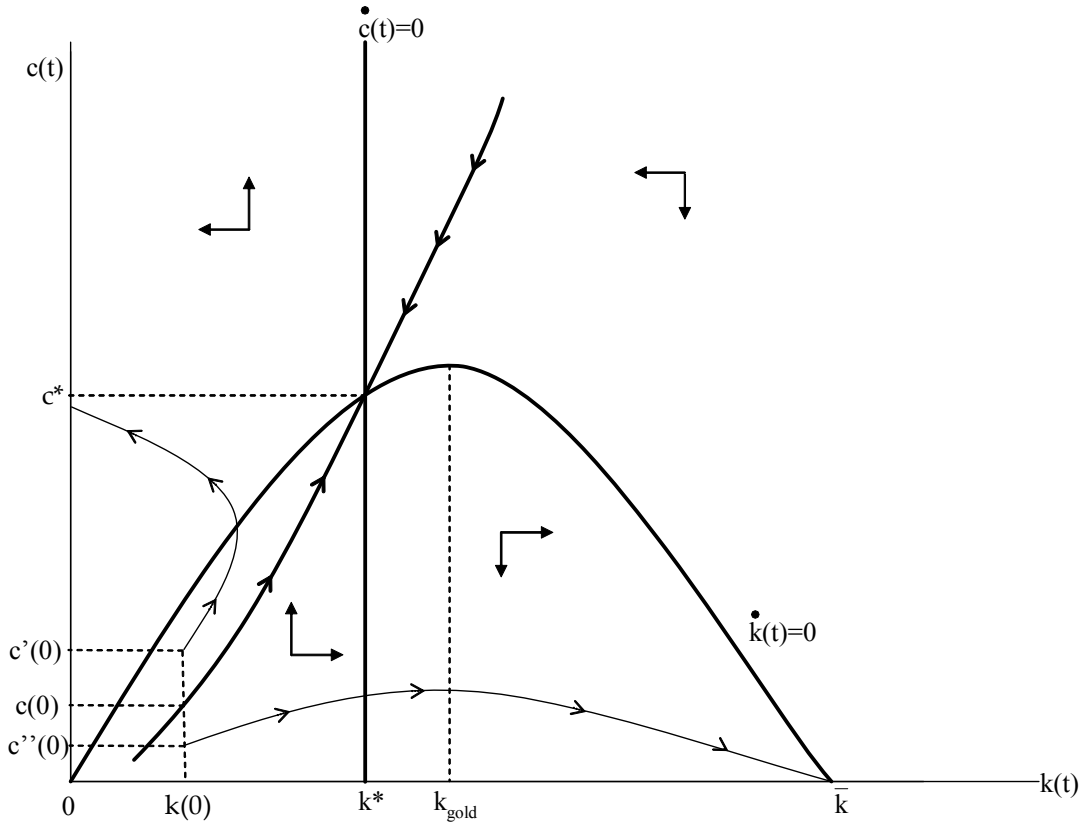

FIGURE 8.1. Transitional dynamics in the baseline neoclassical growth model.

The second strategy is the most popular in the literature and involves ruling out all paths other than the stable arm in Figure 8.1. In particular, Figure 8.1 makes it clear that all points away from the stable arm diverge, and eventually reach zero consumption or zero capital stock. For instance, if the initial level of consumption were below the stable arm, for example, at $c''(0)$, then consumption would reach zero in finite time, thus capital would accumulate continuously until the maximum level of capital (reached with zero consumption), $\bar{k} > k_{gold}$. It can be verified that $f'\left(\bar{k}\right) < \delta + n$ (see Exercise 8.12), which implies that

$$\lim_{t\to\infty}\left[k(t)\exp\left(-\int_0^t\left(f'\left(k(s)\right)-\delta-n\right)ds\right)\right]=\bar{k}\lim_{t\to\infty}\left[\exp\left(-\int_0^t\left(f'\left(\bar{k}\right)-\delta-n\right)ds\right)\right]>0,$$

violating the transversality condition, (8.28) or (8.39). This implies that paths starting below the stable arm cannot be part of an equilibrium. Next, suppose that initial consumption starts above this stable arm, say at $c'(0)$. In this case, the capital stock would reach 0 in finite time, while household consumption implied by (8.38) would remain positive (see Exercise 8.12).[3] But this violates feasibility and establishes that initial values of consumption above this

[3]There is the same technical problem here as the one pointed out in the context of the q-theory of investment in Section 7.8 in the previous chapter; when c reaches 0, the necessary conditions no longer apply.

stable arm cannot be part of the equilibrium (or the optimal growth solution). This line of argument then also leads to the conclusion that the transitional dynamics in the neoclassical growth model will involve the initial consumption per capita jumping to $c(0)$ on the stable arm, and then (k, c) monotonically travel along this arm towards the steady state.

Finally, the third way of establishing the same result is by linearizing the two differential equations characterizing the equilibrium path, (8.37) and (8.38). Linearizing these equations around the steady state (k^*, c^*) and suppressing time dependence, we obtain

$$\begin{aligned}\dot{k} &= \text{constant} + \left(f'(k^*) - n - \delta\right)(k - k^*) - c \\ \dot{c} &= \text{constant} + \frac{c^* f''(k^*)}{\varepsilon_u(c^*)}(k - k^*).\end{aligned}$$

Moreover, from (8.34), $f'(k^*) - \delta = \rho$, so the eigenvalues of this two-equation system are given by the values of ξ that solve the following quadratic form:

$$\det \begin{pmatrix} \rho - n - \xi & -1 \\ \frac{c^* f''(k^*)}{\varepsilon_u(c^*)} & 0 - \xi \end{pmatrix} = 0.$$

It is straightforward to verify that, since $c^* f''(k^*)/\varepsilon_u(c^*) < 0$, there are two real eigenvalues, one negative and one positive. This implies that there exists a one-dimensional stable curve converging to the steady state, exactly as the stable arm in the above figure (see Exercise 8.21). Therefore, not surisingly, the local analysis leads to the same conclusion as the global analysis. However, the local analysis can only establish local stability, whereas the above analysis established global stability.

8.6. Neoclassical Growth in Discrete Time

It is useful to briefly discuss the baseline neoclassical growth model in discrete time to show the similarity of both the mechanics and the insights to the continuous-time analysis. Discrete time equilibrium growth models will be discussed in greater detail in Chapter 17, when we introduce uncertainty.

For now let us suppose that there is no population growth, so that $c(t)$ denotes per capita consumption, the representative household inelastically supplies one unit of labor, and as usual, $\beta \in (0, 1)$ is the discount factor. The representative household then maximizes

$$U(0) \equiv \sum_{t=0}^{\infty} \beta^t u(c(t))$$

subject to the budget constraint

$$a(t+1) = (1 + r(t))\, a(t) + w(t) - c(t),$$

where $a(t)$ is the asset holdings of the individual at time t, $w(t)$ is the equilibrium wage rate, which is also equal to the labor income of the representative household, who supplies one unit of labor to the market, and $r(t)$ is the rate of return on assets holdings at time t. As in the continuous-time model, this flow budget constraint needs to be augmented with a

no-Ponzi condition. With a similar reasoning to that which led to (8.14) in Section 8.1, this condition takes the form

$$\lim_{t\to\infty}\left[a(t)\prod_{s=1}^{t-1}\frac{1}{1+r(s)}\right]\geq 0, \tag{8.40}$$

and ensures that the present discounted value of the representative household's asymptotic debt is nonnegative (see Exercise 8.23). The representative household's transversality condition then implies that $a(t)$ cannot limit to a negative value, therefore the stronger form of the no-Ponzi condition, which must hold an equilibrium, is

$$\lim_{t\to\infty}\left[a(t)\prod_{s=1}^{t-1}\frac{1}{1+r(s)}\right]=0. \tag{8.41}$$

The production side of the economy is identical to that in the continuous-time model. Specifically, the rental rate of capital $R(t)$ and the wage rate $w(t)$ are given by (8.5) and (8.6). Moreover, given depreciation at the geometric rate $\delta>0$ in discrete time, the rate of return on assets, $r(t)$, is again given by (8.26). Next, straightforward application of the results in Chapter 6 implies that the representative household will choose a consumption path that satisfies the Euler equation

$$u'(c(t))=\beta(1+r(t+1))u'(c(t+1)). \tag{8.42}$$

The reader will recall from the analysis of optimal growth in discrete time in Section 6.8 in Chapter 6 that this is identical to the Euler equation for the optimal growth problem (since from (8.26) $r(t)=f'(k(t))-\delta$). To establish the equivalence of the competitive equilibrium and the optimal growth paths, we only need to show that the no-Ponzi condition (8.41) implies the transversality condition in the optimal growth problem (6.48) and vice versa. To see this, let us again use (8.26) to rewrite (6.48) as

$$\lim_{t\to\infty}\left[\beta^t\left(f'(k(t))+(1-\delta)\right)u'(c(t))k(t)\right]=0.$$

Next recursively substituting from (8.42) for t, $t-1$, and so on, and using (8.26), this is equivalent to

$$\begin{aligned}
&\lim_{t\to\infty}\left[\beta^{t-1}u'(c(t-1))k(t)\right]=0\\
&\lim_{t\to\infty}\left[\beta^{t-2}\frac{1}{1+r(t-1)}u'(c(t-2))k(t)\right]=0\\
&\lim_{t\to\infty}\left[\beta^{t-3}\frac{1}{(1+r(t-1))(1+r(t-2))}u'(c(t-3))k(t)\right]=0\\
&.....\\
&\lim_{t\to\infty}\left[k(t)\prod_{s=0}^{t-1}\frac{1}{1+r(s)}\right]=0,
\end{aligned}$$

where the last line cancels out $u'(c(0))$, which is strictly positive by assumption. Since from market clearing $a(t)=k(t)$, this condition is the same as (8.41) and thus establishes that the competitive equilibrium and optimal growth paths coincide. Naturally, this result is not

surprising, since the neoclassical growth model satisfies the conditions of the First and Second Welfare Theorems, Theorems 5.6 and 5.7. Nevertheless, the explicit derivation here shows how this equivalence manifests itself in the context of this workhorse model of growth theory and macroeconomics.

Given the equivalence between the equilibrium and optimal growth paths, Proposition 6.3 from Section 6.8 in Chapter 6 also characterizes the dynamics of the equilibrium allocation. In particular, this proposition implies that starting with any initial level of capital stock $k(0) > 0$, the competitive equilibrium path of the neoclassical growth model will monotonically converge to the unique steady-state allocation. This result therefore highlights the parallel between the general insights from the discrete-time and continuous-time models.

8.7. Technological Change and the Canonical Neoclassical Model

The above analysis was for the neoclassical growth model without any technological change. As with the basic Solow model, the neoclassical growth model does not account for the long-run growth experience of the world economy without some type of exogenous technological change. Therefore, the more interesting version of this model is the one that incorporates technological change. I now analyze the neoclassical model with exogenous technological change.

The production function is extended to:

$$Y(t) = F[K(t), A(t) L(t)], \tag{8.43}$$

where

$$A(t) = \exp(gt) A(0).$$

Notice that the production function (8.43) imposes purely labor-augmenting—Harrod-neutral—technological change. This is because Theorem 2.6 in Chapter 2 still applies and implies that balanced growth is only possible with technological change that is asymptotically labor-augmenting. As in that chapter, I simplify the analysis by assuming that technological change is labor-augmenting throughout.

Assumptions 1, 2 and 3 on the production and utility functions are still in effect. Assumption $4'$ will be strengthened further in order to ensure finite discounted utility in the presence of sustained economic growth.

The constant returns to scale feature again enables us to work with normalized variables. Now let us define

$$\begin{aligned} \hat{y}(t) &\equiv \frac{Y(t)}{A(t) L(t)} \\ &= F\left[\frac{K(t)}{A(t) L(t)}, 1\right] \\ &\equiv f(k(t)), \end{aligned}$$

where

$$k(t) \equiv \frac{K(t)}{A(t) L(t)} \tag{8.44}$$

is the effective capital-labor ratio, written incorporating labor-augmenting technology in the denominator. Naturally, this is similar to the way that the effective capital-labor ratio was defined in the basic Solow growth model.

In addition to the assumptions on technology, we also need to impose a further assumption on preferences in order to ensure balanced growth. As in the basic Solow model, balanced growth is defined as a pattern of growth consistent with the *Kaldor facts* of constant capital-output ratio and capital share in national income. These two observations together also imply that the rental rate of return on capital, $R(t)$, has to be constant, which, from (8.10), implies that $r(t)$ has to be constant. Let us again refer to an equilibrium path that satisfies these conditions as a balanced growth path (BGP). Balanced growth also requires that consumption and output grow at a constant rate. The Euler equation implies that

$$\frac{\dot{c}(t)}{c(t)} = \frac{1}{\varepsilon_u(c(t))}(r(t) - \rho).$$

If $r(t) \to r^*$, then $\dot{c}(t)/c(t) \to g_c$ is only possible if $\varepsilon_u(c(t)) \to \varepsilon_u$, that is, if the elasticity of marginal utility of consumption is asymptotically constant. Therefore, balanced growth is only consistent with utility functions that have asymptotically constant elasticity of marginal utility of consumption. Since this result is important, I state it as a proposition:

Proposition 8.5. *Balanced growth in the neoclassical model requires that asymptotically (as $t \to \infty$) all technological change is purely labor-augmenting and the elasticity of intertemporal substitution, $\varepsilon_u(c(t))$, tends to a constant ε_u.*

The next example shows the family of utility functions with constant intertemporal elasticity of substitution, which are also those with a constant coefficient of relative risk aversion.

Example 8.1. **(CRRA Utility)** Recall that the Arrow-Pratt coefficient of relative risk aversion for a twice differentiable concave utility function $U(c)$ is

$$\mathcal{R} = -\frac{U''(c)\,c}{U'(c)}.$$

The constant relative risk aversion (CRRA) utility function satisfies the property that $\mathcal{R}$ is constant. Now integrating both sides of the previous equation, setting $\mathcal{R}$ to a constant, implies that the family of CRRA utility functions is given by

$$U(c) = \begin{cases} \frac{c^{1-\theta}-1}{1-\theta} & \text{if } \theta \neq 1 \text{ and } \theta \geq 0 \\ \ln c & \text{if } \theta = 1 \end{cases},$$

with the coefficient of relative risk aversion is given by θ (see Exercise B.9 in Appendix Chapter B for a formal derivation). In writing this expression, I separated the case where $\theta = 1$, since $\left(c^{1-\theta} - 1\right)/(1-\theta)$ is undefined at $\theta = 1$. However, it can be shown that $\ln c$ is indeed the right limit when $\theta \to 1$ (see Exercise 5.4).

With time-separable utility functions, the inverse of the elasticity of intertemporal substitution (defined in eq. (8.29)) and the coefficient of relative risk aversion are identical.

Therefore, the family of CRRA utility functions are also those with constant elasticity of intertemporal substitution.

Now to link this utility function to the Gorman preferences discussed in Chapter 5, let us consider a slightly different problem in which an individual has preferences defined over the consumption of N commodities $\{c_1, ..., c_N\}$ given by

$$U(\{c_1, ..., c_N\}) = \begin{cases} \sum_{j=1}^{N} \frac{c_j^{1-\theta}}{1-\theta} & \text{if } \theta \neq 1 \text{ and } \theta \geq 0 \\ \sum_{j=1}^{N} \ln c_j & \text{if } \theta = 1 \end{cases}. \tag{8.45}$$

Suppose also that this individual faces a price vector $\mathbf{p} = (p_1, ..., p_N)$ and has income y, so that his budget constraint can be expressed as

$$\sum_{j=1}^{N} p_j c_j \leq y. \tag{8.46}$$

Maximizing utility subject to this budget constraint leads to the following indirect utility function

$$v(p,y) = \frac{y^{\frac{\sigma-1}{\sigma}}}{\left[\sum_{j=1}^{N} p_j^{1-\sigma}\right]^{1/\sigma}}$$

(see Exercise 5.6). Although this indirect utility function does not satisfy the Gorman form in Theorem 5.2, a monotonic transformation thereof does (simply raise it to the power $\sigma/(\sigma-1)$).

This establishes that CRRA utility functions are within the Gorman class, and if all individuals have CRRA utility functions, then we can aggregate their preferences and represent them as if they belonged to a single individual.

Now consider a dynamic version of these preferences (defined over infinite horizon):

$$U = \begin{cases} \sum_{t=0}^{\infty} \beta^t \frac{c(t)^{1-\theta}-1}{1-\theta} & \text{if } \theta \neq 1 \text{ and } \theta \geq 0 \\ \sum_{t=0}^{\infty} \beta^t \ln c(t) & \text{if } \theta = 1 \end{cases}.$$

The important feature of these preferences in growth theory is not that the coefficient of relative risk aversion is constant, but that the intertemporal elasticity of substitution is constant (since most of growth models do not feature uncertainty, but involve intertemporal allocation of consumption). However, as illustrated in Exercise 5.2 in Chapter 5, with time-separable utility functions the coefficient of relative risk aversion and the inverse of the intertemporal elasticity of substitution are identical. The intertemporal elasticity of substitution is particularly important in growth models because it regulates how willing individuals are to substitute consumption over time, thus their savings and consumption behavior. In view of this, it may be more appropriate to refer to CRRA preferences as "constant intertemporal elasticity of substitution" preferences. Nevertheless, we follow the standard convention in the literature and stick to the term CRRA.

Given the restriction that balanced growth is only possible with preferences featuring a constant elasticity of intertemporal substitution, let us start with the CRRA instantaneous

utility function

$$u\left(c\left(t\right)\right)=\begin{cases}\frac{c(t)^{1-\theta}-1}{1-\theta} & \text{if } \theta\neq 1 \text{ and } \theta\geq 0\\ \ln c(t) & \text{if } \theta=1\end{cases},$$

where the elasticity of marginal utility of consumption, ε_u, is given by the constant θ. When $\theta=0$, these represent linear preferences, whereas when $\theta=1$, they correspond to log preferences. As $\theta\to\infty$, these preferences become infinitely risk-averse, and infinitely unwilling to substitute consumption over time.

More specifically, the economy now admits a representative household with CRRA preferences

$$\int_0^{\infty}\exp\left(-(\rho-n)t\right)\frac{c\left(t\right)^{1-\theta}-1}{1-\theta}dt, \tag{8.47}$$

where $c\left(t\right)\equiv C\left(t\right)/L\left(t\right)$ is per capita consumption. I refer to this model, with labor-augmenting technological change and CRRA preference as given by (8.47) as the *canonical model*, since it is the model used in almost all applications of the neoclassical growth model. In this model, the representative household's problem is given by the maximization of (8.47) subject (8.8) and (8.14). Once again using the necessary conditions from Theorem 7.13, the Euler equation of the representative household is obtained as

$$\frac{\dot{c}\left(t\right)}{c\left(t\right)}=\frac{1}{\theta}\left(r\left(t\right)-\rho\right). \tag{8.48}$$

Let us first characterize the steady-state equilibrium in this model with technological progress. Since with technological progress there will be growth in per capita income, $c\left(t\right)$ will grow. Instead, in analogy with $k\left(t\right)$, let us define

$$\begin{aligned}\tilde{c}\left(t\right) &\equiv \frac{C\left(t\right)}{A\left(t\right)L\left(t\right)}\\ &\equiv \frac{c\left(t\right)}{A\left(t\right)}.\end{aligned}$$

This normalized consumption level will remain constant along the BGP. In particular,

$$\begin{aligned}\frac{\dot{\tilde{c}}\left(t\right)}{\tilde{c}\left(t\right)} &\equiv \frac{\dot{c}\left(t\right)}{c\left(t\right)}-g\\ &= \frac{1}{\theta}\left(r\left(t\right)-\rho-\theta g\right).\end{aligned}$$

Moreover, for the accumulation of capital stock,

$$\dot{k}\left(t\right)=f\left(k\left(t\right)\right)-\tilde{c}\left(t\right)-\left(n+g+\delta\right)k\left(t\right),$$

where recall that $k\left(t\right)\equiv K\left(t\right)/A\left(t\right)L\left(t\right)$.

The transversality condition, in turn, can be expressed as

$$\lim_{t\to\infty}\left[k\left(t\right)\exp\left(-\int_0^t\left[f'\left(k\left(s\right)\right)-g-\delta-n\right]ds\right)\right]=0. \tag{8.49}$$

In addition, the equilibrium interest rate, $r(t)$, is still given by (8.26). Moreover, since in steady state (BGP) $\tilde{c}(t)$ must remain constant, $r(t) = \rho + \theta g$,which implies that

$$f'(k^*) = \rho + \delta + \theta g. \tag{8.50}$$

This equation pins down the steady-state value of the normalized capital ratio k^* uniquely, in a way similar to the model without technological progress. The level of normalized consumption is then given by

$$\tilde{c}^* = f(k^*) - (n + g + \delta) k^*, \tag{8.51}$$

while per capita consumption grows at the rate g.

The only additional complication in this case is that because there is now sustained growth, the transversality condition becomes more demanding. In particular, substituting (8.50) into (8.49),

$$\lim_{t \to \infty} \left[k(t) \exp\left(- \int_0^t [\rho - (1-\theta) g - n] \, ds \right) \right] = 0,$$

which can only hold if the integral within the exponent goes to zero, that is, only if $\rho - (1-\theta) g - n > 0$. This implies that to ensure a well-defined solution to the household maximization problem and a well-defined competitive equilibrium, we need to modify Assumption 4′ to:

ASSUMPTION 4. ***(Discounting with Technological Progress)***

$$\rho - n > (1-\theta) g.$$

Note that this assumption strengthens Assumption 4′ when $\theta < 1$. Recall that in steady state $r = \rho + \theta g$ and the growth rate of output is $g + n$. Therefore, Assumption 4 is equivalent to requiring that $r > g + n$. Assumption 4 will emerge as a necessary condition for the transversality condition to be satisfied in many different models. It is also related to issues of "dynamic efficiency" discussed in the next chapter.

At this point, we can also use a similar reasoning to that in Section 8.2 and establish that, given Assumption 4, the sufficiency conditions in Theorem 7.14 are satisfied, so that the solution to the household maximization problem derived above indeed corresponds to a global maximum (see Exercise 8.18). In this light, the following is an immediate generalization of Proposition 8.2:

PROPOSITION 8.6. *Consider the neoclassical growth model with labor-augmenting technological progress at the rate g and preferences given by (8.47). Suppose that Assumptions 1, 2, 3 and 4 hold. Then, there exists a unique BGP with a normalized capital to effective labor ratio of k^*, given by (8.50), and output per capita and consumption per capita grow at the rate g.*

The steady-state (BGP) capital-labor ratio is no longer independent of the utility function of the representative household, since now the steady-state capital-labor ratio, k^* given by

(8.50), depends on the elasticity of marginal utility (or the inverse of the intertemporal elasticity of substitution), θ. This is because there is now positive growth in output per capita, and thus in consumption per capita. Since individuals face an upward-sloping consumption profile, their willingness to substitute consumption today for consumption tomorrow determines how much they will accumulate and thus the equilibrium effective capital-labor ratio.

Perhaps the most important implication of Proposition 8.6 is that, while the steady-state effective capital-labor ratio, k^*, is determined endogenously, the steady-state growth rate of the economy is given exogenously and is equal to the rate of labor-augmenting technological progress, g. Therefore, the neoclassical growth model, like the basic Solow growth model, endogenizes the capital-labor ratio, but not the growth rate of the economy. The advantage of the neoclassical growth model is that the capital-labor ratio and the equilibrium level of (normalized) output and consumption are determined by the preferences of the individuals rather than an exogenously fixed saving rate. This also enables us to compare equilibrium and optimal growth (and in this case conclude that the competitive equilibrium is Pareto optimal and any Pareto optimum can be decentralized). But the determination of the rate of growth of the economy is still outside the scope of analysis.

A similar analysis to before also leads to a generalization of Proposition 8.4.

PROPOSITION 8.7. *Consider the neoclassical growth model with labor-augmenting technological progress at the rate g and preferences given by (8.47). Suppose that Assumptions 1, 2, 3 and 4 hold. Then, there exists a unique equilibrium path of normalized capital and consumption, $(k(t), \tilde{c}(t))$ converging monotonically to the unique steady-state $(k^*, \tilde{c}^*)$ with k^* given by (8.50). Moreover, if $k(0) < k^*$, then $k(t) \uparrow k^*$ and $\tilde{c}(t) \uparrow \tilde{c}^*$, whereas if $k(0) > k^*$, then $k(t) \downarrow k^*$ and $\tilde{c}(t) \downarrow \tilde{c}^*$.*

PROOF. See Exercise 8.19. □

It is also useful to briefly look at an example with Cobb-Douglas technology.

EXAMPLE 8.2. Consider the model with CRRA utility and labor-augmenting technological progress at the rate g. Assume that the production function is given by $F(K, AL) = K^{\alpha}(AL)^{1-\alpha}$, so that

$$f(k) = k^{\alpha},$$

and thus $r = \alpha k^{\alpha-1} - \delta$. In this case, suppressing time dependence to simplify notation, the Euler equation written in terms of normalized consumption becomes

$$\frac{\dot{\tilde{c}}}{\tilde{c}} = \frac{1}{\theta}\left(\alpha k^{\alpha-1} - \delta - \rho - \theta g\right),$$

and the accumulation equation can be written as

$$\frac{\dot{k}}{k} = k^{\alpha-1} - \delta - g - n - \frac{\tilde{c}}{k}.$$

Now define $z \equiv \tilde{c}/k$ and $x \equiv k^{\alpha-1}$, which implies that $\dot{x}/x = (\alpha - 1)\dot{k}/k$. Therefore,

$$\frac{\dot{x}}{x} = -(1-\alpha)(x - \delta - g - n - z), \tag{8.52}$$

and also

$$\frac{\dot{z}}{z} = \frac{\dot{\tilde{c}}(t)}{\tilde{c}(t)} - \frac{\dot{k}}{k},$$

which implies that

$$\begin{aligned} \frac{\dot{z}}{z} &= \frac{1}{\theta}(\alpha x - \delta - \rho - \theta g) - x + \delta + g + n + z \\ &= \frac{1}{\theta}((\alpha - \theta)x - (1-\theta)\delta + \theta n) - \frac{\rho}{\theta} + z. \end{aligned} \tag{8.53}$$

The two differential equations (8.52) and (8.53) together with the initial condition $x(0)$ and the transversality condition completely determine the dynamics of the system. In Exercise 8.22, you are asked to complete this example for the special case in which $\theta \to 1$ (log preferences).

8.8. The Role of Policy

In the above model, the rate of growth of per capita consumption and output per worker (per capita) are determined exogenously, by the growth rate of labor-augmenting technological progress. The level of income, on the other hand, depends on the intertemporal elasticity of substitution, $1/\theta$, the discount rate, ρ, the depreciation rate, δ, the population growth rate, n, and naturally the form of the production function $f(\cdot)$.

Returning to the proximate causes of cross-country differences in income per capita and growth, this model would give us a way of understanding those differences only in terms of preference and technology parameters. As already discussed in Chapter 4, we would also like to link the proximate causes of economic growth to potential fundamental causes. The intertemporal elasticity of substitution and the discount rate can be viewed as potential determinants of economic growth related to cultural or geographic factors. However, an explanation for cross-country and over-time differences in economic growth based on differences or changes in preferences is unlikely to be satisfactory. A more appealing direction may be to link the incentives to accumulate physical capital (and later to accumulate human capital and technology) to the institutional environment of an economy. I discuss how institutions might affect various investment decisions in Part 8. For now, it is useful to focus on a particularly simple way in which institutional differences might affect investment decisions, which is through differences in policies. To do this, let us extend the above framework in a simple way and introduce linear tax policy. Suppose that returns on capital net of depreciation are taxed at the rate τ and the proceeds of this are redistributed back to the consumers. In that case, the capital accumulation equation, in terms of normalized capital, still remains as above:

$$\dot{k}(t) = f(k(t)) - \tilde{c}(t) - (n + g + \delta)k(t),$$

but the net interest rate faced by households now changes to:

$$r(t) = (1-\tau)\left(f'(k(t)) - \delta\right),$$

because of the taxation of capital returns. The growth rate of normalized consumption is then obtained from the consumer Euler equation, (8.48), as

$$\begin{aligned}\frac{\dot{\tilde{c}}(t)}{\tilde{c}(t)} &= \frac{1}{\theta}\left(r(t)-\rho-\theta g\right)\\ &= \frac{1}{\theta}\left(\left(1-\tau\right)\left(f'\left(k\left(t\right)\right)-\delta\right)-\rho-\theta g\right).\end{aligned}$$

An identical argument to that above immediately implies that the steady-state capital to effective labor ratio is given by

$$f'\left(k^*\right)=\delta+\frac{\rho+\theta g}{1-\tau}. \tag{8.54}$$

This equation shows the effects of taxes on steady-state capital to effective labor ratio and output per capita. A higher tax rate τ increases the right-hand side of (8.54), and since from Assumption 1, $f'(\cdot)$ is decreasing, it reduces k^*. Therefore, higher taxes on capital have the effect of depressing capital accumulation and reducing income per capita. This shows one channel through which policy (and thus institutional) differences might affect economic performance. Similar results apply if, instead of being imposed on the returns from capital, taxes were imposed on the amount of investment (see next section). However, I have not so far offered a reason why some countries may tax capital at a higher rate than others, and this is a topic that will be discussed later. Before doing this, the next section discusses how large these effects can be and whether they could account for the differences in cross-country incomes.

8.9. Comparative Dynamics

Comparative dynamics are a little different in the neoclassical growth model than those in the basic Solow model. Recall that while comparative statics refer to changes in steady state in response to changes in parameters, comparative dynamics look at how the entire equilibrium path of variables changes in response to a change in policy or parameters. Since the purpose here is to give a sense of how these results are different, I briefly consider the model with capital taxation discussed in the previous section and consider the effects of a change in the tax rate τ. Suppose that there is population growth at the rate n and labor-augmenting technological progress at the rate g and that capital is taxed at the rate τ. Suppose also that the economy is initially in a steady state represented by $(k^*, \tilde{c}^*)$. Now imagine that the capital tax rate declines from τ to $\tau' < \tau$. How does the equilibrium path change?

We know from the analysis so far that at the new tax rate $\tau' > 0$, there exists a unique steady state equilibrium that is saddle path stable. Let this steady state be denoted by $(k^{**}, \tilde{c}^{**})$. Therefore, the equilibrium will ultimately tend to this new steady-state equilibrium. Moreover, since $\tau' < \tau$, we know that the new steady-state effective capital-labor ratio has to be greater than k^*, that is, $k^{**} > k^*$ (while the equilibrium growth rate will remain unchanged). Figure 8.2 shows diagrammatically how the comparative dynamics work out

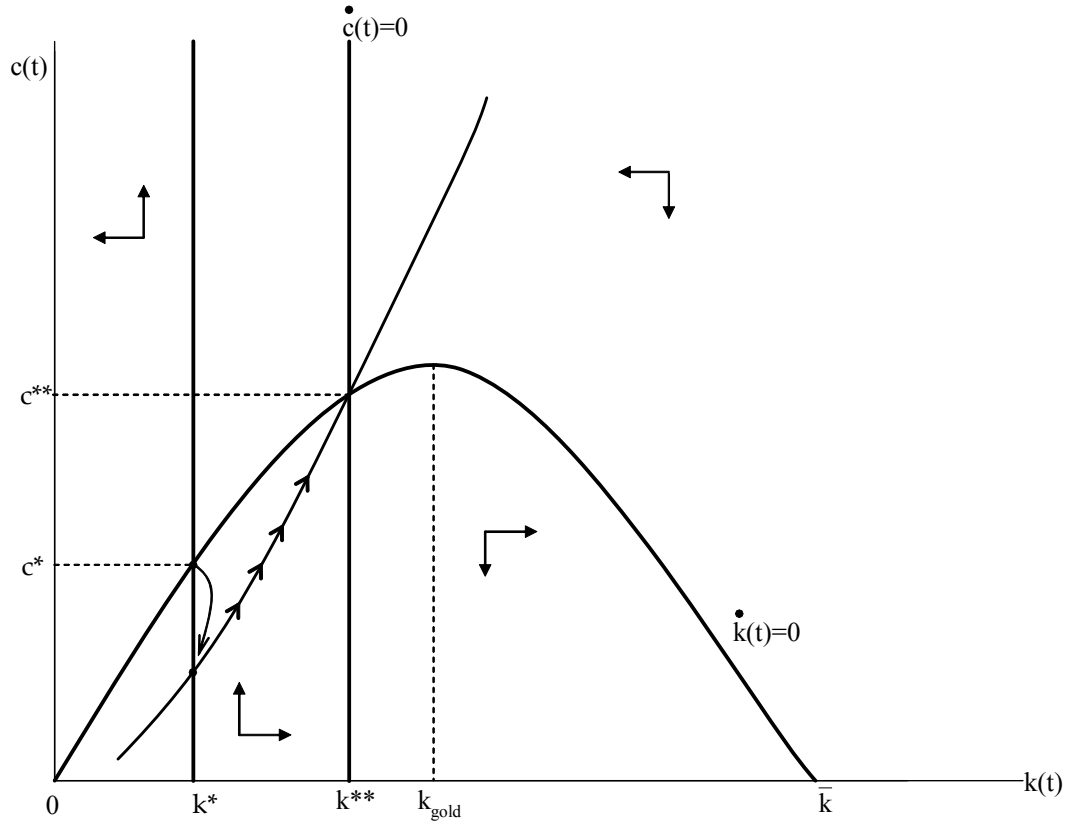

FIGURE 8.2. The dynamic response of capital and consumption to a decline in the tax rate from τ to $\tau' < \tau$.

(using c instead of $\tilde{c}$ to simplify the figure). This figure is drawn under the assumption that the change in the tax rate is unanticipated and occurs at some date T. At this point, the curve corresponding to $\dot{\tilde{c}}(t)/\tilde{c}(t) = 0$ shifts to the right and together with this, the laws of motion represented by the phase diagram change (in the figure, the arrows represent the dynamics of the economy after the change). It can be seen that following this decline in the tax rate, the previous steady-state level of consumption, $\tilde{c}^*$, is above the stable arm of the new dynamical system. Therefore, consumption must drop immediately to reach the new stable arm, so that capital can accumulate towards its new steady-state level. This is shown in the figure with the arc representing the jump in consumption immediately following the decline in the tax rate. Following this initial reaction, consumption slowly increases along the stable arm to a higher level of (normalized) consumption. Therefore, a decline in the tax rate leads to a temporary decline in consumption, associated with a long-run increase in consumption. The overall level of normalized consumption will necessarily increase, since the intersection between the curve for $\dot{\tilde{c}}(t)/\tilde{c}(t) = 0$ and the inverse U-shaped curve for $\dot{k}/k = 0$ will necessarily be to the left side of k_{gold}.

Comparative dynamics in response to changes in other parameters, including the rate of labor-augmenting technological progress g, the rate of population growth n, the discount rate ρ, and other aspects of the utility function, can also be analyzed similarly. Exercise 8.26 asks you to work through the comparative dynamics in response to a change in the rate of labor-augmenting technological progress, g, and in response to an anticipated future change in τ.

8.10. A Quantitative Evaluation

As a final exercise, let us investigate whether the quantitative implications of the neoclassical growth model are reasonable. For this purpose, consider a world consisting of a collection $\mathcal{J}$ of closed neoclassical economies (with all the caveats of ignoring technological, trade and financial linkages across countries, already discussed in Chapter 3; see also Chapter 19). Suppose that each country $j \in \mathcal{J}$ admits a representative household with identical preferences, given by

$$\int_0^{\infty} \exp\left(-\rho t\right) \frac{C_j\left(t\right)^{1-\theta}-1}{1-\theta} dt. \tag{8.55}$$

Let us assume that there is no population growth, so that C_j is both total and per capita consumption. Equation (8.55) imposes that all countries have the same discount rate ρ (see Exercise 8.30).

All countries also have access to the same production technology given by the Cobb-Douglas production function

$$Y_j\left(t\right) = K_j\left(t\right)^{1-\alpha} \left(AH_j\left(t\right)\right)^{\alpha}, \tag{8.56}$$

with H_j representing the exogenously given stock of effective labor (human capital). The accumulation equation is

$$\dot{K}_j\left(t\right) = I_j\left(t\right) - \delta K_j\left(t\right).$$

The only difference across countries is in the budget constraint for the representative household, which takes the form

$$\left(1+\tau_j\right) I_j\left(t\right) + C_j\left(t\right) \leq Y_j\left(t\right), \tag{8.57}$$

where τ_j is the (constant) tax on investment. This tax varies across countries, for example because of policies or differences in institutions/property rights enforcement. Notice that $1+\tau_j$ is also the relative price of investment goods (relative to consumption goods): one unit of consumption goods can only be transformed into $1/\left(1+\tau_j\right)$ units of investment goods.

The right-hand side of the budget constraint (8.57) is still $Y_j\left(t\right)$, which implicitly assumes that $\tau_j I_j\left(t\right)$ is wasted, rather than being redistributed back to the representative household. This is without any major consequence, since, as noted in Theorem 5.2 above, CRRA preferences as in (8.55) have the nice feature that they can be exactly aggregated across individuals, so we do not have to worry about the distribution of income in the economy.

The competitive equilibrium can be characterized as the solution to the maximization of (8.55) subject to (8.57) and the capital accumulation equation. With the same steps as above, the Euler equation of the representative household is

$$\frac{\dot{C}_j(t)}{C_j(t)} = \frac{1}{\theta}\left(\frac{(1-\alpha)}{(1+\tau_j)}\left(\frac{AH_j(t)}{K_j(t)}\right)^{\alpha} - \delta - \rho\right).$$

Consider the steady state. Because A is constant, the steady state requires that $\dot{C}_j(t)/C_j(t) = 0$ for all $j \in \mathcal{J}$. This immediately implies that

$$K_j(t) = \frac{(1-\alpha)^{1/\alpha} AH_j(t)}{\left[(1+\tau_j)(\rho+\delta)\right]^{1/\alpha}}.$$

So countries with higher taxes on investment will have a lower capital stock in steady state. Equivalently, they will also have lower capital per worker, or a lower capital output ratio (using (8.56) the capital output ratio is simply $K/Y = (K/AH)^{\alpha}$).

Now substituting this into (8.56), and comparing two countries with different taxes (but the same human capital), and denoting the steady-state income level of a country with a tax rate equal to τ by $Y(\tau)$, we obtain the relative steady-state income differences as

$$\frac{Y(\tau)}{Y(\tau')} = \left(\frac{1+\tau'}{1+\tau}\right)^{\frac{1-\alpha}{\alpha}}. \tag{8.58}$$

This equation therefore summarizes the intuitively obvious notion that countries that tax investment, either directly or indirectly, at a higher rate will be poorer. More interestingly, this equation can be used for a quantitative evaluation of how large the effects of these types of policies might be. The advantage of using the neoclassical growth model for quantitative evaluation relative to the Solow growth model is that the extent to which different types of distortions (here captured by the tax rates on investment) will affect income and capital accumulation is determined endogenously. In contrast, in the Solow growth model, what matters is the saving rate, so to evaluate the effect of policy on cross-country income differences, we would need to link taxes or distortions to savings (or to other determinants of income per capita such as technology). This can be done by using other sources of evidence to estimate the effect of these distortions on savings, though this is in general not a trivial task.

How large are the effects of tax distortions captured by eq. (8.58)? Put differently, can the neoclassical growth model combined with policy differences account for quantitatively large cross-country income differences?

The advantage of the current approach is its parsimony. As eq. (8.58) shows, only differences in τ across countries and the value of the parameter α matter for this comparison. Recall also that a plausible value for α is $2/3$, since this is the share of labor income in national product which, with Cobb-Douglas production function, is equal to α, so this parameter can be easily mapped to data.

Where can we obtain estimates of differences in τ's across countries? There is no obvious answer to this question. A popular approach in the literature is to obtain estimates of τ from the relative price of investment goods (as compared to consumption goods), motivated

by the fact that in eq. (8.57) τ corresponds to a distortion directly affecting investment expenditures. Therefore, we may expect τ to have an effect on the market price of investment goods. Data from the Penn World tables suggest that there is a large amount of variation in the relative price of investment goods. For example, countries with the highest relative price of investment goods have relative prices almost eight times as high as countries with the lowest relative price.

Then, using $\alpha = 2/3$, eq. (8.58) implies that the income gap between two such countries should be approximately threefolds:

$$\frac{Y(\tau)}{Y(\tau')} \approx 8^{1/2} \approx 3.$$

Therefore, differences in capital-output ratios or capital-labor ratios caused by taxes or tax-type distortions, even by very large differences in taxes or distortions, are unlikely to account for the large differences in income per capita that we observe in practice. This is of course not surprising and parallels our discussion of the Mankiw-Romer-Weil approach in Chapter 3. In particular, recall that the discussion in Chapter 3 showed that differences in income per capita across countries are unlikely to be accounted for by differences in capital per worker alone. Instead, to explain such large differences in income per capita across countries, we need sizable differences in the efficiency with which these factors are used. Such differences do not feature in this model. Therefore, the simplest neoclassical model does not generate sufficient differences in capital-labor ratios to explain cross-country income differences.

Nevertheless, many economists have tried (and still try) to use versions of the neoclassical model to go further. The motivation is simple. If instead of using $\alpha = 2/3$, we take $\alpha = 1/3$, the ratio of incomes in the two countries would be

$$\frac{Y(\tau)}{Y(\tau')} \approx 8^{2} \approx 64.$$

Therefore, if there is any way of increasing the responsiveness of capital or other factors to distortions, the predicted differences across countries can be made much larger. How could we have a model in which $\alpha = 1/3$? Such a model must have additional accumulated factors, while still keeping the share of labor income in national product roughly around 2/3. One possibility might be to include human capital (see Chapter 10 below). Nevertheless, the discussion in Chapter 3 showed that human capital differences appear to be insufficient to explain much of the income per capita differences across countries. Another possibility is to introduce other types of capital or perhaps technology that responds to distortions in the same way as capital. While these are all logically possible, a serious analysis of these issues requires models of endogenous technology, which will be our focus in the next part of the book.

8.11. Extensions

There are many empirically and theoretically relevant extensions of the neoclassical growth model. I do not present these here for the sake of brevity. But the most important

ones are presented as exercises. In particular, Exercise 8.31 endogenizes the labor supply decisions of households by introducing leisure in the utility function. The model presented in this exercise is particularly important, since it corresponds to the version of the neoclassical growth model most often employed in short-run and medium-run macroeconomic analyses. This exercise also shows that further restrictions on the form of the utility function need to be imposed in order to preserve balanced growth in this case. Exercise 8.32 studies government expenditures and taxation further. Exercise 8.34 looks at the behavior of the basic neoclassical growth model with a free capital account, representing borrowing and lending opportunities for the economy at some exogenously given international interest rate r^*. Exercise 8.35 combines the costs of adjustments in investment as in the q-theory with the basic neoclassical model. Finally, Exercise 8.36 looks at a version of the neoclassical model with multiple sectors.

8.12. Taking Stock

This chapter presented arguably the most important model in macroeconomics; the one-sector neoclassical growth model. Recall that our study of the basic models of economic growth started in Chapter 2, with the Solow growth model. We saw that while this model gives a number of important insights, it treats much of the mechanics of economic growth as a "blackbox". Growth can only be generated by technological progress (unless we are in the special AK model without diminishing returns to capital), but technological progress is outside the model. The next important element in determining cross-country differences in income is the saving rate, but in the Solow growth model this was also taken as exogenous. The major contribution of the current chapter has been to open the blackbox of capital accumulation by specifying the preferences of households. Consequently, we can link the saving rates to preferences and also to technology and prices in the economy. Moreover, as Exercise 8.37 shows, the implications of policy on equilibrium quantities are different in the neoclassical model than in the Solow growth model with exogenously given saving rates. Another major advantage of the neoclassical growth model is that because preferences are explicitly specified, equilibrium and optimal growth can be compared.

Perhaps the most important contribution of this model is that it paves the way for further analysis of capital accumulation, human capital and endogenous technological progress, which will be our topic in the next few chapters (starting with human capital in Chapter 10). Therefore, this chapter is the first, and perhaps conceptually the most important, step towards a systematic study of economic growth. It provides us with the mathematical and conceptual tools that are necessary for modeling capital accumulation, human capital accumulation and technological change endogenously.

Did our study of the neoclassical growth model generate new insights about the sources of cross-country income differences and economic growth relative to the Solow growth model? The answer here is largely no. While the current model is an important milestone in our study of the mechanics of economic growth, as with the Solow growth model, the focus is on

the proximate causes of these differences—we are still looking at differences in saving rates, investments in human capital and technology, perhaps as determined by preferences and other dimensions of technology (such as the rate of labor-augmenting technological change). It is therefore important to bear in mind that this model, by itself, does not enable us to answer questions about the fundamental causes of economic growth. What it does, however, is to clarify the nature of the economic decisions so that we are in a better position to ask such questions.

8.13. References and Literature

The neoclassical growth model goes back to Frank Ramsey's (1928) classic treatment and for that reason is often referred to as the "Ramsey model". Ramsey's model is very similar to standard neoclassical growth model, except that it did not feature discounting. Another early optimal growth model was presented by John von Neumann (1935), focusing more on the limiting behavior of the dynamics in a linear model. The current version of the neoclassical growth model is most closely related to the analysis of optimal growth by David Cass (1965) and Tjalling Koopmans (1960, 1965). An excellent discussion of optimal growth is provided in Arrow and Kurz's (1970) volume.

All growth and macroeconomic textbooks cover the neoclassical growth model. Sargent and Ljungqvist (2004, Chapter 14) provides an introductory treatment of the neoclassical growth model in discrete time. Barro and Sala-i-Martin (2004, Chapter 2) provides a detailed treatment focusing on the continuous-time economy. Blanchard and Fisher (1989, Chapter 2) and Romer (2006, Chapter 2) also present the continuous-time version of the neoclassical growth model. These books use the necessary conditions implied by the Maximum Principle, including the strong version of the transversality condition, and characterize utility-maximizing consumption behavior. The implicit assumption is that any other paths of consumption and assets are either less desirable or violate the no-Ponzi condition. As discussed in the previous chapter, some more care is necessary in applying the Maximum Principle in the presence of the no-Ponzi condition and also it is important to verify that the necessary conditions in fact characterize a global optimum. The difficulty is that these necessary conditions do not apply at the boundaries, thus the argument for ruling out paths that reach zero consumption is not fully rigorous. Instead of focusing on necessary conditions, I used the sufficient conditions in Theorem 7.14 from the previous chapter. This approach is both simpler and more powerful.

Ricardian Equivalence discussed in Exercise 8.33 was first proposed by Barro (1974). It is further discussed in Chapter 9.

A systematic quantitative evaluation of the effects of policy differences is provided in Chari, Kehoe and McGrattan (1997). These authors follow Jones (1995) in emphasizing differences in the relative prices of investment goods (compared to consumption goods) in the Penn Worlds tables and interpret these as due to taxes and other distortions. This interpretation is not without any problems. In particular, in the presence of international

trade, these relative price differences will reflect other technological factors or possible factor proportion differences (see Chapter 19, and also Acemoglu and Ventura, 2002, and Hsieh and Klenow, 2006). Parente and Prescott (1994) use an extended version of the neoclassical growth model (where the "stock of technology," which is costly to adopt from the world frontier, is interpreted as a capital good) to perform similar quantitative exercises. Other authors have introduced yet other accumulable factors in order to increase the elasticity of output to distortions (that is, to reduce the α parameter above). Pete Klenow has dubbed these various accumulable factors introduced in the models to increase this elasticity the "mystery capital" to emphasize the fact that while they may help the quantitative match of the neoclassical-type models, they are not directly observable in the data.

8.14. Exercises

Exercise 8.1. Consider the consumption allocation decision of an infinitely-lived household with (a continuum of) $L(t)$ members at time t, with $L(0) = 1$. Suppose that the household has total consumption $C(t)$ to allocate at time t. The household has "utilitarian" preferences with instantaneous utility function $u(c)$ and discount the future at the rate $\rho > 0$.

(1) Show that the problem of the household can be written as

$$\max \int_0^{\infty} \exp(-\rho t) \left[\int_0^{L(t)} u(c_i(t)) \, di \right] dt,$$

subject to

$$\int_0^{L(t)} c_i(t) \, di \leq C(t),$$

and subject to the budget constraint of

$$\dot{\mathcal{A}}(t) = r(t) \mathcal{A}(t) + W(t) - C(t),$$

where i denotes a generic member of the household, $\mathcal{A}(t)$ is the total asset holding of the household, $r(t)$ is the rate of return on assets and $W(t)$ is total labor income.

(2) Show that as long as $u(\cdot)$ is strictly concave, this problem becomes

$$\max \int_0^{\infty} \exp(-(\rho - n)t) \, u(c(t)) \, dt,$$

subject to

$$\dot{a}(t) = (r(t) - n) a(t) + w(t) - c(t),$$

where $w(t) \equiv W(t)/L(t)$ and $a(t) \equiv \mathcal{A}(t)/L(t)$. Provide an intuition for this transformed problem.

Exercise 8.2. Consider the maximization of (8.3) subject to (8.8) without any other constraints.

(1) Show that for any candidate consumption plan $[c(t)]_{t=0}^{\infty}$, there exists another consumption plan $[c'(t)]_{t=0}^{\infty}$ that satisfies the flow budget constraint (8.8), involves $c'(t) > c(t)$ and yields strictly higher utility.

(2) Using the argument in part 1, show that the household will choose asset levels $a(t)$ becoming arbitrarily negative for all t. [Hint: this problem does not have a sequence of consumption and asset levels that reach the maximum value of the objective function; so here you should simply show that $a(t)$ arbitrarily negative for all t approaches the maximum value of the objective function (which may be, but does not need to be, $+\infty$)].

(3) Explain why an allocation with these features would violate feasibility.

Exercise 8.3. Start with the law of motion of the total assets of the households, given by (8.7). Using this, show that if the household starts with $\mathcal{A}(t) = 0$ at time t and chooses zero consumption thereafter, then it will asymptotically generate asset holdings of

$$\bar{W}(t) \equiv \int_t^{\infty} w(s) L(s) \exp\left(-\int_t^s r(z)\, dz\right) ds.$$

Explain why the natural debt limit requires that $\mathcal{A}(t) \geq -\bar{W}(t)$. Using the definition of $\mathcal{A}(t)$ and the fact that $L(t)$ grows at the rate n, derive (8.11). Relate this natural debt limit to its discrete-time analog, (6.41) in Chapter 6.

Exercise 8.4. Show that the relaxed natural debt limit (8.12) implies that the original debt limit (8.11) holds for all t.

Exercise 8.5. Derive (8.8) from (8.13).

Exercise 8.6. Derive (8.15) from (8.14) and (8.18). [Hint: use (8.17) to substitute for $\mu(t)$ in (8.18).

Exercise 8.7. Verify that Theorem 7.13 can be applied to the household maximization problem. In particular, following the same steps as in Section 7.7 in the previous chapter, show that Assumption 7.1 is satisfied when $c(t) \in [\varepsilon, +\infty)$ and when $r(t) \geq n$ for all t. Then, show that the restriction to $c \in [\varepsilon, +\infty)$ has no effect on the results and that any equilibrium must involve $r(t) \geq n$ for all t.

Exercise 8.8. (1) Show that the natural debt limit (8.12) with $\hat{a} > -\infty$ implies no-Ponzi condition (8.14), but (8.14) does not imply (8.12) with $\hat{a} > -\infty$. [Hint: consider an economy where $w(t) = \exp(gt) w(0)$ and $r(t) = r$ for all t].

(2) Show that in the neoclassical growth model with technological progress studied in Section 8.7, it is not possible to apply Theorems 7.13 or 7.14 using the natural debt limit.

Exercise 8.9. Show that $\lim_{t\to\infty} r(t) > \rho$ and $\lim_{t\to\infty} w(t) = w \geq 0$ are sufficient to guarantee that $\hat{a} > -\infty$ in the natural debt limit (8.12).

Exercise 8.10. (1) Show that imposing the natural debt limit with $\hat{a} > -\infty$ leads to the same household behavior as in the maximization problem with the no-Ponzi condition. [Hint: consider the maximization of (8.3) subject to (8.8) and (8.11) with $\hat{a} > -\infty$ and show that this problem satisfies the conditions of Theorems 7.13 and 7.14. Then, derive the necessary conditions and compare them to those in Section 8.2].

(2) Verify that in equilibrium $\hat{a} > -\infty$.

Exercise 8.11. Use an argument similar to that of the proof of Mangasarian's sufficiency result, Theorem 7.5, to show that because the representative household's Hamiltonian in the neoclassical growth model, $H(t, a, c, \mu)$, is concave in (a, c) and strictly concave in c, the path of consumption and assets given by (8.16)-(8.18) is the unique solution to the household's maximization problem.

Exercise 8.12.
(1) In the analysis of transitional dynamics, show that if initial consumption, $c(0)$, started above the stable arm, the capital stock would reach 0 in finite time. Explain why this would violate feasibility. [Hint: be specific on what you assume about the necessary conditions when $c = 0$].
(2) Show that if the initial level of consumption starts below the stable arm, then capital-labor ratio converges to $\bar{k}$, where $f'(\bar{k}) < \delta + n$. [Hint: use the facts that $\bar{k} > k_{gold}$ and that $f'(k_{gold}) = \delta + n$].

Exercise 8.13. Consider a variant of the neoclassical model (with constant population growth at the rate n) in which preferences are given by

$$\max \int_0^{\infty} \exp(-\rho t) u(c(t)) dt,$$

and there is population growth at the constant rate n. How does this affect the equilibrium? How does the transversality condition need to be modified? What is the relationship between the rate of population growth, n, and the steady-state capital labor ratio k^*?

Exercise 8.14. Prove Proposition 8.3.

Exercise 8.15. Explain why the steady state capital-labor ratio k^* does not depend on the form of the utility function without technological progress, but depends on the intertemporal elasticity of substitution when there is positive technological progress.

Exercise 8.16.
(1) Show that the steady-state saving rate s^* defined in (8.36) is decreasing in ρ, so that lower discount rates lead to higher steady-state savings.
(2) Show that in contrast to the Solow model, the saving rate s^* can never be so high that a decline in savings (or an increase in ρ) can raise the steady-state level of consumption per capita.

Exercise 8.17. In the dynamics of the basic neoclassical growth model, depicted in Figure 8.1, prove that the $\dot{c} = 0$ locus intersects the $\dot{k} = 0$ locus always to the left of k_{gold}. Based on this analysis, explain why the modified golden rule capital-labor ratio, k^*, given by (8.34) differs from k_{gold}.

Exercise 8.18. Consider the neoclassical model with technological progress studied in Section 8.7. Show that when Assumption 4 is satisfied, the household maximization problem, that of maximizing (8.47) subject to (8.8) and (8.14) satisfies the sufficiency conditions in Theorem 7.14. What happens if Assumption 4 is not satisfied?

Exercise 8.19. Prove that, as stated in Proposition 8.7, in the neoclassical model with labor-augmenting technological change and the standard assumptions, starting with $k(0) > 0$, there exists a unique equilibrium path where normalized consumption and capital-labor ratio monotonically converge to the BGP. [Hint: use Figure 8.1].

Exercise 8.20. Consider a neoclassical economy, with a representative household with preferences at time $t = 0$:

$$U(0) = \int_0^\infty \exp(-\rho t) \frac{c(t)^{1-\theta} - 1}{1-\theta}.$$

There is no population growth and labor is supplied inelastically. Assume that the aggregate production function is given by $Y(t) = F[A(t) K(t), L(t)]$ where F satisfies the standard assumptions (constant returns to scale, differentiability, Inada conditions).

(1) Define a competitive equilibrium for this economy.
(2) Suppose that $A(t) = A(0)$ for all t and characterize the steady-state equilibrium. Explain why the steady-state capital-labor ratio is independent of θ.
(3) Now assume that $A(t) = \exp(gt) A(0)$, and show that a BGP (with constant capital share in national income and constant and equal rates of growth of output, capital and consumption) exists only if F takes the Cobb-Douglas form, $Y(t) = (A(t) K(t))^\alpha (L(t))^{1-\alpha}$.
(4) Characterize the BGP in the Cobb-Douglas case. Derive the common growth rate of output, capital and consumption. Explain why the (normalized) steady-state capital-labor ratio now depends on θ.

Exercise 8.21. Consider the baseline neoclassical model with no technological progress.

(1) Show that in the neighborhood of the steady state k^*, the law of motion of $k(t) \equiv K(t)/L(t)$ can be written as

$$\log k(t) = \log k^* + \eta_1 \exp(\xi_1 t) + \eta_2 \exp(\xi_2 t),$$

where ξ_1 and ξ_2 are the eigenvalues of the linearized system.
(2) Compute these eigenvalues show that one of them, say ξ_2, is positive.
(3) What does this imply about the value of η_2?
(4) How is the value of η_1 determined?
(5) What determines the speed of adjustment of $k(t)$ towards its steady-state value k^*?

Exercise 8.22. Derive closed-form equations for the solution to the differential equations of transitional dynamics presented in Example 8.2 with log preferences.

Exercise 8.23. Derive eq. (8.40) from the $T-1$-period lifetime budget constraint of the representative household. In particular, write this budget constraint as

$$\begin{aligned} &\sum_{t=0}^{T-1} \left(\prod_{s=0}^{t-1} \frac{1}{1+r(s)} \right) c(t) + \left(\prod_{s=0}^{t-1} \frac{1}{1+r(s)} \right) a(T) \\ \leq &\sum_{t=0}^{T-1} \left(\prod_{s=0}^{t-1} \frac{1}{1+r(s)} \right) w(t) + a(0). \end{aligned}$$

Explain why this is the correct form of the budget constraint.

By taking the limit as $T \to \infty$, show that this constraint will imply the infinite-horizon lifetime budget constraint only if (8.40) is satisfied.

Exercise 8.24. Consider the discrete-time version of the neoclassical growth model. Suppose that the economy admits a representative household with log preferences ($\theta = 1$ in terms of (8.47)) and the production function is Cobb-Douglas. Assume also that $\delta = 1$, so that there is full depreciation. Characterize the steady-state equilibrium and derive a difference equation that explicitly characterizes the behavior of capital stock away from the steady state.

Exercise 8.25. Again in the discrete-time version of the neoclassical growth model, suppose that there is labor-augmenting technological progress at the rate g,

$$A(t+1) = (1+g) A(t).$$

For simplicity, suppose that there is no population growth.

(1) Prove that balanced growth requires preferences to take the CRRA form

$$U(0) = \begin{cases} \sum_{t=0}^{\infty} \beta^t \frac{c(t)^{1-\theta}-1}{1-\theta} & \text{if } \theta \neq 1 \text{ and } \theta \geq 0 \\ \sum_{t=0}^{\infty} \beta^t \ln c(t) & \text{if } \theta = 1 \end{cases}.$$

(2) Assuming this form of preferences, prove that there exists a unique steady-state equilibrium in which effective capital-labor ratio remains constant.

(3) Prove that this steady-state equilibrium is globally stable and convergence to this steady-state starting from a non-steady-state level of effective capital-labor ratio is monotonic.

Exercise 8.26. (1) Analyze the comparative dynamics of the basic model in response to unanticipated increase in the rate of labor-augmenting technological progress will increase to $g' > g$. Does consumption increase or decrease upon impact?

(2) Analyze the comparative dynamics in response to the announcement at time T that at some future date $T' > T$ the tax rate will decline to $\tau' < \tau$. Does consumption increase or decrease at time T. Explain.

Exercise 8.27. Consider the basic neoclassical growth model with technological change and CRRA preferences (8.47). Explain why $\theta > 1$ ensures that the transversality condition is always satisfied.

Exercise 8.28. Consider the basic neoclassical growth model with CRRA preferences, but with consumer heterogeneity in initial asset holdings (you may assume no technological change if you wish). In particular, there is a set $\mathcal{H}$ of household and household $h \in \mathcal{H}$ starts with initial assets $a_h(0)$. Households are otherwise identical.

(1) Characterize the competitive equilibrium of this economy and show that it is identical to the corresponding representative household economy, with the representative household starting with assets $a(0) = |\mathcal{H}|^{-1} \int_{\mathcal{H}} a_h(0)\, dh$, where $|\mathcal{H}|$ is the measure (number) of households in this economy. Interpret this result and relate it to the aggregation theorems discussed in Chapter 5.

(2) Show that if, instead of the natural debt limit or the no-Ponzi condition, we impose $a_h(t) \geq 0$ for all $h \in \mathcal{H}$ and for all t, then a different equilibrium allocation may result. In light of this finding, interpret the appropriability of using the no-borrowing constraint instead of the natural debt limit.

Exercise 8.29. Consider a variant of the neoclassical economy with preferences given by

$$U(0) = \int_0^\infty \exp(-\rho t) \frac{(c(t) - \gamma)^{1-\theta} - 1}{1-\theta}$$

with $\gamma > 0$. There is no population growth. Assume that the production function is given by $Y(t) = F[K(t), A(t) L(t)]$, which satisfies all the standard assumptions and $A(t) = \exp(gt) A(0)$.

(1) Interpret the utility function.
(2) Define the competitive equilibrium for this economy.
(3) Characterize the equilibrium of this economy. Does a BGP with positive growth in consumption exist? Why or why not?
(4) Derive a parameter restriction ensuring that the standard transversality condition is satisfied.
(5) Characterize the transitional dynamics of the economy.

Exercise 8.30. Consider a world consisting of a collection of closed neoclassical economies $\mathcal{J}$. Each $j \in \mathcal{J}$ has access to the same neoclassical production technology and admits a representative household with preferences $(1-\theta)^{-1} \int_0^\infty \exp(-\rho_j t) \left(c_j^{1-\theta} - 1\right) dt$. Characterize the cross-country differences in income per capita in this world economy. What is the effect of the 10% difference in discount factor (e.g., a difference between a discount rate of 0.02 versus 0.022) on steady-state per capita income differences? [Hint: use the fact that the capital share of income is about 1/3].

Exercise 8.31. Consider the standard neoclassical growth model augmented with labor supply decisions. In particular, there is a total population normalized to 1, and all individuals have utility function

$$U(0) = \int_0^\infty \exp(-\rho t) u(c(t), 1 - l(t)),$$

where $l(t) \in (0,1)$ is labor supply. In a symmetric equilibrium, employment $L(t)$ is equal to $l(t)$. Assume that the production function is given by $Y(t) = F[K(t), A(t) L(t)]$, which satisfies all the standard assumptions and $A(t) = \exp(gt) A(0)$.

(1) Define a competitive equilibrium.
(2) Set up the current-value Hamiltonian that each household solves taking wages and interest rates as given, and determine first-order conditions for the allocation of consumption over time and leisure-labor trade off.
(3) Set up the current-value Hamiltonian for a planner maximizing the utility of the representative household, and derive the necessary conditions for an optimal solution.
(4) Show that the two problems are equivalent given competitive markets.
(5) Show that unless the utility function is asymptotically equal to

$$u(c(t), 1 - l(t)) = \begin{cases} \frac{Ac(t)^{1-\theta}}{1-\theta} h(1 - l(t)) & \text{for } \theta \neq 1, \\ A \log c(t) + Bh(1 - l(t)) & \text{for } \theta = 1, \end{cases}$$

for some $h(\cdot)$ with $h'(\cdot) > 0$, there will not exist a BGP with constant and equal rates of consumption and output growth, and a constant level of labor supply.

Exercise 8.32. Consider the standard neoclassical growth model with a representative household with preferences

$$U(0) = \int_0^\infty \exp(-\rho t)\left[\frac{c(t)^{1-\theta}-1}{1-\theta} + G(t)\right],$$

where $G(t)$ is a public good financed by government spending. Assume that the production function is given by $Y(t) = F[K(t), L(t)]$, which satisfies all the standard assumptions, and the budget set of the representative household is $C(t) + I(t) \leq Y(t)$, where $I(t)$ is private investment. Assume that $G(t)$ is financed by taxes on investment. In particular, the capital accumulation equation is

$$\dot{K}(t) = (1-\tau(t))\, I(t) - \delta K(t),$$

and the fraction $\tau(t)$ of the private investment $I(t)$ is used to finance the public good, that is, $G(t) = \tau(t) I(t)$.

Take the process of tax rates $[\tau(t)]_{t=0}^{\infty}$ as given.

(1) Define a competitive equilibrium.
(2) Set up the individual maximization problem and characterize consumption and investment behavior.
(3) Assuming that $\lim_{t\to\infty} \tau(t) = \tau$, characterize the steady state.
(4) What value of τ maximizes the steady-state utility of the representative household. Starting away from the state state, is this also the tax rate that would maximize the initial utility level? Why or why not?

Exercise 8.33. Consider the neoclassical growth model with a government that needs to finance a flow expenditure of G. Suppose that government spending does not affect utility and that the government can finance this expenditure by using lump-sum taxes (that is, some amount $\mathcal{T}(t)$ imposed on each household at time t regardless of their income level and capital holdings) and debt, so that the government budget constraint takes the form

$$\dot{b}(t) = r(t)\, b(t) + g - \mathcal{T}(t),$$

where $b(t)$ denotes its debt level. The no-Ponzi condition for the government is

$$\lim_{t\to\infty}\left[b(t)\exp\left(-\int_0^t (r(s)-n)\, ds\right)\right] = 0.$$

Prove the following *Ricardian equivalence* result: any process of lump-sum taxes $[\mathcal{T}(t)]_{t=0}^{\infty}$ that satisfy the government's budget constraint (together with the no-Ponzi condition) leads to the same equilibrium process of capital-labor ratio and consumption. Interpret this result.

Exercise 8.34. Consider the baseline neoclassical growth model with no population growth and no technological change, and preferences given by the standard CRRA utility function (8.47). Assume, however, that the representative household can borrow and lend at the exogenously given international interest rate r^*. Characterize the steady state equilibrium

and transitional dynamics in this economy. Show that if the economy starts with less capital than its steady state level it will immediately jump to the steady state level by borrowing internationally. How will the economy repay this debt?

Exercise 8.35. Modify the neoclassical economy (without technological change) by introducing cost of adjustment in investment as in the q-theory of investment studied in the previous chapter. Characterize the steady-state equilibrium and the transitional dynamics. How do the implications of this model differ from those of the baseline neoclassical model?

Exercise 8.36. * Consider a version of the neoclassical model that admits a representative household with preferences given by (8.47), no population growth and no technological progress. The main difference from the standard model is that there are multiple capital goods. In particular, suppose that the production function of the economy is given by

$$Y(t) = F(K_1(t), ..., K_M(t), L(t)),$$

where K_m denotes the m^{th} type of capital and L is labor. F is homogeneous of degree 1 in all of its variables. Capital in each sector accumulates in the standard fashion, with

$$\dot{K}_m(t) = I_m(t) - \delta_m K_m(t),$$

for $m = 1, ..., M$. The resource constraint of the economy is

$$C(t) + \sum_{m=1}^{M} I_m(t) \leq Y(t)$$

for all t.

(1) Write budget constraint of the representative household in this economy. Show that this can be done in two alternative and equivalent ways; first, with M separate assets, and second with only a single asset that is a claim to all of the capital in the economy.
(2) Define equilibrium and BGP allocations.
(3) Characterize the BGP by specifying the profit-maximizing decision of firms in each sector and the dynamic optimization problem of consumers.
(4) Write down the optimal growth problem in the form of a multidimensional current-value Hamiltonian and show that the optimum growth problem coincides with the equilibrium growth problem. Interpret this result.
(5) Characterize the transitional dynamics in this economy. Define and discuss the appropriate notion of saddle-path stability and show that the equilibrium is always saddle-path stable and the equilibrium dynamics can be reduced to those in the one-sector neoclassical growth model.
(6) Characterize the transitional dynamics under the additional assumption that investment is irreversible in each sector, i.e., $I_m(t) \geq 0$ for all t and each $m = 1, ..., M$.

Exercise 8.37. Contrast the effects of taxing capital income at the rate τ in the Solow growth model and the neoclassical growth model. Show that capital income taxes have no

effect in the former, while they depress the effective capital-labor ratio in the latter. Explain why there is such a difference.

CHAPTER 9

Growth with Overlapping Generations

A key feature of the neoclassical growth model analyzed in the previous chapter is that it admits a representative household. This model is useful as it provides us with a tractable framework for the analysis of capital accumulation. Moreover, it enables us to appeal to the First and Second Welfare Theorems to establish the equivalence between equilibrium and optimum growth problems. In many situations, however, the assumption of a representative household is not appropriate. One important set of circumstances where we may wish to depart from this assumption is in the analysis of an economy in which new households arrive (or are born) over time. The arrival of new households in the economy is not only a realistic feature, but it also introduces a range of new economic interactions. In particular, decisions made by older "generations" will affect the prices faced by younger "generations". These economic interactions have no counterpart in the neoclassical growth model. They are most succinctly captured in the *overlapping generations models* introduced and studied by Paul Samuelson and then later Peter Diamond.

These models are useful for a number of reasons; first, they capture the potential interaction of different generations of individuals in the marketplace; second they provide a tractable alternative to the infinite-horizon representative agent models; third, some of their key implications are different from those of the neoclassical growth model; fourth, the dynamics of capital accumulation and consumption in some special cases of these models will be quite similar to the basic Solow model rather than the neoclassical model; and finally these models generate new insights about the role of national debt and Social Security in the economy.

I start with an illustration of why the First Welfare Theorem cannot be applied in overlapping generations models. I then discuss the baseline overlapping generations model and and a number of applications of this framework. Finally, I present the overlapping generations model in continuous time. This latter model, originally developed by Menahem Yaari and Olivier Blanchard and also referred to as the *perpetual youth* model, is a tractable alternative to the basic overlapping generations model and also has a number of different implications. It will also be used in the context of human capital investments in the next chapter.

9.1. Problems of Infinity

This section illustrates why the First Welfare Theorem does not apply to overlapping generations models using an abstract general equilibrium economy introduced by Karl Shell. This model is interesting in part because it is closely related to the baseline overlapping generations model of Samuelson and Diamond, which will be presented in the next section.

Consider the following static economy with a countably infinite number of households, each denoted by $i \in \mathbb{N}$, and a countably infinite number of commodities, denoted by $j \in \mathbb{N}$. Assume that all households behave competitively (alternatively, we can assume that there are M households of each type, and M is a large number). Household i has preferences given by:

$$u_i = c_i^i + c_{i+1}^i,$$

where c_j^i denotes the consumption of the jth type of commodity by household i. These preferences imply that household i enjoys the consumption of the commodity with the same index as its own index and the next indexed commodity (i.e., if an individual's index is 3, she only derives utility from the consumption of goods indexed 3 and 4, and so on).

The endowment vector $\boldsymbol{\omega}$ of the economy is as follows: each household has one unit endowment of the commodity with the same index as its index. Let us choose the price of the first commodity as the numeraire, so that $p_0 = 1$. A competitive equilibrium in this economy is defined in the usual manner (e.g., Definition 5.1 in Chapter 5

The following proposition characterizes a competitive equilibrium. Exercise 9.1 asks you to prove that this is the unique competitive equilibrium in this economy.

Proposition 9.1. *In the above-described economy, the price vector $\bar{\mathbf{p}}$ such that $\bar{p}_j = 1$ for all $j \in \mathbb{N}$ is a competitive equilibrium price vector and induces an equilibrium with no trade, denoted by $\bar{\mathbf{x}}$.*

Proof. At $\bar{\mathbf{p}}$, each household has income equal to 1. Therefore, the budget constraint of household i can be written as

$$c_i^i + c_{i+1}^i \leq 1.$$

This implies that consuming own endowment is optimal for each household, establishing that the price vector $\bar{\mathbf{p}}$ and no trade, $\bar{\mathbf{x}}$, constitute a competitive equilibrium. □

However, the competitive equilibrium in Proposition 9.1 is not Pareto optimal. To see this, consider the following alternative allocation, $\tilde{\mathbf{x}}$. Household $i = 0$ consumes one unit of good $j = 0$ and one unit of good $j = 1$, and household $i > 0$ consumes one unit of good $i + 1$. In other words, household $i = 0$ consumes its own endowment and that of household 1, while all other households, indexed $i > 0$, consume the endowment of the neighboring household, $i + 1$. In this allocation, all households with $i > 0$ are as well off as in the competitive equilibrium $(\bar{\mathbf{p}}, \bar{\mathbf{x}})$, and individual $i = 0$ is strictly better-off. This establishes:

Proposition 9.2. *In the above-described economy, the competitive equilibrium at $(\bar{\mathbf{p}}, \bar{\mathbf{x}})$ is not Pareto optimal.*

So why does the First Welfare Theorem not apply in this economy? Recall that the first version of this theorem, Theorem 5.5, was stated under the assumption of a finite number of commodities, whereas there are an infinite number of commodities here. Clearly, the source of the problem must be related to the infinite number of commodities. The extended version of the First Welfare Theorem, Theorem 5.6, covers the case with an infinite number of

commodities, but only under the assumption that $\sum_{j=0}^{\infty} p_j^* < \infty$, where p_j^* refers to the price of commodity j in the competitive equilibrium in question. It can be immediately verified that this assumption is not satisfied in the current example, since the competitive equilibrium in question features $p_j^* = 1$ for all $j \in \mathbb{N}$, so that $\sum_{j=0}^{\infty} p_j^* = \infty$. As discussed in Chapter 5, when the sum of prices is equal to infinity, there can be feasible allocations for the economy as a whole that Pareto dominate the competitive equilibrium. The economy discussed here gives a simple example where this happens.

If the failure of the First Welfare Theorem were a specific feature of this abstract (perhaps artificial) economy, it would not have been of great interest to us. However, this abstract economy is very similar (in fact "isomorphic") to the baseline overlapping generations model. Therefore, the source of inefficiency (Pareto suboptimality of the competitive equilibrium) in this economy will be the source of potential inefficiencies in the baseline overlapping generations model.

It is also useful to recall that, in contrast to Theorem 5.6, the Second Welfare Theorem, Theorem 5.7, did not make use of the assumption that $\sum_{j=0}^{\infty} p_j^* < \infty$ even when the number of commodities was infinite. Instead, this theorem made use of the convexity of preferences, consumption sets and production possibilities sets. So one might conjecture that in this model, which is clearly an exchange economy with convex preferences and convex consumption sets, Pareto optima must be decentralizable by some redistribution of endowments (even though some competitive equilibria may be Pareto suboptimal). This is in fact true, and the following proposition shows how the Pareto optimal allocation $\tilde{\mathbf{x}}$ described above can be decentralized as a competitive equilibrium:

PROPOSITION 9.3. *In the above-described economy, there exists a reallocation of the endowment vector $\boldsymbol{\omega}$ to $\tilde{\boldsymbol{\omega}}$, and an associated competitive equilibrium $(\bar{\mathbf{p}}, \tilde{\mathbf{x}})$ that is Pareto optimal where $\tilde{\mathbf{x}}$ is as described above, and $\bar{\mathbf{p}}$ is such that $\bar{p}_j = 1$ for all $j \in \mathbb{N}$.*

PROOF. Consider the following reallocation of the endowment vector $\boldsymbol{\omega}$: the endowment of household $i \geq 1$ is given to household $i-1$. Consequently, at the new endowment vector $\tilde{\boldsymbol{\omega}}$, household $i = 0$ has one unit of good $j = 0$ and one unit of good $j = 1$, while all other households i have one unit of good $i+1$. At the price vector $\bar{\mathbf{p}}$, household 0 has a budget set

$$c_0^0 + c_1^1 \leq 2,$$

thus chooses $c_0^0 = c_1^0 = 1$. All other households have budget sets given by

$$c_i^i + c_{i+1}^i \leq 1,$$

thus it is optimal for each household $i > 0$ to consume one unit of the good c_{i+1}^i, which is within its budget set and gives as high utility as any other allocation within his budget set, establishing that $\tilde{\mathbf{x}}$ is a competitive equilibrium. □

9.2. The Baseline Overlapping Generations Model

I now discuss the baseline two-period overlapping generation economy.

9.2.1. Demographics, Preferences and Technology. In this economy, time is discrete and runs to infinity. Each individual lives for two periods. For example, all individuals born at time t live for dates t and $t+1$. For now let us assume a general (separable) utility function for individuals born at date t, of the form

$$U(t) = u(c_1(t)) + \beta u(c_2(t+1)), \tag{9.1}$$

where $u : \mathbb{R}_+ \to \mathbb{R}$ satisfies the conditions in Assumption 3 (from Chapter 8) $c_1(t)$ denotes the consumption of the individual born at time t when young (at date t), and $c_2(t+1)$ is the individual's consumption when old (at date $t+1$). Also $\beta \in (0,1)$ is the discount factor. Factor markets are competitive. Individuals can only work in the first period of their lives and supply one unit of labor inelastically, earning the equilibrium wage rate $w(t)$.

Let us also assume that there is exponential population growth, so that total population is

$$L(t) = (1+n)^t L(0). \tag{9.2}$$

The production side of the economy is the same as before, characterized by a set of competitive firms, and is represented by a standard constant returns to scale aggregate production function, satisfying Assumptions 1 and 2 (from Chapter 2);

$$Y(t) = F(K(t), L(t)).$$

To simplify the analysis let us assume that $\delta = 1$, so that capital fully depreciates after use (see Exercise 9.3). Thus, again defining $k \equiv K/L$, the (gross) rate of return to saving, which equals the rental rate of capital, is given by

$$1 + r(t) = R(t) = f'(k(t)), \tag{9.3}$$

where $f(k) \equiv F(k,1)$ is the standard per capita production function. As usual, the wage rate is

$$w(t) = f(k(t)) - k(t) f'(k(t)). \tag{9.4}$$

9.2.2. Consumption Decisions. Let us start with the individual consumption decisions. Savings by an individual of generation t, $s(t)$, is determined as a solution to the following maximization problem

$$\max_{c_1(t), c_2(t+1), s(t)} u(c_1(t)) + \beta u(c_2(t+1))$$

subject to

$$c_1(t) + s(t) \leq w(t)$$

and

$$c_2(t+1) \leq R(t+1)\, s(t),$$

where I am using the convention that old individuals rent their savings of time t as capital to firms at time $t+1$, so they receive the gross rate of return $R(t+1) = 1 + r(t+1)$. The second constraint incorporates the notion that individuals will only spend money on their own end of life consumption (since there is no altruism or bequest motive). There is no need

to introduce the additional constraint that $s(t) \geq 0$, since negative savings would violate the second-period budget constraint (given that $c_2(t+1) \geq 0$).

Since the utility function $u(\cdot)$ is strictly increasing (Assumption 3), both constraints will hold as equalities. Therefore, the first-order condition for a maximum can be written in the familiar form of the consumption Euler equation (for the discrete-time problem, recall Chapter 6),

$$u'(c_1(t)) = \beta R(t+1) u'(c_2(t+1)). \tag{9.5}$$

Moreover, since the problem of each individual is strictly concave, this Euler equation is sufficient to characterize an optimal consumption path given market prices. Combining this equation with the budget constraint, we obtain the following implicit function that determines savings per person as

$$s(t) = s(w(t), R(t+1)), \tag{9.6}$$

where $s : \mathbb{R}^2_+ \to \mathbb{R}$ is strictly increasing in its first argument and may be increasing or decreasing in its second argument (see Exercise 9.4). Total savings in the economy will be equal to

$$S(t) = s(t) L(t),$$

where $L(t)$ denotes the size of generation t, who are saving for time $t+1$. Since capital depreciates fully after use and all new savings are invested in the only productive asset of the economy, capital, the law of motion of the capital stock is given by

$$K(t+1) = L(t) s(w(t), R(t+1)). \tag{9.7}$$

9.2.3. Equilibrium. A competitive equilibrium in the overlapping generations economy can be defined as follows:

DEFINITION 9.1. *A **competitive equilibrium** can be represented by sequences of aggregate capital stocks, household consumption and factor prices, $\{K(t), c_1(t), c_2(t), R(t), w(t)\}_{t=0}^{\infty}$, such that the factor price sequence $\{R(t), w(t)\}_{t=0}^{\infty}$ is given by (9.3) and (9.4), individual consumption decisions $\{c_1(t), c_2(t)\}_{t=0}^{\infty}$ are given by (9.5) and (9.6), and the aggregate capital stock, $\{K(t)\}_{t=0}^{\infty}$, evolves according to (9.7).*

A steady-state equilibrium is defined in the usual fashion, as an equilibrium in which the capital-labor ratio $k \equiv K/L$ is constant.

To characterize the equilibrium, divide (9.7) by labor supply at time $t+1$, $L(t+1) = (1+n) L(t)$, to obtain the capital-labor ratio as

$$k(t+1) = \frac{s(w(t), R(t+1))}{1+n}.$$

Now substituting for $R(t+1)$ and $w(t)$ from (9.3) and (9.4),

$$k(t+1) = \frac{s(f(k(t)) - k(t) f'(k(t)), f'(k(t+1)))}{1+n} \tag{9.8}$$

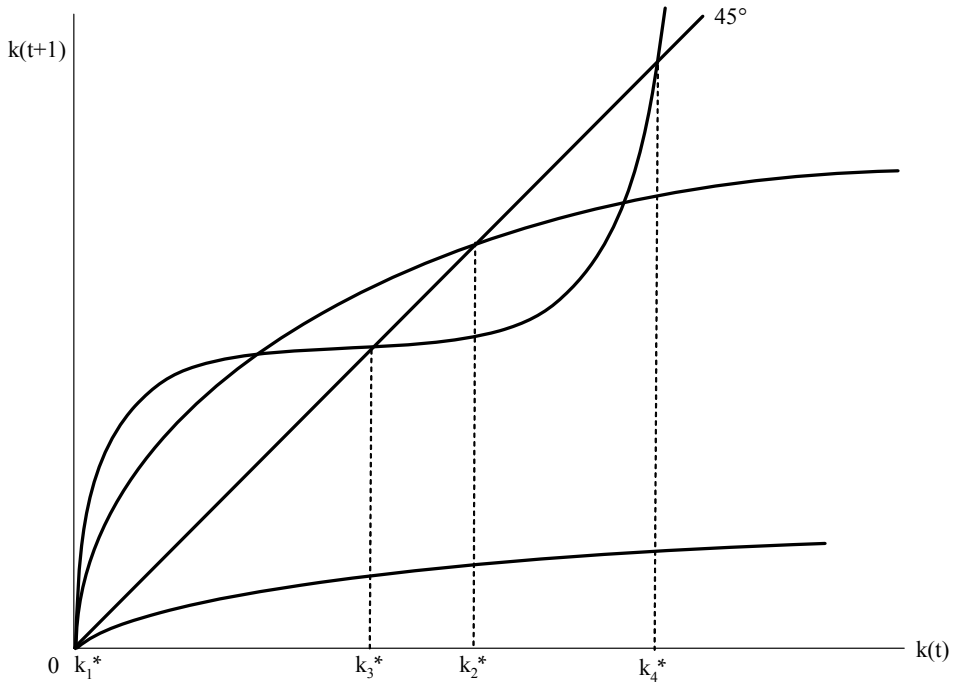

FIGURE 9.1. Various types of steady-state equilibria in the baseline overlapping generations model.

as the fundamental law of motion of the overlapping generations economy. A steady state is given by a solution to this equation such that $k(t+1) = k(t) = k^*$, that is,

$$k^* = \frac{s\left(f\left(k^*\right) - k^* f'\left(k^*\right), f'\left(k^*\right)\right)}{1+n} \tag{9.9}$$

Since the savings function $s(\cdot,\cdot)$ can take any form, the difference equation (9.8) can lead to quite complicated dynamics, and multiple steady states are possible. The next figure shows some potential forms that eq. (9.8) can take. The figure illustrates that the overlapping generations model can lead to a unique stable equilibrium, to multiple equilibria, or to an equilibrium with zero capital stock. In other words, without putting more structure on utility and/or production functions, the model makes few predictions.

9.2.4. Restrictions on Utility and Production Functions. In this subsection, I characterize the steady-state equilibrium and transition dynamics when further assumptions are imposed on the utility and production functions. In particular, let us suppose that the utility functions take the familiar CRRA form:

$$U(t) = \frac{c_1(t)^{1-\theta} - 1}{1-\theta} + \beta\left(\frac{c_2(t+1)^{1-\theta} - 1}{1-\theta}\right), \tag{9.10}$$

where $\theta > 0$ and $\beta \in (0,1)$. Furthermore, assume that technology is Cobb-Douglas, so that

$$f(k) = k^{\alpha}$$

The rest of the environment is as described above. The CRRA utility simplifies the first-order condition for consumer optimization and implies

$$\frac{c_2(t+1)}{c_1(t)} = (\beta R(t+1))^{1/\theta}.$$

Once again, this expression is the discrete-time consumption Euler equation from Chapter 6, now for the CRRA utility function. This Euler equation can be alternatively expressed in terms of savings as

$$s(t)^{-\theta} \beta R(t+1)^{1-\theta} = (w(t) - s(t))^{-\theta}, \tag{9.11}$$

which gives the following equation for the saving rate:

$$s(t) = \frac{w(t)}{\psi(t+1)}, \tag{9.12}$$

where

$$\psi(t+1) \equiv [1 + \beta^{-1/\theta} R(t+1)^{-(1-\theta)/\theta}] > 1,$$

which ensures that savings are always less than earnings. The impact of factor prices on savings is summarized by the following derivatives:

$$\begin{aligned} s_w &\equiv \frac{\partial s(t)}{\partial w(t)} = \frac{1}{\psi(t+1)} \in (0,1), \\ s_r &\equiv \frac{\partial s(t)}{\partial R(t+1)} = \left(\frac{1-\theta}{\theta}\right) (\beta R(t+1))^{-1/\theta} \frac{s(t)}{\psi(t+1)}. \end{aligned}$$

Since $\psi(t+1) > 1$, we also have that $0 < s_w < 1$. Moreover, in this case $s_r > 0$ if $\theta < 1$, $s_r < 0$ if $\theta > 1$, and $s_r = 0$ if $\theta = 1$. The relationship between the rate of return on savings and the level of savings reflects the counteracting influences of income and substitution effects. The case of $\theta = 1$ (log preferences) is of special importance and is often used in many applied models. This special case is sufficiently common that it may deserve to be called the *canonical overlapping generations model* and will be analyzed separately in the next section.

In the current somewhat more general context, eq. (9.8) implies

$$\begin{aligned} k(t+1) &= \frac{s(t)}{(1+n)} \\ &= \frac{w(t)}{(1+n)\psi(t+1)}, \end{aligned} \tag{9.13}$$

or more explicitly,

$$k(t+1) = \frac{f(k(t)) - k(t) f'(k(t))}{(1+n)\left[1 + \beta^{-1/\theta} f'(k(t+1))^{-(1-\theta)/\theta}\right]}. \tag{9.14}$$

The steady state then involves a solution to the following implicit equation:

$$k^* = \frac{f(k^*) - k^* f'(k^*)}{(1+n)\left[1 + \beta^{-1/\theta} f'(k^*)^{-(1-\theta)/\theta}\right]}.$$

Now using the Cobb-Douglas formula, the steady state is obtained as the solution to the equation

$$(1+n)\left[1+\beta^{-1/\theta}\left(\alpha(k^*)^{\alpha-1}\right)^{(\theta-1)/\theta}\right]=(1-\alpha)(k^*)^{\alpha-1}. \tag{9.15}$$

For simplicity, define $R^* \equiv \alpha(k^*)^{\alpha-1}$ as the marginal product of capital in steady-state, in which case, eq. (9.15) can be rewritten as

$$(1+n)\left[1+\beta^{-1/\theta}\left(R^*\right)^{(\theta-1)/\theta}\right]=(1-\alpha)\,R^*/\alpha. \tag{9.16}$$

The steady-state value of R^*, and thus of k^*, can now be determined from eq. (9.16), which always has a unique solution. Let us next investigate the stability of this steady state. Substitute for the Cobb-Douglas production function in (9.14) to obtain

$$k\left(t+1\right)=\frac{\left(1-\alpha\right)k\left(t\right)^{\alpha}}{\left(1+n\right)\left[1+\beta^{-1/\theta}\left(\alpha k(t+1)^{\alpha-1}\right)^{-(1-\theta)/\theta}\right]}. \tag{9.17}$$

Using (9.17), the following proposition can be proved:[1]

Proposition 9.4. *In the overlapping-generations model with two-period lived households, Cobb-Douglas technology and CRRA preferences, there exists a unique steady-state equilibrium with the capital-labor ratio k^* given by (9.15) and as long as $\theta \geq 1$, this steady-state equilibrium is globally stable for all $k\left(0\right)>0$.*

Proof. See Exercise 9.5. □

In this particular (well-behaved) case, equilibrium dynamics are very similar to the basic Solow model, and are shown in Figure 9.2 (more specifically, the equilibrium dynamics in this figure are drawn for the canonical overlapping generations model, which will be presented in the next section). The figure shows that convergence to the unique steady-state capital-labor ratio, k^*, is monotonic. In particular, starting with an initial capital-labor ratio $k\left(0\right)<k^*$, the overlapping generations economy steadily accumulates more capital and converge to k^*. Starting with $k'\left(0\right)>k^*$, the equilibrium involves lower and lower levels of capital-labor ratio, ultimately converging to k^*.

9.3. The Canonical Overlapping Generations Model

Even the model with CRRA utility and Cobb-Douglas production function is relatively messy. For this reason, many of the applications of the overlapping generations model use an even more specific utility function, log preferences (or equivalently $\theta=1$ in terms of the CRRA preferences of the last section). Log preferences are particularly useful in this context, since they ensure that income and substitution effects exactly cancel out, so that changes in the interest rate (and therefore changes in the capital-labor ratio of the economy) have no effect on the saving rate.

[1] In this proposition and throughout the rest of this chapter, I again ignore the trivial steady state with $k=0$.

Since this version of the model is sufficiently common, it may deserve to be called the canonical overlapping generations model and will be the focus of this section. Another interesting feature of this model is that the structure of the equilibrium is essentially identical to the basic Solow model in Chapter 2.

Suppose that the utility of the household and generation t is given by

$$U(t) = \log c_1(t) + \beta \log c_2(t+1), \tag{9.18}$$

where as before $\beta \in (0,1)$ (even though $\beta \geq 1$ could be allowed here without any change in the analysis). The aggregate production technology is again Cobb-Douglas, that is, $f(k) = k^{\alpha}$. The consumption Euler equation now becomes even simpler:

$$\frac{c_2(t+1)}{c_1(t)} = \beta R(t+1)$$

and implies that savings should satisfy the equation

$$s(t) = \frac{\beta}{1+\beta} w(t), \tag{9.19}$$

which corresponds to a constant saving rate, equal to $\beta/(1+\beta)$, out of labor income for each individual. This constant saving rate makes this model very similar to the baseline Solow growth model of Chapter 2.

Now combining (9.19) with the capital accumulation eq. (9.8),

$$\begin{aligned} k(t+1) &= \frac{s(t)}{(1+n)} \\ &= \frac{\beta w(t)}{(1+n)(1+\beta)} \\ &= \frac{\beta(1-\alpha)[k(t)]^{\alpha}}{(1+n)(1+\beta)}, \end{aligned}$$

where the second line uses (9.19) and the last line uses the fact that, given competitive factor markets, the wage rate is equal to $w(t) = (1-\alpha)[k(t)]^{\alpha}$.

It is straightforward to verify that there exists a unique steady state with capital-labor ratio given by

$$k^* = \left[\frac{\beta(1-\alpha)}{(1+n)(1+\beta)}\right]^{\frac{1}{1-\alpha}}. \tag{9.20}$$

Moreover, starting with any $k(0) > 0$, equilibrium dynamics are identical to those of the basic Solow model and monotonically converge to k^*. This is illustrated in Figure 9.2 and stated in the next proposition:

Proposition 9.5. *In the canonical overlapping generations model with log preferences and Cobb-Douglas technology, there exists a unique steady state, with capital-labor ratio k^* given by (9.20). Starting with any $k(0) \in (0, k^*)$, equilibrium dynamics are such that $k(t) \uparrow k^*$, and starting with any $k'(0) > k^*$, equilibrium dynamics involve $k(t) \downarrow k^*$.*

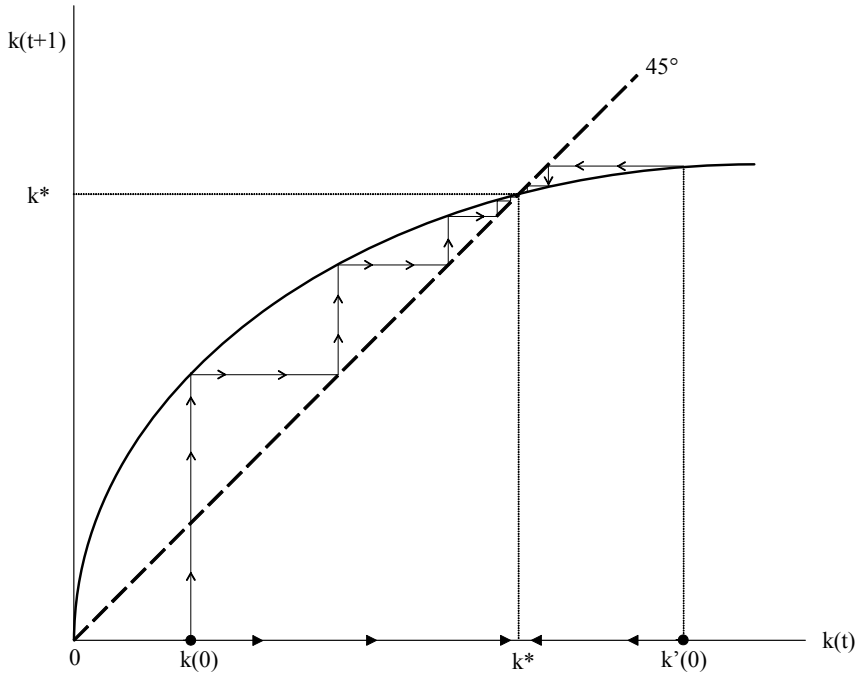

FIGURE 9.2. Equilibrium dynamics in the canonical overlapping generations model.

Exercise 9.6 asks you to introduce technological progress into this canonical model and to conduct a range of comparative static exercises. Exercise 9.7 asks you to analyze the same economy without the Cobb-Douglas technology assumption.

9.4. Overaccumulation and Pareto Optimality of Competitive Equilibrium in the Overlapping Generations Model

Let us now return to the general problem and compare the overlapping-generations equilibrium to the choice of a social planner wishing to maximize a weighted average of all generations' utilities. In particular, suppose that the social planner maximizes

$$\sum_{t=0}^{\infty} \beta_S^t U(t),$$

where β_S is the discount factor of the social planner, which reflects how she values the utilities of different generations. Substituting from (9.1), this implies that the planner will maximize

$$\sum_{t=0}^{\infty} \beta_S^t \left(u\left(c_1(t)\right) + \beta u\left(c_2(t+1)\right)\right)$$

subject to the resource constraint

$$F\left(K(t), L(t)\right) = K(t+1) + L(t) c_1(t) + L(t-1) c_2(t).$$

Dividing this by $L(t)$ and using (9.2), the resource constraint can be written in per capita terms as

$$f(k(t)) = (1+n)k(t+1) + c_1(t) + \frac{c_2(t)}{1+n}.$$

The social planner's maximization problem then implies the following first-order necessary condition:

$$u'(c_1(t)) = \beta f'(k(t+1))u'(c_2(t+1)).$$

Since $R(t+1) = f'(k(t+1))$, this is identical to (9.5). This result is not surprising; the social planner prefers to allocate consumption of a given individual in exactly the same way as the individual himself would do; there are no "market failures" in the over-time allocation of consumption at given prices.

However, the social planner's and the competitive economy's allocations across generations will differ, since the social planner is giving different weights to different generations as captured by the parameter β_S. In particular, it can be shown that the socially planned economy will converge to a steady state with capital-labor ratio k^S such that

$$\beta_S f'\left(k^S\right) = 1+n,$$

which is similar to the modified golden rule in the context of the Ramsey growth model in discrete time (cf., Chapter 6). In particular, the steady-state level of capital-labor ratio k^S chosen by the social planner does not depend on preferences (on the utility function $u(\cdot)$) and does not even depend on the individual rate of time preference, β. Clearly, k^S will typically differ from the steady-state value of the competitive economy, k^*, given by (9.9).

More interesting is the question of whether the competitive equilibrium is Pareto optimal. The example in Section 9.1 suggests that it may not be. Exactly as in that example, we cannot use the First Welfare Theorem, Theorem 5.6, because there is an infinite number of commodities and the sum of their prices is not necessarily less than infinity.

In fact, the competitive equilibrium is not in general Pareto optimal. The simplest way of seeing this is that the steady state level of capital stock, k^*, given by (9.9), can be so high that it is in fact greater than k_{gold}. Recall that k_{gold} is the golden rule level of capital-labor ratio that maximizes the steady-state level of consumption (recall, for example, Figure 8.1 in Chapter 8 for the discussion in Chapter 2). When $k^* > k_{gold}$, reducing savings can increase consumption for every generation.

More specifically, note that in steady state

$$\begin{aligned} f(k^*) - (1+n)k^* &= c_1^* + (1+n)^{-1} c_2^* \\ &\equiv c^*, \end{aligned}$$

where the first line follows by national income accounting, and the second defines c^* as the total steady-state consumption. Therefore

$$\frac{\partial c^*}{\partial k^*} = f'(k^*) - (1+n),$$

and k_{gold} is defined as

$$f'(k_{gold}) = 1 + n.$$

Now if $k^* > k_{gold}$, then $\partial c^*/\partial k^* < 0$, so reducing savings can increase (total) consumption for everybody. If this is the case, the economy is referred to as *dynamically inefficient*—it involves overaccumulation. Another way of expressing dynamic inefficiency is that

$$r^* < n,$$

that is, the steady-state (net) interest rate $r^* = R^* - 1$ is less than the rate of population growth. Recall that in the infinite-horizon Ramsey economy, the transversality condition (which follows from individual optimization) required that $r > g + n$, therefore, dynamic inefficiency could never arise in this Ramsey economy. Dynamic inefficiency arises because of the heterogeneity inherent in the overlapping generations model, which removes the transversality condition.

In particular, suppose we start from steady state at time T with $k^* > k_{gold}$. Consider the following variation where the capital stock for next period is reduced by a small amount. In particular, change next period's capital stock by $-\Delta k$, where $\Delta k > 0$, and from then on, imagine that we immediately move to a new steady state (which is clearly feasible). This implies the following changes in consumption levels:

$$\begin{aligned} \Delta c(T) &= (1+n)\Delta k > 0 \\ \Delta c(t) &= -\left(f'(k^* - \Delta k) - (1+n)\right)\Delta k \text{ for all } t > T. \end{aligned}$$

The first expression reflects the direct increase in consumption due to the decrease in savings. In addition, since $k^* > k_{gold}$, for small enough Δk, $f'(k^* - \Delta k) - (1+n) < 0$, thus $\Delta c(t) > 0$ for all $t \geq T$, and explains the second expression. The increase in consumption for each generation can be allocated equally during the two periods of their lives, thus necessarily increasing the utility of all generations. This variation clearly creates a Pareto improvement in which all generations are better-off. This establishes:

Proposition 9.6. *In the baseline overlapping-generations economy, the competitive equilibrium is not necessarily Pareto optimal. More specifically, whenever $r^* < n$, the economy is dynamically inefficient. In this case, it is possible to reduce the capital stock starting from the competitive steady state and increase the consumption level of all generations.*

As the above derivation makes it clear, Pareto inefficiency of the competitive equilibrium is intimately linked with *dynamic inefficiency.* Dynamic inefficiency, the rate of interest being less than the rate of population growth, is not a theoretical curiosity. Exercise 9.8 shows that dynamic inefficiency can arise under reasonable circumstances.

Loosely speaking, the intuition for dynamic inefficiency can be given as follows. Individuals who live at time t face prices determined by the capital stock with which they are working. This capital stock is the outcome of actions taken by previous generations. Therefore, there is a *pecuniary externality* from the actions of previous generations affecting the welfare of the current generation. Pecuniary externalities are typically second-order and do not matter for

welfare (in a sense this could be viewed as the essence of the First Welfare Theorem). This ceases to be the case, however, when there are an infinite stream of newborn agents joining the economy. These agents are affected by the pecuniary externalities created by previous generations, and it is possible to rearrange accumulation decisions and consumption plans in such a way that these pecuniary externalities can be exploited.

A complementary intuition for dynamic inefficiency, which will be particularly useful in the next section, is as follows. Dynamic inefficiency arises from overaccumulation, which, in turn, is a result of the fact that the current young generation needs to save for old age. However, the more they save, the lower is the rate of return to capital and this may encourage them to save even more. Once again, the effect of the savings by the current generation on the future rate of return to capital is a pecuniary externality on the next generation. We may reason that this pecuniary externality should not lead to Pareto suboptimal allocations, as in the equilibria of standard competitive economies with a finite number of commodities and households. But this reasoning is no longer correct when there are an infinite number of commodities and an infinite number of households. This second intuition also suggests that if, somehow, alternative ways of providing consumption to individuals in old age were introduced, the overaccumulation problem could be solved or at least ameliorated. This is the topic of the next section.

9.5. Role of Social Security in Capital Accumulation

I now briefly discuss how Social Security can be introduced as a way of dealing with overaccumulation in the overlapping-generations model. Let us first consider a fully-funded system, in which the young make contributions to the Social Security system and their contributions are paid back to them in their old age. The alternative is an unfunded system or a *pay-as-you-go* Social Security system, where transfers from the young directly go to the current old. We will see that, as is typically presumed, pay-as-you-go (unfunded) Social Security discourages aggregate savings. However, when there is dynamic inefficiency, discouraging savings may lead to a Pareto improvement.

9.5.1. Fully Funded Social Security. In a fully funded Social Security system, the government at date t raises some amount $d(t)$ from the young, for example, by compulsory contributions to their Social Security accounts. These funds are invested in the only productive asset of the economy, the capital stock, and the workers receive the returns, given by $R(t+1)\,d(t)$, when they are old. This implies that the individual maximization problem under a fully funded social security system becomes

$$\max_{c_1(t),c_2(t+1),s(t)} u(c_1(t)) + \beta u(c_2(t+1))$$

subject to

$$c_1(t) + s(t) + d(t) \leq w(t)$$

and

$$c_2(t+1) \leq R(t+1)(s(t) + d(t)),$$

for a given choice of $d(t)$ by the government. Notice that now the total amount invested in capital accumulation is $s(t) + d(t) = (1+n)k(t+1)$.

It is also no longer the case that individuals will always choose $s(t) > 0$, since they have the income from Social Security. Therefore this economy can be analyzed under two alternative assumptions, with the constraint that $s(t) \geq 0$ and without.

It is clear that as long as $s(t)$ is free, whatever the sequence of feasible Social Security payments $\{d(t)\}_{t=0}^{\infty}$, the competitive equilibrium applies. When $s(t) \geq 0$ is imposed as a constraint, then the competitive equilibrium applies if given the sequence $\{d(t)\}_{t=0}^{\infty}$, the privately-optimal saving sequence $\{s(t)\}_{t=0}^{\infty}$ is such that $s(t) > 0$ for all t. Consequently:

PROPOSITION 9.7. *Consider a fully funded Social Security system in the above-described environment whereby the government collects $d(t)$ from young individuals at date t.*

(1) *Suppose that $s(t) \geq 0$ for all t. If given the feasible sequence $\{d(t)\}_{t=0}^{\infty}$ of Social Security payments, the utility-maximizing sequence of savings $\{s(t)\}_{t=0}^{\infty}$ is such that $s(t) > 0$ for all t, then the set of competitive equilibria without Social Security are the set of competitive equilibria with Social Security.*
(2) *Without the constraint $s(t) \geq 0$, given any feasible sequence $\{d(t)\}_{t=0}^{\infty}$ of Social Security payments, the set of competitive equilibria without Social Security are the set of competitive equilibria with Social Security.*

PROOF. See Exercise 9.10. □

This is very intuitive: the $d(t)$ taken out by the government is fully offset by a decrease in $s(t)$ as long as individuals were performing enough savings (or always when there are no constraints to force positive savings privately). Exercise 9.11 shows that even when there is the restriction that $s(t) \geq 0$, a funded Social Security program cannot lead to a Pareto improvement.

9.5.2. Unfunded Social Security. The situation is different with unfunded Social Security. Now the government collects $d(t)$ from the young at time t and distributes this to the current old with per capita transfer $b(t) = (1+n)d(t)$ (which takes into account that there are more young than old because of population growth). Therefore, the individual maximization problem becomes

$$\max_{c_1(t),c_2(t+1),s(t)} u(c_1(t)) + \beta u(c_2(t+1))$$

subject to

$$c_1(t) + s(t) + d(t) \leq w(t)$$

and

$$c_2(t+1) \leq R(t+1)s(t) + (1+n)d(t+1),$$

for a given feasible sequence of Social Security payment levels $\{d(t)\}_{t=0}^{\infty}$.

What this implies is that the rate of return on Social Security payments is n rather than $r(t+1) = R(t+1) - 1$, because unfunded Social Security is a pure transfer system. Only

$s(t)$—rather than $s(t)$ plus $d(t)$ as in the funded scheme—goes into capital accumulation. This is the basis of the claim that unfunded Social Security systems discourage aggregate savings. Of course, it is possible that $s(t)$ will change in order to compensate this effect, but such an offsetting change does not typically take place. Consequently, unfunded Social Security reduces capital accumulation. Discouraging capital accumulation can have negative consequences for growth and welfare. In fact, the empirical evidence in Chapters 1-4 suggests that there are many societies in which the level of capital accumulation is suboptimally low. In contrast, in the current model reducing aggregate savings and capital accumulation may be a good thing when the economy exhibits dynamic inefficiency (and overaccumaltion). This leads to the following proposition.

Proposition 9.8. *Consider the above-described overlapping generations economy and suppose that the decentralized competitive equilibrium is dynamically inefficient. Then, there exists a feasible sequence of unfunded Social Security payments* $\{d(t)\}_{t=0}^{\infty}$ *which will lead to a competitive equilibrium starting from any date* t *that Pareto dominates the competitive equilibrium without Social Security.*

Proof. See Exercise 9.13. □

Unfunded Social Security reduces the overaccumulation and improves the allocation of resources. The similarity between the way in which unfunded Social Security achieves a Pareto improvement in this proposition and the way in which the Pareto optimal allocation was decentralized in the example economy of Section 9.1 is apparent. In essence, unfunded Social Security is transferring resources from future generations to initial old generation, and when designed appropriately, it can do so without hurting future generations. Once again, this depends on dynamic inefficiency; when there is *no* dynamic inefficiency, any transfer of resources (and any unfunded Social Security program) would make some future generation worse-off. You are asked to prove this result in Exercise 9.14.

9.6. Overlapping Generations with Impure Altruism

Section 5.3 in Chapter 5 demonstrated that altruism within families (for example of parents towards their offspring) can lead to a structure of preferences identical to those of the representative household in the neoclassical growth model. In contrast, this section has so far ignored altruistic preferences in order to emphasize the effect of finite lives and the economic implications of the arrival of new agents in the economy. As briefly noted in Section 5.3, the exact form of altruism within a family matters for whether the representative household would provide a good approximation to the preference side of the economy. In particular, a potentially empirically relevant form of altruism is one in which parents care about certain dimensions of the consumption vector of their offspring instead of their total utility. These types of preferences are often referred to as "impure altruism" to distinguish it from the pure altruism discussed in Section 5.3. A particular type of impure altruism, commonly referred to as "warm glow preferences", plays an important role in many growth models because of its

tractability. Warm glow preferences assume that parents derive utility from (the warm glow of) their bequest, rather than the utility or the consumption of their offspring. This class of preferences turn out to constitute another very tractable alternative to the neoclassical growth and the baseline overlapping generations models. It has some clear parallels to the canonical overlapping generations model of last section, since it will also lead to equilibrium dynamics very similar to those of the Solow growth model. Given the importance of this class of preferences in many applied growth models, it is useful to review them briefly. These preferences will also be used in the next chapter and again in Chapter 21.

Suppose that the production side of the economy is given by the standard neoclassical production function, satisfying Assumptions 1 and 2. Let us write this in per capita form as $f(k)$.

The economy is populated by a continuum of individuals of measure 1. Each individual lives for two periods, childhood and adulthood. In the second period of his life, each individual begets an offspring, works and then his life comes to an end. For simplicity, let us assume that there is no consumption in childhood (or that this is incorporated in the parent's consumption). There are no new households, so population is constant at 1. Each individual supplies 1 unit of labor inelastically during his adulthood.

Let us assume that preferences of individual (i,t), who reaches adulthood at time t, are

$$\log\left(c_i\left(t\right)\right)+\beta\log\left(b_i\left(t\right)\right), \tag{9.21}$$

where $c_i(t)$ denotes the consumption of this individual and $b_i(t)$ is bequest to his offspring. Log preferences are assumed to simplify the analysis (see Exercise 9.19). The offspring starts the following period with the bequest, rents this out as capital to firms, supplies labor, begets his own offspring, and makes consumption and bequests decisions. I also assume that capital fully depreciates after use.

This formulation implies that the maximization problem of a typical individual can be written as

$$\max_{c_i(t),b_i(t)}\log\left(c_i\left(t\right)\right)+\beta\log\left(b_i\left(t\right)\right), \tag{9.22}$$

subject to

$$c_i\left(t\right)+b_i\left(t\right)\leq y_i\left(t\right)\equiv w\left(t\right)+R\left(t\right)b_i\left(t-1\right), \tag{9.23}$$

where $y_i(t)$ denotes the income of this individual,

$$w\left(t\right)=f\left(k\left(t\right)\right)-k\left(t\right)f'\left(k\left(t\right)\right) \tag{9.24}$$

is the equilibrium wage rate,

$$R\left(t\right)=f'\left(k\left(t\right)\right) \tag{9.25}$$

is the rate of return on capital and $b_i(t-1)$ is the bequest received by this individual from his own parent.

The total capital-labor ratio at time $t+1$ is given by aggregating the bequests of all adults at time t:

$$k(t+1) = \int_0^1 b_i(t)\, di, \tag{9.26}$$

which exploits the fact that the total measure of workers is 1, so that the capital stock and capital-labor ratio are identical.

An equilibrium in this economy is a somewhat more complicated object than before, because we may want to keep track of the consumption and bequest levels of all individuals. Let us denote the distribution of consumption and bequests across households at time t by $[c_i(t)]_{i\in[0,1]}$ and $[b_i(t)]_{i\in[0,1]}$, and let us assume that the economy starts with the distribution of wealth (bequests) at time t given by $[b_i(0)]_{i\in[0,1]}$, which satisfies $\int_0^1 b_i(0)\, di > 0$.

Definition 9.2. *A **competitive equilibrium** in this overlapping generations economy with warm glow preferences is a sequence of consumption and bequest levels for each household, $\left\{[c_i(t)]_{i\in[0,1]}, [b_i(t)]_{i\in[0,1]}\right\}_{t=0}^{\infty}$, that solve (9.22) subject to (9.23), a sequence of capital-labor ratios, $\{k(t)\}_{t=0}^{\infty}$, given by (9.26) with some initial distribution of bequests $[b_i(0)]_{i\in[0,1]}$, and sequences of factor prices, $\{w(t), R(t)\}_{t=0}^{\infty}$, that satisfy (9.24) and (9.25).*

The solution of (9.22) subject to (9.23) is straightforward because of the log preferences, and gives

$$\begin{aligned} b_i(t) &= \frac{\beta}{1+\beta} y_i(t) \\ &= \frac{\beta}{1+\beta}\left[w(t) + R(t)\, b_i(t-1)\right], \end{aligned} \tag{9.27}$$

for all i and t. This equation shows that individual bequest levels will follow non-trivial dynamics. Since $b_i(t)$ determines the asset holdings of individual i of generation t, it can alternatively be interpreted as his "wealth" level. Consequently, this economy will feature a distribution of wealth that will evolve endogenously over time. This evolution will depend on factor prices. To obtain factor prices, let us aggregate bequests to obtain the capital-labor ratio of the economy via eq. (9.26). Integrating (9.27) across all individuals,

$$\begin{aligned} k(t+1) &= \int_0^1 b_i(t)\, di \\ &= \frac{\beta}{1+\beta}\int_0^1 \left[w(t) + R(t)\, b_i(t-1)\right] di \\ &= \frac{\beta}{1+\beta} f(k(t)). \end{aligned} \tag{9.28}$$

The last equality follows from the fact that $\int_0^1 b_i(t-1)\, di = k(t)$ and because by Euler's Theorem, Theorem 2.1, $w(t) + R(t)\, k(t) = f(k(t))$.

Consequently, aggregate equilibrium dynamics in this economy are straightforward and again closely resemble those in the baseline Solow growth model. Moreover, it is worth noting that these aggregate dynamics do *not* depend on the distribution of bequests or income across

households (we will see that this is no longer true when there are other imperfections in the economy as in Chapter 21).

Now, solving for the steady-state equilibrium capital-labor ratio from (9.28),

$$k^* = \frac{\beta}{1+\beta} f(k^*), \tag{9.29}$$

which is uniquely defined and strictly positive in view of Assumptions 1 and 2. Moreover, equilibrium dynamics are again given by Figure 9.2 and involve monotonic convergence to this unique steady state.

A complete characterization of the equilibrium can now be obtained by looking at the dynamics of bequests. It turns out that different types of bequests dynamics are possible along the transition path. More can be said regarding the limiting distribution of wealth and bequests. In particular, we know that $k(t) \to k^*$, so the ultimate bequest dynamics are given by steady-state factor prices. Let these be denoted by $w^* = f(k^*) - k^* f'(k^*)$ and $R^* = f'(k^*)$. Then, once the economy is in the neighborhood of the steady-state capital-labor ratio, k^*, individual bequest dynamics are given by

$$b_i(t) = \frac{\beta}{1+\beta} \left[w^* + R^* b_i(t-1) \right].$$

When $R^* < (1+\beta)/\beta$, starting from any level $b_i(t)$ will converge to a unique bequest (wealth) level given by

$$b^* = \frac{\beta w^*}{1+\beta(1-R^*)}. \tag{9.30}$$

Moreover, it can be verified that the steady-state equilibrium must involve $R^* < (1+\beta)/\beta$. This follows from the fact that in steady state

$$\begin{aligned} R^* &= f'(k^*) \\ &< \frac{f(k^*)}{k^*} \\ &= \frac{1+\beta}{\beta}, \end{aligned}$$

where the second line exploits the strict concavity of $f(\cdot)$ and the last line uses the definition of the steady-state capital-labor ratio, k^*, from (9.29). The following proposition summarizes this analysis:

PROPOSITION 9.9. *Consider the overlapping generations economy with warm glow preferences described above. In this economy, there exists a unique competitive equilibrium. In this equilibrium the aggregate capital-labor ratio is given by (9.28) and monotonically converges to the unique steady-state capital-labor ratio k^* given by (9.29). The distribution of bequests and wealth ultimately converges towards full equality, with each individual having a bequest (wealth) level of b^* given by (9.30) with $w^* = f(k^*) - k^* f'(k^*)$ and $R^* = f'(k^*)$.*

9.7. Overlapping Generations with Perpetual Youth

A key feature of the baseline overlapping generation model is that individuals have finite lives and know exactly when their lives will come to an end. An alternative way of modeling finite lives is along the lines of the "Poisson death model" or the *perpetual youth model* introduced in Section 5.3 of Chapter 5. Let us start with the discrete-time version of that model. Recall that in that model each individual is potentially infinitely lived, but faces a probability $\nu \in (0,1)$ that his life will come to an end at every date (and these probabilities are independent). Recall from eq. (5.12) that the expected utility of an individual with a "pure" discount factor β is given by

$$\sum_{t=0}^{\infty} \left(\beta \left(1-\nu\right)\right)^t u\left(c\left(t\right)\right),$$

where $u\left(\cdot\right)$ is a standard instantaneous utility function, satisfying Assumption 3, with the additional normalization that $u\left(0\right)=0$. Since the probability of death is ν and is independent across periods, the expected lifetime of an individual in this model can be written as (see Exercise 9.15):

$$\text{Expected life} = \nu + 2\left(1-\nu\right)\nu + 3\left(1-\nu\right)^2\nu + ... = \frac{1}{\nu} < \infty. \tag{9.31}$$

This equation captures the fact that with probability ν the individual will have a total life of length 1, with probability $\left(1-\nu\right)\nu$, he will have a life of length 2, and so on. This model is referred to as the perpetual youth model, since even though each individual has a finite expected life, all individuals who have survived up to a certain date have exactly the same expectation of further life. Therefore, individuals who survive in this economy are "perpetually young"; their age has no effect on their future longevity and has no predictive power on how many more years they will live for.

Individual i's flow budget constraint can be written as

$$a_i\left(t+1\right) = \left(1+r\left(t+1\right)\right)a_i\left(t\right) - c_i\left(t\right) + w\left(t\right) + z_i\left(t\right), \tag{9.32}$$

which is similar to the standard flow budget constraint, for example (6.51) in Chapter 6. Recall that the gross rate of return on savings is $1+r\left(t+1\right)$, with the timing convention reflecting that assets at time t are rented out as capital at time $t+1$. The only difference from the standard budget constraint is the additional term, $z_i\left(t\right)$, which reflects transfers to the individual. The reason why these transfers are introduced is as follows: since individuals face an uncertain time of death, there may be "accidental bequests". In particular, individuals will typically come to the end of their lives while their asset positions are positive. When this happens, one possibility is that the accidental bequests might be collected by the government and redistributed equally across all households in the economy. In this case, $z_i\left(t\right)$ would represent these receipts for individual i. However, this would require that we impose a constraint of the form $a_i\left(t\right) \geq 0$, in order to prevent individuals from accumulating debts by the time their life comes to an end.

An alternative, which avoids this additional constraint and makes the model more tractable, has been proposed and studied by Menahem Yaari and Olivier Blanchard. This alternative involves introducing life-insurance or annuity markets, where competitive life insurance firms make payments to individuals (as a function of their asset levels) in return for receiving their positive assets when they die. The term $z(t)$ captures these annuity payments. In particular, imagine the following type of life insurance contract: a company would make a payment equal to $z(a(t))$ to an individual as a function of his asset holdings during every period in which he is alive.[2] When the individual dies, all his assets go to the insurance company. The fact that the payment level $z(a(t))$ depends only on the asset holdings of the individual and not on his age is a consequence of the perpetual youth assumption—the conditional expectation of further life is independent of when the individual was born and in fact, it is independent of everything else in the model. The profits of a particular insurance company contracting with an individual with asset holding equal to $a(t)$, at time t will be

$$\pi(a,t) = -(1-\nu) z(a) + \nu a.$$

With free entry, insurance companies should make zero expected profits (in terms of net present discounted value), which requires that $\pi(a(t),t) = 0$ for all t and a, thus

$$z(a(t)) = \frac{\nu}{1-\nu} a(t). \tag{9.33}$$

The other important element of the model is the evolution of the demographics. Since each agent faces a probability of death equal to ν at every date, there is a natural force towards decreasing population. We assume, however, that there are also new agents who are born at every date. Differently from the basic neoclassical growth model, I assume here that these new agents are not born into a dynasty; instead, they become separate households themselves. We assume that when the population at time t is $L(t)$, there are $nL(t)$ new households born. Consequently, the evolution of total population is given by

$$L(t+1) = (1+n-\nu) L(t), \tag{9.34}$$

with the initial value $L(0) = 1$, where

$$n > \nu,$$

so that there is positive population growth. Throughout this section, we ignore technological progress.

Perpetual youth, together with exponential population growth, leads to a simple pattern of demographics in this economy. In particular, it is easy to verify that at some point in time $t > 0$, there will be $n(1+n-\nu)^{t-1}$ one-year-olds, $n(1+n-\nu)^{t-2}(1-\nu)$ two-year-olds, $n(1+n-\nu)^{t-3}(1-\nu)^2$ three-year-olds, and so on (see Exercise 9.21).

The production side of the economy is standard and it is represented by an aggregate production function satisfying Assumptions 1 and 2, $F(K(t), L(t))$. Suppose that capital

[2]The reader might note that this is the opposite of the most common type of life insurance contract where individuals make payments in order for their families to receive payments after their death. These types of insurance contracts are not useful in the current model, since individuals do not have offspring or are not altruistic towards them.

depreciates at the rate δ. Factor markets are competitive and factor prices are determined in the usual fashion. The rental return of capital at time t is again given by $R(t) = f'(k(t))$, so that the net return on saving is $r(t+1) = f'(k(t)) - \delta$, and the wage rate is $w(t) = f(k(t)) - k(t) f'(k(t))$.

An allocation in this economy is similar to an allocation in the neoclassical growth model and involves time paths for the aggregate capital stock, wage rates and rental rates of capital, $\{K(t), w(t), R(t)\}_{t=0}^{\infty}$. However, it is no longer sufficient to specify aggregate consumption, since the level of consumption is not the same for all individuals. Instead, individuals born at different times will have accumulated different amounts of assets and will consume different amounts. Let us denote the consumption at date t of a household born at date $\tau \leq t$ by $c(t \mid \tau)$. An allocation must now specify the entire sequence $\{c(t \mid \tau)\}_{t=0,\tau\leq t}^{\infty}$. Using this notation and the life insurance contracts introduced by (9.33), the flow budget constraint of an individual of generation τ can be written as:

$$a(t+1 \mid \tau) = \left(1 + r(t+1) + \frac{\nu}{1-\nu}\right) a(t \mid \tau) - c(t \mid \tau) + w(t). \tag{9.35}$$

A competitive equilibrium in this economy can then be defined as follows:

DEFINITION 9.3. *A **competitive equilibrium** consists of paths of capital stock, wage rates and rental rates of capital, $\{K(t), w(t), R(t)\}_{t=0}^{\infty}$, and paths of consumption for each generation, $\{c(t \mid \tau)\}_{t=0,\tau\leq t}^{\infty}$, such that each individual maximizes utility and the time path of factor prices, $\{w(t), R(t)\}_{t=0}^{\infty}$, is such that given the time path of capital stock and labor $\{K(t), L(t)\}_{t=0}^{\infty}$, all markets clear.*

In addition to the competitive factor prices, the key equation is the consumer Euler equation for an individual of generation τ at time t. Taking into account that the gross rate of return on savings is $1 + r(t+1) + \nu/(1-\nu)$ and that the effective discount factor of the individual is $\beta(1-\nu)$, this Euler equation can be written as

$$u'(c(t \mid \tau)) = \beta\left[(1 + r(t+1))(1-\nu) + \nu\right] u'(c(t+1 \mid \tau)). \tag{9.36}$$

This equation looks similar to be standard consumption Euler equation, for example as in Chapter 6. It only differs from the equation there because it applies separately to each generation τ and because the term ν, the probability of death facing each individual, features in this equation. Note, however, that when both r and ν are small

$$(1+r)(1-\nu) + \nu \approx 1 + r,$$

and the terms involving ν disappear. In fact, the reason why these terms are present is because of the discrete-time nature of the current model. In the next section, I present the continuous-time version of the perpetual youth model, where the approximation in the previous equation is exact. Moreover, the continuous-time model will allow us to obtain closed-form solutions for aggregate consumption and capital stock dynamics. Therefore, this model gives one example of a situation in which continuous-time methods turn out to be

more appropriate than discrete-time methods (whereas the baseline overlapping generations model required discrete time).

Recall that in the neoclassical model without technological progress, the consumer Euler equation admitted a simple solution because consumption had to be equal across dates for the representative household. This is no longer the case in the perpetual youth model, since different generations will have different levels of assets and may satisfy eq. (9.36) with different growth rates of consumption depending on the form of the utility function $u(\cdot)$.

To simplify the analysis, let us now suppose that the utility function takes the logarithmic form,

$$u(c) = \log c.$$

In that case, (9.36) simplifies to

$$\frac{c(t+1 \mid \tau)}{c(t \mid \tau)} = \beta\left[(1 + r(t+1))(1-\nu) + \nu\right], \tag{9.37}$$

and implies that the growth rate of consumption must be equal for all generations. Using this observation, it is possible to characterize the behavior of the aggregate capital stock, though this turns out to be much simpler in continuous time. For this reason, I now turn to the continuous-time version of this model (details on the discrete-time model with log preferences are covered in Exercise 9.22).

9.8. Overlapping Generations in Continuous Time

9.8.1. Demographics, Technology and Preferences. Consider a continuous-time version of the perpetual youth model. Suppose that each individual faces a Poisson death rate of $\nu \in (0, \infty)$. Suppose also that individuals have logarithmic preferences and a pure discount rate of $\rho > 0$. As demonstrated in Exercise 5.7 in Chapter 5, this implies that individual i will maximize the objective function

$$\int_0^\infty \exp\left(-(\rho+\nu)t\right) \log c_i(t)\, dt. \tag{9.38}$$

Demographics in this economy are similar to those in the discrete-time perpetual youth model of the previous section. In particular, expected further life of an individual is independent of when he was born and is equal to

$$\frac{1}{\nu} < \infty.$$

This is both the life expectancy at birth and the expected further life of an individual who has survived up to a certain point. Next, let population at time t be $L(t)$. Then, the Poisson death rate implies that a total flow of $\nu L(t)$ individuals will die at time t. Once again, there is arrival of new households at the exponential rate $n > \nu$, so that aggregate population dynamics are given by

$$\dot{L}(t) = (n-\nu) L(t). \tag{9.39}$$

It can also be computed that at time t the mass of individuals of cohort born at time $\tau < t$ is given by

$$L(t \mid \tau) = \exp(-\nu(t-\tau) + (n-\nu)\tau). \tag{9.40}$$

In this equation and throughout the section, suppose that at $t = 0$, the economy starts with a population of $L(0) = 1$ who are all newborn at that point. Equation (9.40) is derived in Exercise 9.23.

As in the previous section, it is sufficient to specify the consumption behavior and the budget constraints for each cohort. In particular, the flow budget constraint for cohort τ at time t is

$$\dot{a}(t \mid \tau) = r(t)\, a(t \mid \tau) - c(t \mid \tau) + w(t) + z(a(t \mid \tau) \mid t, \tau),$$

where again $z(a(t \mid \tau) \mid t, \tau)$ is the transfer payment or annuity payment at time t to an individual born at time τ holding assets $a(t \mid \tau)$. Let us follow Yaari and Blanchard and again assume complete annuity markets, with free entry. Now the instantaneous profits of a life insurance company providing such annuities at time t for an individual born at time τ with assets $a(t \mid \tau)$ is

$$\pi(a(t \mid \tau) \mid t, \tau) = \nu a(t \mid \tau) - z(a(t \mid \tau) \mid t, \tau),$$

since the individual will die and leave his assets to the life insurance company at the flow rate ν. The zero-profits condition now implies that

$$z(a(t \mid \tau) \mid t, \tau) = \nu a(t \mid \tau).$$

Substituting this into the flow budget constraint above, we obtain the more useful expression

$$\dot{a}(t \mid \tau) = (r(t) + \nu)\, a(t \mid \tau) - c(t \mid \tau) + w(t). \tag{9.41}$$

Let us assume that the production side is given by the per capita aggregate production function $f(k)$ satisfying Assumptions 1 and 2, where k is the aggregate capital-labor ratio. Capital is assumed to depreciate at the rate δ. Factor prices are given by the usual expressions

$$R(t) = f'(k(t)) \text{ and } w(t) = f(k(t)) - k(t) f'(k(t)), \tag{9.42}$$

and as usual $r(t) = R(t) - \delta$. The law of motion of in the capital-labor ratio is given by

$$\dot{k}(t) = f(k(t)) - (n - \nu + \delta)k(t) - c(t), \tag{9.43}$$

where $c(t)$ is average consumption per capita, given by

$$\begin{aligned} c(t) &= \frac{\int_{-\infty}^{t} c(t \mid \tau) L(t \mid \tau)\, d\tau}{\int_{-\infty}^{t} L(t \mid \tau)\, d\tau} \\ &= \frac{\int_{-\infty}^{t} c(t \mid \tau) L(t \mid \tau)\, d\tau}{L(t)}, \end{aligned}$$

where recall that $L(t \mid \tau)$ is the size of the cohort born at τ at time t.

9.8.2. Equilibrium. A competitive equilibrium is again defined in a straightforward manner.

DEFINITION 9.4. *A **competitive equilibrium** consists of paths of capital stock, wage rates and rental rates of capital,* $[K(t), w(t), R(t)]_{t=0}^{\infty}$ *and paths of consumption for each generation,* $[c(t \mid \tau)]_{t=0,\tau\leq t}^{\infty}$*, such that each individual maximizes (9.38) subject to (9.41), and the time path of prices,* $\{w(t), R(t)\}_{t=0}^{\infty}$*, are given by (9.42), and the capital-labor ratio evolves according to (9.43).*

Let us start with consumer optimization. The maximization of (9.38) subject to (9.41) gives the usual Euler equation

$$\frac{\dot{c}(t \mid \tau)}{c(t \mid \tau)} = r(t) - \rho, \tag{9.44}$$

where $\dot{c}(t \mid \tau) \equiv \partial c(t \mid \tau)/\partial t$. Notice that, in contrast to the discrete-time version of this equation, (9.37), the probability (flow rate) of death, ν, does not feature here, since it exactly cancels out (the rate of return on assets is $r(t) + \nu$ and the effective discount factor is $\rho + \nu$, so that their difference is equal to $r(t) - \rho$).

The transversality condition for an individual of cohort τ can be written as

$$\lim_{t\to\infty} \exp\left(-\left(\bar{r}(t,\tau) + \nu\right)\right) a(t \mid \tau) = 0, \tag{9.45}$$

where

$$\bar{r}(t,\tau) \equiv \frac{1}{t-\tau}\int_{\tau}^{t} r(s)\, ds$$

is the average interest rate between dates τ and t as in eq. (8.30) in Chapter 8, and the transversality condition here is the analog of eq. (8.31) there. The transversality condition, (9.45), requires the net present discounted value of the assets in the very far future of an individual born at the time τ discounted back to this time to be equal to 0.

Combining (9.44) together with (9.41) and (9.45) gives the following consumption "function" for an individual of cohort τ (see Exercise 9.24):

$$c(t \mid \tau) = (\rho + \nu)\left[a(t \mid \tau) + \omega(t)\right]. \tag{9.46}$$

This linear form of the consumption function is a particularly attractive feature of logarithmic preferences. The term in square brackets is the asset and human wealth of the individual, with the second term defined as

$$\omega(t) = \int_{t}^{\infty} \exp\left(-\left(\bar{r}(s,t) + \nu\right)\right) w(s)\, ds.$$

This term clearly represents the net present discounted value of future wage earnings of an individual discounted to time t. It is independent of τ, since the future expected earnings of all individuals are the same regardless of when they are born. The additional discounting with ν in this term arises because individuals will die at this rate and thus lose future earnings from then on.

Equation (9.46) implies that each individual consumes a fraction of this wealth equal to his effective discount rate, $\rho+\nu$. Now integrating this across cohorts, using the fact that the size of the cohort τ at time t is $\exp\left(-\nu\left(t-\tau\right)+\left(n-\nu\right)\tau\right)$, per capita consumption is

$$c\left(t\right)=\left(\rho+\nu\right)\left(a\left(t\right)+\omega\left(t\right)\right), \tag{9.47}$$

where $a\left(t\right)$ is average assets per capita. Since the only productive asset in this economy is capital, we also have that $a\left(t\right)=k\left(t\right)$. Finally, differentiating (9.47),

$$\dot{c}\left(t\right)=\left(\rho+\nu\right)\left(\dot{a}\left(t\right)+\dot{\omega}\left(t\right)\right). \tag{9.48}$$

The law of motion of assets per capita can be written as

$$\dot{a}\left(t\right)=\left(r\left(t\right)-\left(n-\nu\right)\right)a\left(t\right)+w\left(t\right)-c\left(t\right).$$

This equation is intuitive. Aggregate wealth, $a\left(t\right)L\left(t\right)$, increases because of the returns to capital at the rate $r\left(t\right)$ and also because of the wage income, $w\left(t\right)L\left(t\right)$. Out of this, total consumption of $c\left(t\right)L\left(t\right)$ needs to be subtracted. Finally, since $L\left(t\right)$ grows at the rate $n-\nu$, this reduces the rate of growth of assets per capita. Human wealth per capita, on the other hand, satisfies

$$\left(r\left(t\right)+\nu\right)\omega\left(t\right)=\dot{\omega}\left(t\right)+w\left(t\right).$$

The intuition for this equation comes from the Hamilton-Jacobi-Bellman equations discussed in Chapter 7. We can think of $\omega\left(t\right)$ as the value of an asset with a claim to the future earnings of a typical individual. The required rate of return on this is $r\left(t\right)+\nu$, which takes into account that the individual will lose his future earnings stream at the rate ν when he dies. The return on this asset is equal to its capital gains, $\dot{\omega}\left(t\right)$, and dividends, $w\left(t\right)$. Now substituting for $\dot{a}\left(t\right)$ and $\dot{\omega}\left(t\right)$ from these two equations into (9.48):

$$\begin{aligned}
\dot{c}\left(t\right) &= \left(\rho+\nu\right)\left(\left(r\left(t\right)-\left(n-\nu\right)\right)a\left(t\right)+w\left(t\right)-c\left(t\right)+\left(r\left(t\right)+\nu\right)\omega\left(t\right)-w\left(t\right)\right)\\
&= \left(\rho+\nu\right)\left(\left(r\left(t\right)+\nu\right)\left(a\left(t\right)+\omega\left(t\right)\right)-na\left(t\right)-c\left(t\right)\right)\\
&= \left(\rho+\nu\right)\left(\frac{\left(r\left(t\right)+\nu\right)}{\rho+\nu}c\left(t\right)-na\left(t\right)-c\left(t\right)\right)\\
&= \left(r\left(t\right)-\rho\right)c\left(t\right)-\left(\rho+\nu\right)na\left(t\right),
\end{aligned}$$

where the third line uses (9.47). Dividing both sides by $c\left(t\right)$, using the fact that $a\left(t\right)=k\left(t\right)$, and substituting $r\left(t\right)=f'\left(k\left(t\right)\right)-\delta$,

$$\frac{\dot{c}\left(t\right)}{c\left(t\right)}=f'\left(k\left(t\right)\right)-\delta-\rho-\left(\rho+\nu\right)n\frac{k\left(t\right)}{c\left(t\right)}. \tag{9.49}$$

This is similar to the standard Euler equation (under logarithmic preferences), except for the last term. This last term reflects the fact that consumption growth per capita is slowed down by the arrival of new individuals at each instance, who have less wealth than the average individual. Their lower wealth implies lower consumption and reduces average consumption growth in the economy. This intuitively explains why the last term depends on n (the rate of arrival of new individuals) and on k/c (the size of average asset holdings relative to consumption).

The equilibrium path of the economy is completely characterized by the two differential equations, (9.43) and (9.49)—together with an initial condition for $k(0) > 0$ and the transversality condition (9.45) applied to average assets, thus to the capital-labor ratio, $k(t)$. First, a steady-state equilibrium is obtained when both $\dot{k}(t)/k(t)$ and $\dot{c}(t)/c(t)$ are equal to zero, and thus satisfies the following two equations:

$$\frac{c^*}{k^*} = \frac{(\rho+\nu)n}{f'(k^*)-\delta-\rho} \tag{9.50}$$

$$\frac{f(k^*)}{k^*} - (n-\nu+\delta) - \frac{(\rho+\nu)n}{f'(k^*)-\delta-\rho} = 0. \tag{9.51}$$

The second equation pins down a unique positive level of steady-state capital-labor ratio, k^*, ratio (since both $f(k)/k$ and $f'(k)$ are decreasing). Given k^* the first equation pins down a unique level of average consumption per capita, c^*. It can also be verified that at k^*,

$$f'(k^*) > \rho + \delta,$$

so that the capital-labor ratio is lower than the level consistent with the modified golden rule k_{mgr}, given by $f'(k_{mgr}) = \rho + \delta$. Recall that optimal steady-state capital-labor ratio of the neoclassical growth model satisfied the modified golden rule. In comparison, in this economy there is always *underaccumulation.* This contrasts with the baseline overlapping generations model, which potentially led to dynamic inefficiency and overaccumulation.

Let us next turn to equilibrium dynamics. Figure 9.3 plots (9.43) and (9.49), together with the arrows indicating how average consumption per capita and capital-labor ratio change in different regions. Both (9.43) and (9.49) are upward sloping and start at the origin. It is also straightforward to verify that while (9.43) is concave in the k-c space, (9.49) is convex. Thus they have a unique intersection. We also know from the preceding discussion that this unique intersection is at a capital-labor ratio less than that satisfying the modified golden rule, which is marked as k_{mgr} in the figure. Naturally, k_{mgr} itself is less than k_{gold}. The phase diagram also makes it clear that there exists a unique stable arm that is upward sloping in the k-c space. The shape of the stable arm is the same as in the basic neoclassical growth model. If the initial level of consumption is above this stable arm, feasibility is violated, while if it is below, the economy tends towards zero consumption and violates the transversality condition. Consequently, the steady-state equilibrium is globally saddle-path stable; consumption starts along the stable arm, and consumption and the capital-labor ratio monotonically converge to the steady state. Exercise 9.26 asks you to verify local saddle-path stability by linearizing (9.43) and (9.49) around the steady state.

The following proposition summarizes this analysis.

Proposition 9.10. *In the continuous-time perpetual youth model, there exists a unique steady state (k^*, c^*) given by (9.50) and (9.51). The level of capital-labor ratio is less than the level of capital-labor ratio that satisfies the modified golden rule, k_{mgr}. Starting with any $k(0) > 0$, equilibrium dynamics monotonically converge to this unique steady state.*

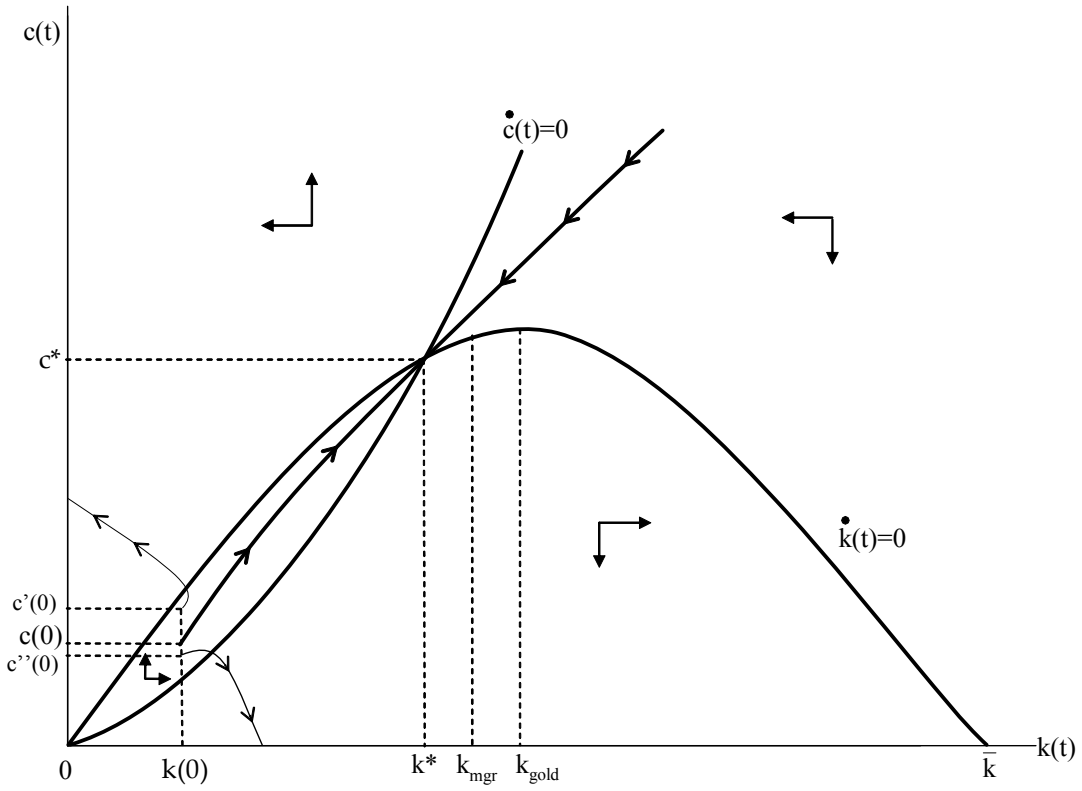

FIGURE 9.3. Steady state and transitional dynamics in the overlapping generations model in continuous time.

Perhaps the most interesting feature of this equilibrium is that, despite finite lives and overlapping generations, there is no overaccumulation. The reason for this is that individuals have constant stream of labor income throughout their lives and thus do not need to save excessively in order to ensure smooth consumption. Is it possible to obtain overaccumulation in the continuous-time perpetual youth model? The answer is yes and is demonstrated by Blanchard (1985). He assumes that each individual starts life with one unit of effective labor and then his effective labor units decline at some positive exponential rate $\zeta > 0$ throughout his life, so that the labor earnings of an individual of generation τ at time t is $\exp(-\zeta(t-\tau))w(t)$, where $w(t)$ is the market wage per unit of effective labor at time t. Consequently, the individual consumption function changes from (9.46) to

$$c(t \mid \tau) = (\rho + \nu)[a(t \mid \tau) + \omega(t \mid \tau)],$$

where now $\omega(t \mid \tau)$ is the human wealth of an individual of generation τ at time t, given by (see Exercise 9.28):

$$\omega(t \mid \tau) = \int_t^{\infty} \exp(-(\bar{r}(t-s)+\nu)) \exp(-\zeta(s-\tau)) w(s)\, ds, \tag{9.52}$$

where $\exp(-\zeta(s-\tau))$ is the correction factor taking into account the decline in effective labor units. Repeating the same steps as before with this new expression for human wealth,

$$\frac{\dot{c}(t)}{c(t)} = f'(k(t)) - \delta - \rho - \zeta - (\rho+\nu)(n+\zeta)\frac{k(t)}{c(t)}, \tag{9.53}$$

while the behavior of capital-labor ratio is still given by (9.43). It can now be shown that for ζ sufficiently large, the steady-state capital-labor ratio k^* can exceed both the modified golden rule level k_{mgr} and the golden rule level, k_{gold} (see Exercise 9.28).

This discussion therefore illustrates that overaccumulation emerges when there are both overlapping generations and a strong motive for saving for the future. Interestingly, it can be shown that what is important is not finite lives per se, but overlapping generations indeed. In particular, Exercise 9.30 shows that when $n = 0$, overaccumulation is not possible so that finite lives is not sufficient for overaccumulation. However, $k^* > k_{gold}$ is possible when $n > 0$ and $\nu = 0$, so that the overlapping generations model with infinite lives can also generate overaccumulation.

9.9. Taking Stock

This chapter has continued our investigation of the mechanics of capital accumulation in dynamic equilibrium models. The main departure from the baseline neoclassical growth model of the last section has been the relaxation of the representative household assumption. The simplest way of accomplishing this is to introduce two-period lived overlapping generations (without pure altruism). In the baseline overlapping generations model of Samuelson and Diamond, each individual lives for two periods, but can only supply labor during the first period of his life. I have also discussed alternative non-representative-household models, in particular, overlapping generations with impure altruism and models of perpetual youth. In models of overlapping generations with impure altruism, individuals transfer resources to their offspring, but they do not care directly about the utility of their offspring and instead derive utility from the "act of giving" or from some subcomponent of the consumption vector of their descendent. In models of perpetual youth, the economy features expected finite life and overlapping generations, but each individual still has an infinite planning horizon, because the time of death is uncertain.

All of these models fall outside the scope of the First Welfare Theorem. As a result, there is no guarantee that the resulting equilibrium path will be Pareto optimal. In fact, the extensive study of the baseline overlapping generations models were partly motivated by the possibility of Pareto suboptimal allocations in such models. We have seen that these equilibria may be "dynamically inefficient" and feature overaccumulation—a steady-state capital-labor ratio greater than the golden rule capital-labor ratio. We have also seen how an unfunded Social Security system can reduce aggregate savings and thus ameliorate the overaccumulation problem. The important role that unfunded Social Security (or national debt) plays in the overlapping generations model has made this model a workhorse for analysis of transfer programs, fiscal policies and generational accounting.

Our analysis of perpetual youth models, especially Yaari and Blanchard's continuous-time perpetual youth model, further clarified the roles of the path of labor income, finite horizons and arrival of new individuals in generating the overaccumulation result. In particular, this model shows that the declining path of labor income is important for the overaccumulation result (in the Samuelson-Diamond model there is an extreme form of this, since there is no labor income in the second period of the life of the individual). But perhaps the more important insight generated by these models is that what matters is not finite horizons per se, but the arrival of new individuals. Overaccumulation and Pareto suboptimality arise because of the pecuniary externalities created on individuals that are not yet in the marketplace.

While overaccumulation and dynamic inefficiency have dominated much of the discussion of overlapping generations models in the literature, one should not overemphasize the importance of dynamic inefficiency. As emphasized in Chapter 1, the major question of economic growth is why so many countries have so little capital for their workers and why the process of economic growth and capital accumulation started only over the past 200 years. It is highly doubtful that overaccumulation is a major problem for most countries in the world.

The models presented in this chapter are very useful for another reason, however. They significantly enrich our arsenal in the study of the mechanics of economic growth and capital accumulation. All three of the major models presented in this chapter, the baseline overlapping generations model, the overlapping generations model with impure altruism, and the perpetual youth model, are tractable and useful vehicles for the study of economic growth in a variety of circumstances. For example, the first two lead to equilibrium dynamics similar to the baseline Solow growth model, but without explicitly imposing an exogenously constant saving rate. The latter model, on the other hand, allows an analysis of equilibrium dynamics similar to the basic neoclassical growth model, but also incorporates finite lives and overlapping generations, which will be essential in many problems, for example, in the context ofn human capital investments studied in the next chapter.

In summary, this chapter has provided us with new modeling tools and new perspectives on the question of capital accumulation, aggregate saving and economic growth. It has not, however, offered new answers to questions of why countries grow (for example, technological progress) and why some countries are much poorer than others (related to the fundamental causes of income differences). Of course, its purpose was not to provide such answers in the first place.

9.10. References and Literature

The baseline overlapping generations model with two-period lived agents is due to Samuelson (1958) and Diamond (1965). A related model appears in French in Maurice Allais' work. Blanchard and Fischer (1989, Chapter 3) provide an excellent textbook treatment of the baseline overlapping generations model. Some textbooks use this setup as the main workhorse macroeconomic model, for example, Azariadis (1993), McCandless and Wallace (1991) and De La Croix and Michel (2002).

The economy studied in Section 9.1 is due to Shell (1974). The source of inefficiency in the overlapping generations model is much discussed in the literature. Shell's (1974) example economy in Section 9.1 provides the clearest intuitive explanation for why the First Welfare Theorem does not apply. A lucid discussion is contained in Bewley (2006).

The model of overlapping generations with impure altruism is due to Andreoni (1989). This model has been used extensively in the economic growth and economic development literatures, especially for the analysis of equilibrium dynamics in the presence of imperfect capital markets. Well-known examples include the models by Aghion and Bolton (1996), Banerjee and Newman (1989, 1994), Galor and Zeira (1993) and Piketty (1996), which will be studied in Chapter 21. I am not aware of an analysis of wealth inequality dynamics with perfect markets in this economy along the lines of the model presented in Section 9.6, even though the analysis is quite straightforward. A similar analysis of wealth inequality dynamics is included in Stiglitz's (1979) model, but he assumes that each household can only use its savings in its own diminishing return technology (thus creating a strong force towards convergence of incomes).

The continuous-time perpetual youth model is due to Yaari (1965) and Blanchard (1985). The discrete-time version of this model was presented to facilitate the transition to the continuous-time version. My treatment of the continuous-time version closely followed Blanchard (1985). Blanchard and Fischer (1989, Chapter 3) and Barro and Sala-i-Martin (2004, Chapter 3) provide alternative textbook treatments. The importance of the path of labor income is emphasized and analyzed in Blanchard (1985). The importance of new arrivals in the market is emphasized and explained in Weil (1989).

Models with overlapping generations and finite lives are used extensively in the analysis of Ricardian Equivalence, introduced in Exercise 8.33 in Chapter 8, is a good approximation to reality. Blanchard (1985) and Bernheim (1987) include extensive discussions of this issue. Another important application of overlapping generations models is to generational accounting, for example, as in the work by Auerbach and Kotlikoff (1987).

9.11. Exercises

Exercise 9.1. Prove that the allocation characterized in Proposition 9.1 is the unique competitive equilibrium. [Hint: first, show that there cannot be any equilibrium with $p_j > p_{j-1}$ for any j. Second, show that even if $p_0 > p_1$, household $i = 0$ must consume only commodity $j = 0$; then inductively, show that this is true for any household].

Exercise 9.2. Consider the following variant of economy with infinite number of commodities and infinite number of individuals presented in Section 9.1. The utility of individual indexed $i = j$ is

$$u(c(j)) + \beta u(c(j+1))$$

where $\beta \in (0,1)$, and each individual has one unit of the good with the same index as his own.

(1) Define a competitive equilibrium for this economy.

(2) Characterize the set of competitive equilibria in this economy.
(3) Characterize the set of Pareto optima in this economy.
(4) Can all Pareto optima be decentralized without changing endowments? Can they be decentralized by changing endowments?

Exercise 9.3. Show that in the model of Section 9.2 the dynamics of capital stock are identical to those derived in the text even when $\delta < 1$.

Exercise 9.4. In the baseline overlapping generations model, verify that savings $s(w, R)$, given by (9.6), are increasing in the first argument, w. Provide conditions on the utility function $u(\cdot)$ such that they are also increasing in their second argument (in the interest rate R).

Exercise 9.5. Prove Proposition 9.4

Exercise 9.6. Consider the canonical overlapping generations model with log preferences

$$\log(c_1(t)) + \beta \log(c_2(t))$$

for each household. Suppose that there is population growth at the rate n. Individuals work only when they are young, and supply one unit of labor inelastically. Production technology is given by

$$Y(t) = A(t) K(t)^{\alpha} L(t)^{1-\alpha},$$

where $A(t+1) = (1+g) A(t)$, with $A(0) > 0$ and $g > 0$.

(1) Define a competitive equilibrium and the steady-state equilibrium.
(2) Characterize the steady-state equilibrium and show that it is globally stable.
(3) What is the effect of an increase in g on the equilibrium?
(4) What is the effect of an increase in β on the equilibrium? Provide an intuition for this result.

Exercise 9.7. Consider the canonical model with log preferences, $\log(c_1(t)) + \beta \log(c_2(t))$, and the general neoclassical technology $F(K, L)$ satisfying Assumptions 1 and 2. Show that multiple steady-state equilibria are possible in this economy.

Exercise 9.8. Consider again the canonical overlapping generations model with log preferences and Cobb-Douglas production function.

(1) Define a competitive equilibrium.
(2) Characterize the competitive equilibrium and derive explicit conditions under which the steady-state equilibrium is dynamically inefficient.
(3) Using plausible numbers discuss whether dynamic inefficiency can arise in "realistic" economies.
(4) Show that when there is dynamic inefficiency, it is possible to construct an unfunded Social Security system which creates a Pareto improvement relative to the competitive allocation.

Exercise 9.9. Consider again the canonical overlapping generations model with log preferences and Cobb-Douglas production function, but assume that individuals now work in both periods of their lives.

(1) Define a competitive equilibrium and the steady-state equilibrium.
(2) Characterize the steady-state equilibrium and the transitional dynamics in this economy.
(3) Can this economy generate overaccumulation?

Exercise 9.10. Prove Proposition 9.7.

Exercise 9.11. Consider the overlapping generations model with fully funded Social Security. Prove that even when the restriction $s(t) \geq 0$ for all t is imposed, no fully funded Social Security program can lead to a Pareto improvement.

Exercise 9.12. Consider an overlapping generations economy with the dynamically inefficient steady-state equilibrium. Show that the government can improve the allocation of resources by introducing national debt. [Hint: suppose that the government borrows from the current young and redistributes to the current old, paying back the current young the following period with another round of borrowing]. Contrast this result with the Ricardian equivalence result in Exercise 8.33 in Chapter 8.

Exercise 9.13. Prove Proposition 9.8.

Exercise 9.14. Consider the baseline overlapping generations model and suppose that the equilibrium is dynamically efficient, i.e., $r^* > n$. Show that any unfunded Social Security system will increase the welfare of the current old generation and reduce the welfare of some future generation.

Exercise 9.15. Derive eq. (9.31).

Exercise 9.16. Consider the overlapping generations model with warm glow preferences in Section 9.6, and suppose that preferences are given by

$$c(t)^{\eta} b(t+1)^{1-\eta},$$

with $\eta \in (0,1)$, instead of eq. (9.21). The production side is the same as in Section 9.6. Characterize the dynamic equilibrium of this economy.

Exercise 9.17. Consider the overlapping generations model with warm glow preferences in Section 9.6, and suppose that preferences are given by $u_1(c_i(t)) + u_2(b_i(t))$, where u_1 and u_2 are strictly increasing and concave functions. The production side is the same as in the text. Characterize a dynamic equilibrium of this economy.

Exercise 9.18. Characterize the aggregate equilibrium dynamics and the dynamics of wealth distribution in the overlapping generations model with warm glow preferences as in Section 9.6, when the per capita production function is given by the Cobb-Douglas form $f(k) = Ak^{\alpha}$. Show that away from the steady state, there can be periods during which wealth inequality increases. Explain why this may be the case.

Exercise 9.19. * Generalize the results in Section 9.6 to an environment in which the preferences of an individual of generation t are given by

$$u(c(t)) + v(b(t)),$$

where $c(t)$ denotes own consumption, $b(t)$ is bequests, and $u(\cdot)$ and $v(\cdot)$ are strictly increasing, differentiable and strictly concave utility functions. Determine conditions on $u(\cdot)$ and

$v(\cdot)$ such that aggregate dynamics are globally stable. Provide conditions on $u(\cdot)$ and $v(\cdot)$ to ensure that asymptotically all individuals tend to the same wealth level.

Exercise 9.20. Show that the steady-state capital labor ratio in the overlapping generations model with impure altruism (of Section 9.6) can lead to overaccumulation, i.e., to $k^* > k_{gold}$.

Exercise 9.21. Prove that given the perpetual youth assumption and population dynamics in eq. (9.34), at time $t > 0$, there will be $n(1+n-\nu)^{t-s}(1-\nu)^{s-1}$ s-year-olds for any $s \in \{1, 2, ..., t-1\}$.

Exercise 9.22. * Consider the discrete-time perpetual youth model discussed in Section 9.7 and assume that preferences are logarithmic. Characterize the steady-state equilibrium and the equilibrium dynamics of the capital-labor ratio.

Exercise 9.23. Consider the continuous-time perpetual youth model of Section 9.8.

(1) Show that given $L(0) = 1$, the initial size of a cohort born at the time $\tau \geq 0$ is $\exp((n-\nu)\tau)$.
(2) Show that the probability that an individual born at the time τ is alive at time $t \geq \tau$ is $\exp(-\nu(t-\tau))$.
(3) Derive eq. (9.40).
(4) Show that this equation would not apply at any finite time if the economy starts at $t = 0$ with an arbitrary age distribution.

Exercise 9.24. Derive eq. (9.46). [Hint: first integrate the flow budget constraint of the individual, (9.41) using the transversality condition (9.45) and then use the Euler equation (9.44)].

Exercise 9.25. Generalize the analysis of the continuous-time perpetual youth model of Section 9.8 to an economy with labor-augmenting technological progress at the rate g. Prove that the steady-state equilibrium is unique and globally (saddle-path) stable. What is the impact of a higher rate of technological progress?

Exercise 9.26. Linearize the differential equations (9.43) and (9.49) around the steady state, (k^*, c^*), and show that the linearized system has one negative and one positive eigenvalue.

Exercise 9.27. Determine the effects of n and ν on the steady-state equilibrium (k^*, c^*) in the continuous-time perpetual youth model of Section 9.8.

Exercise 9.28. (1) Derive eq.'s (9.52) and (9.53).

(2) Show that for ζ sufficiently large, the steady-state equilibrium capital-labor ratio, k^*, can exceed k_{gold}, so that there is overaccumulation. Provide an intuition for this result.

Exercise 9.29. Consider the continuous-time perpetual youth model with a constant flow of government spending G. Suppose that this does not affect consumer utility and that lump-sum taxes $[\mathcal{T}(t)]_{t=0}^{\infty}$ are allowed. Specify the government budget constraint as in Exercise 8.33 in Chapter 8. Prove that contrary to the Ricardian Equivalence result in Exercise 8.33, the sequence of taxes affects the equilibrium path of capital-labor ratio and consumption. Interpret this result and explain the difference between the overlapping generations model and the neoclassical growth model.

Exercise 9.30. * Consider the continuous-time perpetual youth model with labor income declining at the rate $\zeta > 0$.

(1) Show that if $n = 0$, $k^* \leq k_{gold}$ for any $\zeta > 0$.

(2) Show that there exists $\zeta > 0$ sufficiently high such that if $n > 0$ and $\nu = 0$, $k^* > k_{gold}$.

Exercise 9.31. Consider an economy with aggregate production function

$$Y(t) = AK(t)^{1-\alpha} L(t)^{\alpha}.$$

All markets are competitive, the labor supply is normalized to 1, capital fully depreciates after use, and the government imposes a linear tax on capital income at the rate τ and uses the proceeds for government consumption. Consider two specifications of preferences:

- All agents are infinitely lived, with preferences

$$\sum_{t=0}^{\infty} \beta^t \ln c(t).$$

- An overlapping generations model where agents work in the first period, and consume the capital income from their savings in the second period. The preferences of a generation born at time t, defined over consumption when young and old, are given by

$$\ln c^y(t) + \beta \ln c^0(t).$$

(1) Characterize the equilibria in these two economies, and show that in the first economy, taxation reduces output, while in the second, it does not.

(2) Interpret this result, and in the light of this result discuss the applicability of models that try to explain income differences across countries with differences in the rates of capital income taxation.

CHAPTER 10

Human Capital and Economic Growth

This chapter investigates human capital investments and the role of human capital in economic growth and in cross-country income differences. As already noted in Chapter 3, human capital can play a major role in economic growth and cross-country income differences, and our main purpose is to understand which factors affect human capital investments and how these influence the process of economic growth and economic development. Human capital refers to all the attributes of workers that potentially increase their productivity in all or some productive tasks. The term is coined because much of these attributes are accumulated by workers through investments. Human capital theory, developed primarily by Becker (1965) and Mincer (1974), is about the role of human capital in the production process and about the incentives to invest in skills, including pre-labor market investments (in the form of schooling) and on-the-job investments (in the form of training). It would not be an exaggeration to say that this theory is the basis of much of labor economics and plays an equally important role in macroeconomics. The literature on education and other types of human capital investments is vast, so only parts of this literature that are relevant to the main focus of this book will be covered here. There are a number of other important connections between human capital and economic growth, especially related to its effect on technological progress and its role in economic takeoff, which will be discussed later in the book.

10.1. A Simple Separation Theorem

Let us start with the partial equilibrium schooling decisions and establish a simple general result, sometimes referred to as a "separation theorem" for human capital investments. I set up the basic model in continuous time for simplicity.

Consider the schooling decision of a single individual facing exogenously given prices for human capital. Throughout, I assume that there are perfect capital markets. The separation theorem referred to in the title of this section will show that, with perfect capital markets, schooling decisions will maximize the net present discounted value of the individual (I return to human capital investments with imperfect capital markets in Chapter 21). In particular, consider an individual with an instantaneous utility function $u(c)$ that satisfies Assumption 3 above. Suppose that the individual has a planning horizon of T (where $T = \infty$ is allowed), discounts the future at the rate $\rho > 0$ and faces a constant flow rate of death equal to $\nu \geq 0$ (as in the perpetual youth model studied in the previous chapter). Standard arguments imply

that the objective function of this individual at time $t = 0$ is

$$\max \int_0^T \exp\left(-\left(\rho+\nu\right)t\right) u\left(c\left(t\right)\right) dt. \tag{10.1}$$

Now suppose that this individual is born with some human capital $h\left(0\right) \geq 0$. Suppose that his human capital evolves over time according to the differential equation

$$\dot{h}\left(t\right) = G\left(t, h\left(t\right), s\left(t\right)\right), \tag{10.2}$$

where $s\left(t\right) \in \left[0,1\right]$ is the fraction of time that the individual spends for human capital investments (i.e., "schooling") and $G : \mathbb{R}_+^2 \times \left[0,1\right] \to \mathbb{R}_+$ determines how human capital evolves as a function of time, the individual's stock of human capital and schooling decisions. In addition, we may impose a further restriction on schooling decisions, for example,

$$s\left(t\right) \in \mathcal{S}\left(t\right), \tag{10.3}$$

where $\mathcal{S}\left(t\right) \subset \left[0,1\right]$ and captures the fact that all schooling may have to be full-time, i.e., $s\left(t\right) \in \left\{0,1\right\}$, or that there may exist other restrictions on schooling decisions.

The individual is assumed to face an exogenous process of wage per unit of human capital given by $\left[w\left(t\right)\right]_{t=0}^T$, so that his labor earnings at time t are

$$W\left(t\right) = w\left(t\right)\left[1 - s\left(t\right)\right]\left[h\left(t\right) + \omega\left(t\right)\right],$$

where $1 - s\left(t\right)$ is the fraction of time spent supplying labor to the market and $\omega\left(t\right)$ is non-human capital labor that the individual may be supplying to the market at time t. The process of non-human capital labor that the individual can supply to the market, $\left[\omega\left(t\right)\right]_{t=0}^T$, is exogenous. This formulation assumes that the only margin of choice is between market work and schooling (that is, there is no leisure).

Finally, let us assume that the individual faces a constant (flow) interest rate equal to r on his savings (potentially including annuity payments as discussed in the previous chapter). Using the equation for labor earnings, the lifetime budget constraint of the individual can be written as

$$\int_0^T \exp\left(-rt\right) c\left(t\right) dt \leq \int_0^T \exp\left(-rt\right) w\left(t\right)\left[1 - s\left(t\right)\right]\left[h\left(t\right) + \omega\left(t\right)\right] dt. \tag{10.4}$$

The Separation Theorem, which is the subject of this section, can be stated as follows:

Theorem 10.1. ***(Separation Theorem)*** *Suppose that the instantaneous utility function $u\left(\cdot\right)$ is strictly increasing. Then, $\left[\hat{c}\left(t\right), \hat{s}\left(t\right), \hat{h}\left(t\right)\right]_{t=0}^T$ is a solution to the maximization of (10.1) subject to (10.2), (10.3) and (10.4) if and only if $\left[\hat{s}\left(t\right), \hat{h}\left(t\right)\right]_{t=0}^T$ maximizes*

$$\int_0^T \exp\left(-rt\right) w\left(t\right)\left[1 - s\left(t\right)\right]\left[h\left(t\right) + \omega\left(t\right)\right] dt \tag{10.5}$$

subject to (10.2) and (10.3), and $\left[\hat{c}\left(t\right)\right]_{t=0}^T$ maximizes (10.1) subject to (10.4) given $\left[\hat{s}\left(t\right), \hat{h}\left(t\right)\right]_{t=0}^T$. That is, human capital accumulation and supply decisions can be separated *from consumption decisions.*

PROOF. To prove the "only if" part, suppose that $\left[\hat{s}(t), \hat{h}(t)\right]_{t=0}^{T}$ does not maximize (10.5), but there exists $\hat{c}(t)$ such that $\left[\hat{c}(t), \hat{s}(t), \hat{h}(t)\right]_{t=0}^{T}$ is a solution to (10.1). Let the value of (10.5) generated by $\left[\hat{s}(t), \hat{h}(t)\right]_{t=0}^{T}$ be denoted Y. Since $\left[\hat{s}(t), \hat{h}(t)\right]_{t=0}^{T}$ does not maximize (10.5), there exists $[s(t), h(t)]_{t=0}^{T}$ reaching a value of (10.5), $Y' > Y$. Consider $[c(t), s(t), h(t)]_{t=0}^{T}$, where $c(t) = \hat{c}(t) + \varepsilon$. By the hypothesis that $\left[\hat{c}(t), \hat{s}(t), \hat{h}(t)\right]_{t=0}^{T}$ is a solution to (10.1), the budget constraint (10.4) implies

$$\int_0^T \exp(-rt)\,\hat{c}(t)\,dt \leq Y.$$

Let $\varepsilon > 0$ and consider $c(t) = \hat{c}(t) + \varepsilon$ for all t. We have that

$$\begin{aligned}\int_0^T \exp(-rt)\,c(t)\,dt &= \int_0^T \exp(-rt)\,\hat{c}(t)\,dt + \frac{[1-\exp(-rT)]}{r}\varepsilon. \\ &\leq Y + \frac{[1-\exp(-rT)]}{r}\varepsilon.\end{aligned}$$

Since $Y' > Y$, for ε sufficiently small, $\int_0^T \exp(-rt)\,c(t)\,dt \leq Y'$ and thus $[c(t), s(t), h(t)]_{t=0}^{T}$ is feasible. Since $u(\cdot)$ is strictly increasing, $[c(t), s(t), h(t)]_{t=0}^{T}$ is strictly preferred to $\left[\hat{c}(t), \hat{s}(t), \hat{h}(t)\right]_{t=0}^{T}$, leading to a contradiction and proving the "only if" part.

The proof of the "if" part is similar. Suppose that $\left[\hat{s}(t), \hat{h}(t)\right]_{t=0}^{T}$ maximizes (10.5). Let the maximum value be denoted by Y. Consider the maximization of (10.1) subject to the constraint that $\int_0^T \exp(-rt)\,c(t)\,dt \leq Y$. Let $[\hat{c}(t)]_{t=0}^{T}$ be a solution. This implies that if $[c'(t)]_{t=0}^{T}$ is strictly preferred to $[\hat{c}(t)]_{t=0}^{T}$, then $\int_0^T \exp(-rt)\,c'(t)\,dt > Y$. This implies that $\left[\hat{c}(t), \hat{s}(t), \hat{h}(t)\right]_{t=0}^{T}$ must be a solution to the original problem, because any other $[s(t), h(t)]_{t=0}^{T}$ leads to a value of (10.5) $Y' \leq Y$, and if $[c'(t)]_{t=0}^{T}$ is strictly preferred to $[\hat{c}(t)]_{t=0}^{T}$, then $\int_0^T \exp(-rt)\,c'(t)\,dt > Y \geq Y'$ for any Y' associated with any feasible $[s(t), h(t)]_{t=0}^{T}$. □

The intuition for this theorem is straightforward: in the presence of perfect capital markets, the best human capital accumulation decisions are those that maximize the lifetime budget set of the individual. Exercise 10.2 shows that this theorem does not hold when there are imperfect capital markets and also does not generalize to the case where leisure is also an argument of the utility function.

10.2. Schooling Investments and Returns to Education

Let us next turn to the simplest model of schooling decisions in partial equilibrium, which will illustrate the main tradeoffs in human capital investments. The model presented here is a version of Mincer's (1974) seminal contribution. This model also enables a simple mapping from the theory of human capital investments to the large empirical literature on returns to schooling.

Let us first assume that $T = \infty$, which will simplify the expressions. The flow rate of death, ν, is positive, so that individuals have finite expected lives. Suppose that (10.2) is such that the individual has to spend an interval S with $s(t) = 1$—that is, in full-time schooling, and $s(t) = 0$ thereafter. At the end of the schooling interval, the individual will have a schooling level of

$$h(S) = \eta(S),$$

where $\eta(\cdot)$ is a continuous, increasing, differentiable and concave function. For $t \in [S, \infty)$, human capital accumulates over time (as the individual works) according to the differential equation

$$\dot{h}(t) = g_h h(t), \tag{10.6}$$

for some $g_h \geq 0$. Suppose also that wages grow exponentially,

$$\dot{w}(t) = g_w w(t), \tag{10.7}$$

with initial value $w(0) > 0$.

Suppose that

$$g_w + g_h < r + \nu,$$

so that the net present discounted value of the individual is finite. Now using Theorem 10.1, the optimal schooling decision must be a solution to the following maximization problem

$$\max_S \int_S^\infty \exp(-(r+\nu)t)\, w(t)\, h(t)\, dt. \tag{10.8}$$

Using (10.6) and (10.7), this is equivalent to (see Exercise 10.3):

$$\max_S \frac{\eta(S)\, w(0) \exp(-(r+\nu-g_w)S)}{r+\nu-g_h-g_w}. \tag{10.9}$$

Since $\eta(S)$ is concave, the objective function in (10.9) is strictly concave. Therefore, the unique solution to this problem is characterized by the first-order condition

$$\frac{\eta'(S^*)}{\eta(S^*)} = r + \nu - g_w. \tag{10.10}$$

Equation (10.10) shows that higher interest rates and higher values of ν (corresponding to shorter planning horizons) reduce human capital investments, while higher values of g_w increase the value of human capital and thus encourage further investments.

Integrating both sides of this equation with respect to S,

$$\ln \eta(S^*) = \text{constant} + (r + \nu - g_w) S^*. \tag{10.11}$$

Now note that the wage earnings of the worker of age $\tau \geq S^*$ in the labor market at time t will be given by

$$W(S,t) = \exp(g_w t) \exp(g_h(t-S))\, \eta(S).$$

Taking logs and using eq. (10.11) implies that the earnings of the worker will be

$$\ln W(S^*, t) = \text{constant} + (r + \nu - g_w) S^* + g_w t + g_h(t - S^*),$$

where $t - S$ can be thought of as worker experience (time after schooling). If we make a cross-sectional comparison across workers, the time trend term $g_w t$, will also go into the constant, so that we obtain the canonical Mincer equation where, in the cross section, log wage earnings are proportional to schooling and experience. Written differently, we have the following cross-sectional equation

$$\ln W_j = \text{ constant} + \gamma_s S_j + \gamma_e \text{experience}, \tag{10.12}$$

where j refers to individual j. Note however that there is so far no source of heterogeneity that can generate different levels of schooling across individuals. Nevertheless, eq. (10.12) is important, since it is the typical empirical model for the relationship between wages and schooling estimated in labor economics.

The economic insight provided by this equation is quite important; it suggests that the functional form of the Mincerian wage equation is not just a mere coincidence, but has economic content: the opportunity cost of one more year of schooling is foregone earnings. This implies that the benefit has to be commensurate with these foregone earnings, thus should lead to a proportional increase in earnings in the future. In particular, this proportional increase should be at the rate $(r + \nu - g_w)$.

As already discussed in Chapter 3, empirical work using equations of the form (10.12) leads to estimates for γ in the range of 0.06 to 0.10. Equation (10.12) suggests that these returns to schooling are not unreasonable. For example, we can think of the annual interest rate r as approximately 0.10, ν as corresponding to 0.02 that gives an expected life of 50 years, and g_w corresponding to the rate of wage growth holding the human capital level of the individual constant, which should be approximately about 2%. Thus we should expect an estimate of γ around 0.10, which is consistent with the upper range of the empirical estimates.

10.3. The Ben-Porath Model

The baseline Ben-Porath model enriches the model studied in the previous section by allowing human capital investments and non-trivial labor supply decisions throughout the lifetime of the individual. In particular, we now let $s(t) \in [0, 1]$ for all $t \geq 0$. Together with the Mincer eq. (10.12) (and the model underlying this equation presented in the previous section), the Ben-Porath model is the basis of much of labor economics. Here it is sufficient to consider a simple version of this model where the human capital accumulation equation, (10.2), takes the form

$$\dot{h}(t) = \phi(s(t) h(t)) - \delta_h h(t), \tag{10.13}$$

where $\delta_h > 0$ captures "depreciation of human capital," for example because new machines and techniques are being introduced, eroding the existing human capital of the worker. The individual starts with an initial value of human capital $h(0) > 0$. The function $\phi : \mathbb{R}_+ \to \mathbb{R}_+$ is strictly increasing, differentiable and strictly concave. Furthermore, let us simplify the

analysis by assuming that this function satisfies the Inada-type conditions,

$$\lim_{x\to 0}\phi'(x)=\infty \text{ and } \lim_{x\to h(0)}\phi'(x)=0.$$

The latter condition makes sure that we do not have to impose additional constraints to ensure $s(t)\in[0,1]$ (see Exercise 10.6).

Let us also suppose that there is no non-human capital component of labor, so that $\omega(t)=0$ for all t, that $T=\infty$, and that there is a flow rate of death $\nu>0$. Finally, let us assume that the wage per unit of human capital is constant at w and the interest rate is constant and equal to r. We also normalize $w=1$ without loss of any generality.

Again using Theorem 10.1, human capital investments can be determined as a solution to the following problem

$$\max\int_0^\infty \exp(-(r+\nu))(1-s(t))h(t)\,dt$$

subject to (10.13).

Let us next set up the current-value Hamiltonian, which in this case takes the form

$$\mathcal{H}(h,s,\mu)=(1-s(t))h(t)+\mu(t)(\phi(s(t)h(t))-\delta_h h(t)),$$

where we used $\mathcal{H}$ to denote the Hamiltonian to avoid confusion with human capital. The necessary and sufficient conditions for this problem are (see Exercise 10.5):

$$\begin{aligned}
\mathcal{H}_s(h,s,\mu) &= -h(t)+\mu(t)h(t)\phi'(s(t)h(t))=0 \\
\mathcal{H}_h(h,s,\mu) &= (1-s(t))+\mu(t)\left(s(t)\phi'(s(t)h(t))-\delta_h\right) \\
&= (r+\nu)\mu(t)-\dot{\mu}(t) \\
\lim_{t\to\infty}\exp(-(r+\nu)t)\mu(t)h(t) &= 0.
\end{aligned} \tag{10.14}$$

To solve for the optimal path of human capital investments, let us adopt the following transformation of variables:

$$x(t)\equiv s(t)h(t).$$

Instead of $s(t)$ (or $\mu(t)$) and $h(t)$, we can look at the dynamics of the optimal path in $x(t)$ and $h(t)$.

The first necessary condition then implies that

$$1=\mu(t)\phi'(x(t)), \tag{10.15}$$

while the second necessary condition can be expressed as

$$\frac{\dot{\mu}(t)}{\mu(t)}=r+\nu+\delta_h-s(t)\phi'(x(t))-\frac{1-s(t)}{\mu(t)}.$$

Substituting for $\mu(t)$ from (10.15), and simplifying,

$$\frac{\dot{\mu}(t)}{\mu(t)}=r+\nu+\delta_h-\phi'(x(t)). \tag{10.16}$$

The steady-state (stationary) solution of this optimal control problem involves $\dot{\mu}(t) = 0$ and $\dot{h}(t) = 0$, and thus implies that

$$x^* = \phi'^{-1}(r + \nu + \delta_h), \tag{10.17}$$

where $\phi'^{-1}(\cdot)$ is the inverse function of $\phi'(\cdot)$ (which exists and is strictly decreasing since $\phi(\cdot)$ is strictly concave). This equation shows that $x^* \equiv s^* h^*$ will be higher when the interest rate is low, when the life expectancy of the individual is high, and when the rate of depreciation of human capital is low.

To determine s^* and h^* separately, we set $\dot{h}(t) = 0$ in the human capital accumulation eq. (10.13), which gives

$$\begin{aligned} h^* &= \frac{\phi(x^*)}{\delta_h} \\ &= \frac{\phi\left(\phi'^{-1}(r + \nu + \delta_h)\right)}{\delta_h}. \end{aligned} \tag{10.18}$$

Since $\phi'^{-1}(\cdot)$ is strictly decreasing and $\phi(\cdot)$ is strictly increasing, this equation implies that the steady-state solution for the human capital stock is uniquely determined and is decreasing in r, ν and δ_h.

More interesting than the stationary (steady-state) solution to the optimization problem is the time path of human capital investments in this model. Differentiate (10.15) with respect to time to obtain

$$\frac{\dot{\mu}(t)}{\mu(t)} = \varepsilon_{\phi'}(x) \frac{\dot{x}(t)}{x(t)},$$

where

$$\varepsilon_{\phi'}(x) = -\frac{x\phi''(x)}{\phi'(x)} > 0$$

is the elasticity of the function $\phi'(\cdot)$ and is positive since $\phi'(\cdot)$ is strictly decreasing (thus $\phi''(\cdot) < 0$). Combining this equation with (10.16),

$$\frac{\dot{x}(t)}{x(t)} = \frac{1}{\varepsilon_{\phi'}(x(t))}\left(r + \nu + \delta_h - \phi'(x(t))\right). \tag{10.19}$$

Figure 10.1 plots (10.13) and (10.19) in the h-x space. The upward-sloping curve corresponds to the locus for $\dot{h}(t) = 0$, while (10.19) can only be zero at x^*, thus the locus for $\dot{x}(t) = 0$ corresponds to the horizontal line in the figure. The arrows of motion are also plotted in this phase diagram and make it clear that the steady-state solution (h^*, x^*) is globally saddle-path stable, with the stable arm coinciding with the horizontal line for $\dot{x}(t) = 0$. Starting with $h(0) \in (0, h^*)$, $s(0)$ jumps to the level necessary to ensure $s(0) h(0) = x^*$. From then on, $h(t)$ increases and $s(t)$ decreases so as to keep $s(t) h(t) = x^*$. Therefore, the pattern of human capital investments implied by the Ben-Porath model is one of high investment at the beginning of an individual's life followed by lower investments later on.

In our simplified version of the Ben-Porath model this all happens smoothly. In the original Ben-Porath model, which involves the use of other inputs in the production of human capital and finite horizons, the constraint for $s(t) \leq 1$ typically binds early on in the life of the

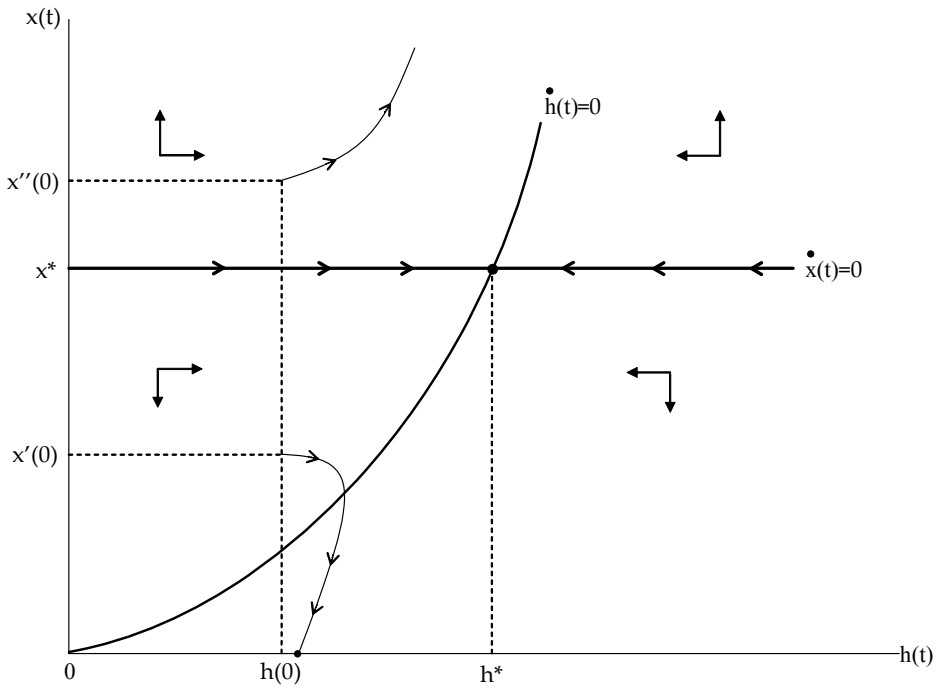

Figure 10.1. Steady state and equilibrium dynamics in the simplified Ben Porath model.

individual, and the interval during which $s(t) = 1$ can be interpreted as full-time schooling. After full-time schooling, the individual starts working ($s(t) < 1$). But even on-the-job, the individual continues to accumulate human capital ($s(t) > 0$), which can be interpreted as spending time in training programs or allocating some of his time on the job to learning rather than production. Moreover, because the horizon is finite, if the Inada conditions were relaxed, the individual could prefer to stop investing in human capital at some point. As a result, the time path of human capital generated by the standard Ben-Porath model may be hump-shaped, with a possibly declining portion at the end (see Exercise 10.7). Instead, the path of human capital (and the earning potential of the individual) in the current model is always increasing as shown in Figure 10.2.

The importance of the Ben-Porath model is twofold. First, it emphasizes that schooling is not the only way in which individuals can invest in human capital and there is a continuity between schooling investments and other investments in human capital. Second, it suggests that in societies where schooling investments are high we may also expect higher levels of on-the-job investments in human capital. Thus there may be systematic mismeasurement of the amount or the quality of human capital across societies.

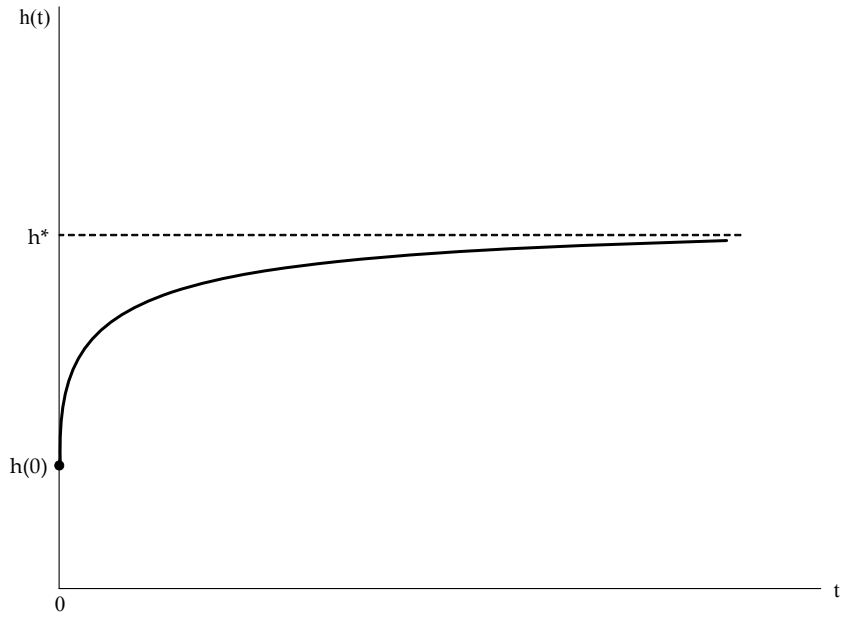

Figure 10.2. Time path of human capital investments in the simplified Ben Porath model.

10.4. Neoclassical Growth with Physical and Human Capital

Our next task is to incorporate human capital investments into the baseline neoclassical growth model. This is useful both to investigate the interactions between physical and human capital, and also to generate a better sense of the impact of differential human capital investments on economic growth. Physical-human capital interactions could potentially be important, since a variety of evidence suggests that physical capital and human capital (capital and skills) are complementary, meaning that greater capital increases the productivity of high human capital workers more than that of low skill workers. This may play an important role in economic growth, for example, by inducing a "virtuous cycle" of investments in physical and human capital. These types of virtue cycles will be discussed in greater detail in Chapter 21. It is instructive to see to what extent these types of complementarities manifest themselves in the neoclassical growth model. The potential for complementarities also raises the issue of "imbalances". If physical and human capital are complementary, the society will achieve the highest productivity when there is a balance between these two different types of capital. However, whether the decentralized equilibrium will ensure such a balance is a question that needs to be investigated.

The impact of human capital on economic growth (and on cross-country income differences) has already been discussed in Chapter 3, in the context of an augmented Solow model,

where the economy was assumed to accumulate physical and human capital with two exogenously given constant saving rates. In many ways, that model was less satisfactory than the baseline Solow growth model, since not only was the aggregate saving rate assumed exogenous, but the relative saving rates in human and physical capital were also taken as given. The neoclassical growth model with physical and human capital investments will enable us to investigate the same set of issues from a different perspective.

Consider the following continuous-time economy admitting a representative household with preferences

$$\int_0^\infty \exp(-\rho t)\, u(c(t))\, dt, \tag{10.20}$$

where the instantaneous utility function $u(\cdot)$ satisfies Assumption 3 (from Chapter 8) and $\rho > 0$. I ignore technological progress and population growth to simplify the discussion. Labor is again supplied inelastically.

Let us follow the specification in Chapter 3 and assume that the aggregate production possibilities frontier of the economy is represented by the aggregate production function:

$$Y(t) = F(K(t), H(t), L(t)),$$

where $K(t)$ is the stock of physical capital, $L(t)$ is total employment, and $H(t)$ represents human capital. Since there is no population growth and labor is supplied inelastically, $L(t) = L$ for all t. This production function is assumed to satisfy Assumptions $1'$ and $2'$ in Chapter 3 (recall that these assumptions generalize Assumptions 1 and 2 to this aggregate production function with three inputs). As already discussed in that chapter, a production function in which "raw" labor and human capital are separate factors of production may be less natural than one in which human capital increases the effective units of labor of workers (as in the analysis of the previous two sections). Nevertheless, this production function allows a simple analysis of neoclassical growth with physical and human capital. As usual, it is more convenient to express all objects in per capita units, thus we write

$$\begin{aligned} y(t) &\equiv \frac{Y(t)}{L} \\ &= f(k(t), h(t)), \end{aligned}$$

where

$$k(t) \equiv \frac{K(t)}{L} \text{ and } h(t) \equiv \frac{H(t)}{L}$$

are the physical and human capital levels per capita. In view of Assumptions $1'$ and $2'$, $f(k,h)$ is strictly increasing, differentiable and jointly strictly concave in both of its arguments. Let us denote its derivatives by f_k, f_h, f_{kh}, and so on Throughout, I assume that physical and human capital are complementary, that is, $f_{kh}(k,h) > 0$ for all $k, h > 0$.

Physical and human capital per capita evolve according to the following two differential equations

$$\dot{k}(t) = i_k(t) - \delta_k k(t), \tag{10.21}$$

and

$$\dot{h}(t) = i_h(t) - \delta_h h(t) \tag{10.22}$$

where $i_k(t)$ and $i_h(t)$ are the investment levels in physical and human capital, while δ_k and δ_h are the depreciation rates of these two capital stocks. The resource constraint for the economy, expressed in per capita terms, is

$$c(t) + i_k(t) + i_h(t) \leq f(k(t), h(t)) \text{ for all } t. \tag{10.23}$$

Since the environment described here is very similar to the standard neoclassical growth model, equilibrium and optimal growth will coincide. For this reason, I focus on the optimal growth problem (the competitive equilibrium is discussed in Exercise 10.13). The optimal growth problem involves the maximization of (10.20) subject to (10.21), (10.22), and (10.23). The solution to this maximization problem can again be characterized by setting up the current-value Hamiltonian. To simplify the analysis, observe first that since $u(c)$ is strictly increasing, (10.23) will always hold as equality. Then, substitute for $c(t)$ using this constraint and write the current-value Hamiltonian as

$$\begin{aligned} \mathcal{H}(k(t), h(t), i_k(t), i_h(t), \mu_k(t), \mu_h(t)) &= u(f(k(t), h(t)) - i_h(t) - i_k(t)) \\ &\quad + \mu_h(t)(i_h(t) - \delta_h h(t)) + \mu_k(t)(i_k(t) - \delta_k k(t)), \end{aligned} \tag{10.24}$$

where now there are two control variables, $i_k(t)$ and $i_h(t)$ and two state variables, $k(t)$ and $h(t)$, as well as two costate variables, $\mu_k(t)$ and $\mu_h(t)$, corresponding to the two constraints, (10.21) and (10.22). The necessary conditions for an optimal solution are

$$\begin{aligned} \mathcal{H}_{i_k}(k(t), h(t), i_k(t), i_h(t), \mu_k(t), \mu_h(t)) &= -u'(c(t)) + \mu_k(t) = 0 \\ \mathcal{H}_{i_h}(k(t), h(t), i_k(t), i_h(t), \mu_k(t), \mu_h(t)) &= -u'(c(t)) + \mu_h(t) = 0 \\ \mathcal{H}_k(k(t), h(t), i_k(t), i_h(t), \mu_k(t), \mu_h(t)) &= f_k(k(t), h(t)) u'(c(t)) - \mu_k(t)\delta_k \\ &= \rho\mu_k(t) - \dot{\mu}_k(t) \\ \mathcal{H}_h(k(t), h(t), i_k(t), i_h(t), \mu_k(t), \mu_h(t)) &= f_h(k(t), h(t)) u'(c(t)) - \mu_h(t)\delta_h \\ &= \rho\mu_h(t) - \dot{\mu}_h(t) \\ \lim_{t\to\infty} \exp(-\rho t)\mu_k(t)k(t) &= 0 \\ \lim_{t\to\infty} \exp(-\rho t)\mu_h(t)h(t) &= 0. \end{aligned}$$

There are two necessary transversality conditions since there are two state variables (and two costate variables). Moreover, it can be shown that $\mathcal{H}(k(t), h(t), i_k(t), i_h(t), \mu_k(t), \mu_h(t))$ is concave given the costate variables $\mu_k(t)$ and $\mu_h(t)$, so that a solution to the necessary conditions indeed gives an optimal path (see Exercise 10.10).

The first two necessary conditions immediately imply that

$$\mu_k(t) = \mu_h(t) = \mu(t).$$

Combining this with the next two conditions gives

$$f_k(k(t), h(t)) - f_h(k(t), h(t)) = \delta_k - \delta_h, \tag{10.25}$$

which (together with $f_{kh} > 0$) implies that there is a one-to-one relationship between physical and human capital, of the form

$$h = \xi(k),$$

where $\xi(\cdot)$ is uniquely defined, strictly increasing and differentiable (with derivative denoted by $\xi'(\cdot)$, see Exercise 10.11).

This observation makes it clear that the model can be reduced to the neoclassical growth model and has exactly the same dynamics as the neoclassical growth model, and thus establishes the following proposition:

Proposition 10.1. *In the neoclassical growth model with physical and human capital investments described above, the optimal path of physical capital and consumption are given as in the one-sector neoclassical growth model, and satisfy the following two differential equations*

$$\frac{\dot{c}(t)}{c(t)} = \frac{1}{\varepsilon_u(c(t))}\left[f_k(k(t), \xi(k(t))) - \delta_k - \rho\right],$$

$$\dot{k}(t) = \frac{1}{1+\xi'(k)}\left[f(k(t), \xi(k(t))) - \delta_h \xi(k(t)) - \delta_k k(t) - c(t)\right],$$

where $\varepsilon_u(c(t)) = -u''(c(t))\, c(t)/u'(c(t))$, *together with the transversality condition* $\lim_{t\to\infty}\left[k(t)\exp\left(-\int_0^t f_k(k(s), \xi(k(s)))\, ds\right)\right] = 0$, *while the level of human capital at time* t *is given by* $h(t) = \xi(k(t))$.

Proof. See Exercise 10.12. □

What is perhaps more surprising, at first, is that eq. (10.25) implies that human and physical capital are always in "balance". Initially, one may have conjectured that an economy that starts with a high stock of physical capital relative to human capital will have a relatively high physical to human capital ratio for an extended period of time. However, Proposition 10.1 and in particular, eq. (10.25) show that this is not the case. The reason for this is that there is no non-negativity constraints on the investment levels. If the economy starts with a high level of physical capital and low level of human capital, at the first instant it will experience a very high level of $i_h(0)$, compensated with a very negative $i_k(0)$, so that at the next instant the physical to human capital ratio will have been brought back to balance. After this, the dynamics of the economy will be identical to those of the baseline neoclassical growth model. Therefore, issues of imbalance will not arise in this version of the neoclassical growth model. This result is an artifact of the fact that there are no non-negativity constraints on physical and human capital investments. The situation is somewhat different when there are such non-negativity or "irreversibility" constraints, that is, if $i_k(t) \geq 0$ and $i_h(t) \geq 0$ for all t. In this case, initial imbalances will persist for a while. In particular, it can be shown that starting with a ratio of physical to human capital stock $(k(0)/h(0))$ that does not satisfy (10.25), the optimal path will involve investment only in one of the two stocks until balance is reached (see Exercise 10.15). Therefore, with irreversibility constraints, some amount of imbalance can arise, but the economy quickly moves towards correcting this imbalance.

Another potential application of the neoclassical growth model with physical and human capital is in the analysis of the impact of policy distortions. Recall the discussion in Section 8.10 in Chapter 8, and suppose that the resource constraint of the economy is modified to

$$c(t) + (1+\tau)(i_k(t) + i_h(t)) \leq f(k(t), h(t)),$$

where $\tau \geq 0$ is a tax affecting both types of investments. Using an analysis parallel to that in Section 8.10, we can characterize the steady-state income ratio of two countries with different policies represented by τ and τ'. In particular, let us suppose that the aggregate production function takes the Cobb-Douglas form

$$\begin{aligned} Y &= F(K, H, L) \\ &= K^{\alpha_k} H^{\alpha_h} L^{1-\alpha_k-\alpha_h}. \end{aligned}$$

In this case, the ratio of income in the two economies with taxes/distortions of τ and τ' is given by (see Exercise 10.16):

$$\frac{Y(\tau)}{Y(\tau')} = \left(\frac{1+\tau'}{1+\tau}\right)^{\frac{\alpha_k+\alpha_h}{1-\alpha_k-\alpha_h}}. \tag{10.26}$$

If we again take α_k to be approximately $1/3$, then the ability of this modified model to account for income differences using tax distortions increases because of the responsiveness of human capital accumulation to these distortions. For example, with $\alpha_k = \alpha_h = 1/3$ and eightfold distortion differences, we would have

$$\frac{Y(\tau)}{Y(\tau')} \approx 8^2 \approx 64,$$

which is a huge difference in economic performance across countries.

Therefore, incorporating human capital into the neoclassical growth model provides one potential way of generating larger income per capita differences. Nevertheless, this result has to be interpreted with caution. First, the large impact of distortions on income per capita here is driven by a very elastic response of human capital accumulation. It is not clear whether human capital investments will indeed respond so much to tax distortions. For instance, if these distortions correspond to differences in corporate taxes or corruption, we would expect them to affect corporations rather than individual human capital decisions. This is of course not to deny that in societies where policies discourage capital accumulation, there are also barriers to schooling and other types of human capital investments. Nevertheless, the impact of these on physical and human capital investments may be quite different. Second, and more important, the large implied elasticity of output to distortions when both physical and human capital are endogenous has an obvious similarity to the Mankiw-Romer-Weil's approach to explaining cross-country differences in terms of physical and human capital stocks. As discussed in Chapter 3, while this is a logical possibility, existing evidence does not support the notion that human capital differences across countries can have such a large impact on income differences. This conclusion equally sheds doubt on the importance of the large contribution of human capital differences induced by policy differences in the current

model. Nevertheless, the conclusions in Chapter 3 were subject to two caveats, which could potentially increase the role of human capital; large human capital externalities and significant differences in the quality of schooling across countries. These issues will be discussed below.

10.5. Capital-Skill Complementarity in an Overlapping Generations Model

The analysis in the previous section suggests that the neoclassical growth model with physical and human capital does not generate significant imbalances between these two different types of capital (unless we impose irreversibilities, in which case it can do so along the transition path). I next investigate possibility of capital-skill imbalances in a simple overlapping generations model with impure altruism, similar to the models introduced in Section 9.6 of the previous chapter. I will see that this class of models also generates only limited capital-skill imbalances. Nevertheless, it provides a simple framework in which labor market frictions can be introduced, and capital-skill imbalances become much more important in the presence of such frictions. We will also use the model in this section to go back to the more natural production function, which features capital and effective units of labor (with human capital-augmenting the effective units of labor), as opposed to the production function used in the previous section with human capital as a third separate factor of production.

The economy is in discrete time and consists of a continuum 1 of dynasties. Each individual lives for two periods, childhood and adulthood. Individual i of generation t works during adulthood at time t, earns labor income equal to $w(t) h_i(t)$, where $w(t)$ is the wage rate per unit of human capital and $h_i(t)$ is the individual's human capital. The individual also earns capital income equal to $R(t) b_i(t-1)$, where $R(t)$ is the gross rate of return on capital and $b_i(t-1)$ is his asset holdings, inherited as bequest from his parent. The human capital of the individual is determined at the beginning of his adulthood by an effort decision. Labor is supplied to the market after this effort decision. At the end of adulthood, after labor and capital incomes are received, the individual decides his consumption and the level of bequest to his offspring.

Preferences of individual i (or of dynasty i) of generation t are given by

$$\eta^{-\eta}(1-\eta)^{-(1-\eta)} c_i(t)^{\eta} b_i(t)^{1-\eta} - \gamma(e_i(t)),$$

where $\eta \in (0,1)$, $c_i(t)$ is own consumption, $b_i(t)$ is the bequest to the offspring, $e_i(t)$ is effort expended for human capital acquisition, and $\gamma(\cdot)$ is a strictly increasing, differentiable and strictly convex cost of effort function. The term $\eta^{-\eta}(1-\eta)^{-(1-\eta)}$ is included as a normalizing factor to simplify the algebra.

The human capital of individual i is given by

$$h_i(t) = a e_i(t), \tag{10.27}$$

where a corresponds to "ability" and increases the effectiveness of effort in generating human capital for the individual. Substituting for $e_i(t)$ in the above expression, the preferences of

individual i of generation t can be written as

$$\eta^{-\eta}(1-\eta)^{-(1-\eta)} c_i(t)^{\eta} b_i(t)^{1-\eta} - \gamma\left(\frac{h_i(t)}{a}\right). \tag{10.28}$$

The budget constraint of the individual is

$$c_i(t) + b_i(t) \leq m_i(t) = w(t) h_i(t) + R(t) b_i(t-1), \tag{10.29}$$

which defines $m_i(t)$ as the current income of individual i at time t consisting of labor earnings, $w(t) h_i(t)$, and asset income, $R(t) b_i(t-1)$ (I use m rather than y, since y will have a different meaning below).

The production side of the economy is given by an aggregate production function

$$Y(t) = F(K(t), H(t)),$$

that satisfies Assumptions 1 and 2, where $H(t)$ is "effective units of labor" or alternatively the total stock of human capital given by,

$$H(t) = \int_0^1 h_i(t)\, di,$$

while $K(t)$, the stock of physical capital, is given by

$$K(t) = \int_0^1 b_i(t-1)\, di.$$

Note also that this specification ensures that capital and skill (K and H) are complements. This is because a production function with two factors and constant returns to scale necessarily implies that the two factors are complements (see Exercise 10.8), that is,

$$\frac{\partial^2 F(K,H)}{\partial K \partial H} \geq 0. \tag{10.30}$$

Furthermore, let us again simplify the notation by assuming capital depreciates fully after use, that is, $\delta = 1$ (see Exercise 10.9).

Since the amount of human capital per worker is an endogenous variable in this economy, it is more useful to define a normalized production function expressing output per unit of human capital rather than the usual per capita production function. In particular, let $\kappa \equiv K/H$ be the capital to human capital ratio (or the "effective capital-labor ratio"), and

$$\begin{aligned} y(t) &\equiv \frac{Y(t)}{H(t)} \\ &= F\left(\frac{K(t)}{H(t)}, 1\right) \\ &= f(\kappa(t)), \end{aligned}$$

where the second line uses the linear homogeneity of $F(\cdot,\cdot)$, while the last line uses the definition of κ. Here we use κ instead of the more usual k, in order to preserve the notation k for capital per worker in the next section. From the definition of κ, the law of motion of

effective capital-labor ratios can be written as

$$\kappa(t) \equiv \frac{K(t)}{H(t)} = \frac{\int_0^1 b_i(t-1)\,di}{\int_0^1 h_i(t)\,di}. \tag{10.31}$$

Factor prices are then given by the usual competitive pricing formulae:

$$R(t) = f'(\kappa(t)) \text{ and } w(t) = f(\kappa(t)) - \kappa(t) f'(\kappa(t)), \tag{10.32}$$

with the only noteworthy feature that $w(t)$ is now wage per unit of human capital, in a way consistent with (10.29).

An equilibrium in this overlapping generations economy is a process of bequest and consumption levels for each individual, $\left\{[h_i(t)]_{i\in[0,1]}, [c_i(t)]_{i\in[0,1]}, [b_i(t)]_{i\in[0,1]}\right\}_{t=0}^{\infty}$, that solve (10.28) subject to (10.29) a sequence of effective capital-labor ratios, $\{\kappa(t)\}_{t=0}^{\infty}$, given by (10.31) with some initial distribution of bequests $[b_i(0)]_{i\in[0,1]}$, and sequences of factor prices, $\{w(t), R(t)\}_{t=0}^{\infty}$, that satisfy (10.32).

The characterization of an equilibrium is simplified by the fact that the solution to the maximization problem of (10.28) subject to (10.29) involves

$$c_i(t) = \eta m_i(t) \text{ and } b_i(t) = (1-\eta) m_i(t), \tag{10.33}$$

and substituting these into (10.28), we obtain the indirect utility function (see Exercise 10.17):

$$m_i(t) - \gamma\left(\frac{h_i(t)}{a}\right), \tag{10.34}$$

which the individual maximizes by choosing $h_i(t)$ and recognizing that $m_i(t) = w(t) h_i(t) + R(t) b_i(t-1)$. The first-order condition of this maximization gives the human capital investment of individual i at time t as

$$aw(t) = \gamma'\left(\frac{h_i(t)}{a}\right). \tag{10.35}$$

Equivalently, inverting this relationship, defining $\gamma'^{-1}(\cdot)$ as the inverse function of $\gamma'(\cdot)$ (which is strictly increasing) and using (10.32), we obtain

$$h_i(t) = h(t) \equiv a\gamma'^{-1}\left[a\left(f(\kappa(t)) - \kappa(t) f'(\kappa(t))\right)\right]. \tag{10.36}$$

An important implication of this equation is that the human capital investment of each individual is identical, and only depends on the effective of capital-labor ratio in the economy. This is a consequence of the specific utility function in (10.28), which ensures that there are no income effects in human capital decisions so that all agents choose the same "income-maximizing" level of human capital (as in Theorem 10.1).

Next, note that since bequest decisions are linear as shown (10.33),

$$\begin{aligned} K(t+1) &= \int_0^1 b_i(t)\,di \\ &= (1-\eta)\int_0^1 m_i(t)\,di \\ &= (1-\eta) f(\kappa(t)) h(t), \end{aligned}$$

where the last line uses the fact that, since all individuals choose the same human capital level given by (10.36), $H(t) = h(t)$, and thus $Y(t) = f(\kappa(t)) h(t)$.

Now combining this with (10.31),

$$\kappa(t+1) = \frac{(1-\eta) f(\kappa(t)) h(t)}{h(t+1)}.$$

Using (10.36), this becomes

$$\begin{aligned} & \kappa(t+1) \gamma'^{-1} \left[a \left(f(\kappa(t+1)) - \kappa(t+1) f'(\kappa(t+1)) \right) \right] \\ = \; & (1-\eta) f(\kappa(t)) \gamma'^{-1} \left[a f(\kappa(t)) - \kappa(t) f'(\kappa(t)) \right]. \end{aligned} \tag{10.37}$$

A steady state, as usual, involves a constant effective capital-labor ratio—$\kappa(t) = \kappa^*$ for all t. Substituting this into (10.37) yields

$$\kappa^* = (1-\eta) f(\kappa^*), \tag{10.38}$$

which defines the unique positive steady-state effective capital-labor ratio, κ^* (since $f(\cdot)$ is strictly concave).

PROPOSITION 10.2. *In the overlapping generations economy with physical and human capital described above, there exists a unique steady state with positive activity, and the physical to human capital ratio is κ^* as given by (10.38).*

This steady-state equilibrium is also typically stable, but some additional conditions need to be imposed on $f(\cdot)$ and $\gamma(\cdot)$ to ensure this (see Exercise 10.18).

An interesting implication of this equilibrium is that, the capital-skill (k-h) complementarity in the production function $F(\cdot,\cdot)$ implies that a certain target level of physical to human capital ratio, κ^*, has to be reached in equilibrium. In other words, physical capital should not be too abundant relative to human capital, and neither should human capital be excessive relative to physical capital. Consequently, this model does not allow equilibrium "imbalances" between physical and human capital either.

A possible and arguably attractive way of introducing such imbalances is to depart from perfectly competitive labor markets. This also turns out to be useful to illustrate how the role of human capital can be quite different in models with imperfect labor markets.

10.6. Physical and Human Capital with Imperfect Labor Markets

In this section, I analyze the implications of labor market frictions that lead to factor prices different from the ones used so far (in particular, in terms of the model of the last section, deviating from the competitive pricing formulea (10.32)). The literature on labor market imperfections is vast and my purpose here is not to provide an overview. For this reason, I adopt the simplest representation. In particular, imagine that the economy is identical to that described in the previous section, except that there is a measure 1 of firms as well as a measure 1 of individuals at any point in time, and each firm can only hire one worker. The production function of each firm is still given by

$$y_j(t) = F(k_j(t), h_i(t)),$$

where $y_j(t)$ refers to the output of firm j, $k_j(t)$ is its capital stock (equivalently capital per worker, since the firm is hiring only one worker), and $h_i(t)$ is the human capital of worker i that the firm has matched with. This production function again satisfies Assumptions 1 and 2. The main departure from the models analyzed so far is the structure for the labor market, which is summarized next.

(1) Firms choose their physical capital level irreversibly (incurring the cost $R(t)k_j(t)$, where $R(t)$ is the market rate of return on capital), and simultaneously workers choose their human capital level irreversibly.

(2) After workers complete their human capital investments, they are *randomly matched* with firms. Random matching here implies that high human capital workers are *not* more likely to be matched with high physical capital firms.

(3) After matching, each worker-firm pair bargains over the division of output between themselves. We assume that they simply divide the output according to some pre-specified rule, and the worker receives total earnings of

$$W_j(k_j(t), h_i(t)) = \lambda F(k_j(t), h_i(t)),$$

for some $\lambda \in (0,1)$.

This is admittedly a very simple and reduced-form specification. Nevertheless, it will be sufficient to emphasize the main economic issues. A more detailed game-theoretic justification for a closely related environment is provided in Acemoglu (1996).

Let us also introduce heterogeneity in the cost of human capital acquisition by modifying (10.27) to

$$h_i(t) = a_i e_i(t),$$

where a_i differs across dynasties (individuals). A high-value of a_i naturally corresponds to an individual who can more effectively accumulate human capital.

An equilibrium is defined similarly except that factor prices are no longer determined by (10.32). Let us start the analysis with the physical capital choices of firms. At the time each firm chooses its physical capital it is unsure about the human capital of the worker he will be facing. Therefore, the expected return of firm j can be written as

$$(1-\lambda)\int_0^1 F(k_j(t), h_i(t))\,di - R(t)k_j(t). \tag{10.39}$$

This expression takes into account that the firm will receive a fraction $1-\lambda$ of the output produced jointly by itself and by the worker that it is matched with. The integration takes care of the fact that the firm does not know which worker it will be matched with and thus we are taking the expectation of $F(k_j(t), h_i(t))$ over all possible human capital levels available in the market (given by $[h_i(t)]_{i\in[0,1]}$). The last term is the cost of making irreversible capital investment at the market price $R(t)$. This investment is made before the firm knows which worker it will be matched with. The important observation is that the objective function in (10.39) is strict concave in $k_j(t)$ given the strict concavity of $F(\cdot,\cdot)$ from Assumption 1.

Therefore, each firm will choose the same level of physical capital, $\hat{k}(t)$, such that

$$(1-\lambda)\int_0^1 \frac{\partial F\left(\hat{k}(t), h_i(t)\right)}{\partial k(t)} di = R(t).$$

Now given this (expected) capital investment by firms, and following (10.34) from the previous section, each worker's objective function can be written as:

$$\lambda F\left(\hat{k}(t), h_i(t)\right) + R(t) b_i(t-1) - \gamma\left(\frac{h_i(t)}{a_i}\right),$$

which substitutes for the income $m_i(t)$ of the worker in terms of his wage earnings and capital income, and introduced the heterogeneity in human capital decisions. This implies the following choice of human capital investment by a worker i:

$$\lambda a_i \frac{\partial F\left(\hat{k}(t), h_i(t)\right)}{\partial h_i(t)} = \gamma'\left(\frac{h_i(t)}{a_i}\right).$$

This equation yields a unique equilibrium human capital investment $\hat{h}_i\left(\hat{k}(t)\right)$ for each i. This human capital investment directly depends on the capital choices of all the firms, $\hat{k}(t)$ (since this affects the marginal product of human capital) and also depends implicitly on a_i. Moreover, given (10.30), $\hat{h}_i\left(\hat{k}(t)\right)$ is strictly increasing in $\hat{k}(t)$. Also, since $\gamma(\cdot)$ is strictly convex, $\hat{h}_i\left(\hat{k}(t)\right)$ is a strictly concave function of $\hat{k}(t)$. Substituting this into the first-order condition of firms,

$$(1-\lambda)\int_0^1 \frac{\partial F\left(\hat{k}(t), \hat{h}_i\left(\hat{k}(t)\right)\right)}{\partial k(t)} di = R(t).$$

Finally, to satisfy market clearing in the capital market, the rate of return to capital, $R(t)$, has to adjust, such that

$$\hat{k}(t) = \int_0^1 b_i(t-1)\, di,$$

which follows from the facts that all firms choose the same level of capital investment and that the measure of firms is normalized to 1. This equation implies that in the closed economy version of the current model, capital per firm is fixed by bequest decisions from the previous period. The main economic forces we would like to emphasize here are seen more clearly when physical capital is not predetermined. For this reason, let us imagine that the economy in question is small and open, so that $R(t) = R^*$ is pinned down by international financial markets (the closed economy version is further discussed in Exercise 10.19). Under this assumption, the equilibrium level of capital per firm is determined by

$$(1-\lambda)\int_0^1 \frac{\partial F\left(\hat{k}, \hat{h}_i\left(\hat{k}\right)\right)}{\partial k} di = R^*. \tag{10.40}$$

Proposition 10.3. *In the open economy version of the model described here, there exists a unique positive level of capital per worker $\hat{k}$ given by (10.40) such that the equilibrium*

capital per worker is always equal to $\hat{k}$*. Given* $\hat{k}$*, the human capital investment of worker* i *is uniquely determined by* $\hat{h}_i\left(\hat{k}\right)$ *such that*

$$\lambda a_i \frac{\partial F\left(\hat{k}, \hat{h}_i\left(\hat{k}\right)\right)}{\partial h} = \gamma'\left(\frac{\hat{h}_i\left(\hat{k}\right)}{a_i}\right). \tag{10.41}$$

Here $\hat{h}_i\left(\hat{k}\right)$ *is increasing in* $\hat{k}$*, and a decline in* R^* *increases* $\hat{k}$ *and* $\hat{h}_i$ *for all* $i \in [0,1]$*.*

In addition to this equilibrium, there also exists a no-activity equilibrium in which $\hat{k} = 0$ *and* $\hat{h}_i = 0$ *for all* $i \in [0,1]$*.*

PROOF. Since $F(k,h)$ exhibits constant returns to scale and $\hat{h}_i\left(\hat{k}\right)$ is a concave function of $\hat{k}$ for each i, $\int_0^1 \left(\partial F\left(\hat{k}, \hat{h}_i\left(\hat{k}\right)\right)/\partial k\right) di$ is decreasing in $\hat{k}$ for a distribution of $[a_i]_{i\in[0,1]}$. Thus $\hat{k}$ is uniquely determined. Given $\hat{k}$, (10.41) determines $\hat{h}_i\left(\hat{k}\right)$ uniquely. Applying the Implicit Function Theorem to (10.41) implies that $\hat{h}_i\left(\hat{k}\right)$ is increasing in $\hat{k}$. Finally, (10.40) implies that a lower R^* increases $\hat{k}$, and from the previous observation $\hat{h}_i$ for all $i \in [0,1]$ increase as well.

The no-activity equilibrium follows, since when all firms choose $\hat{k} = 0$, output is equal to zero and it is best response for workers to choose $\hat{h}_i = 0$, and when $\hat{h}_i = 0$ for all $i \in [0,1]$, $\hat{k} = 0$ is the best response for all firms. □

Observe that there is underinvestment both in human capital and physical capital (this refers to the positive activity equilibrium; clearly, there is even a more severe underinvestment in the no-activity equilibrium). Consider a social planner wishing to maximize output (or one who could transfer resources across individuals in a lump-sum fashion). Suppose that the social planner is restricted by the same random matching technology, so that she cannot allocate workers to firms as she wishes. A similar analysis to above implies that the social planner would also like each firm to choose an identical level of capital per firm, say $\bar{k}$. However, this level of capital per firm will be different than in the competitive equilibrium and she will also choose a different relationship between human capital and physical capital investments. In particular, given $\bar{k}$, she would make human capital decisions to satisfy

$$a_i \frac{\partial F\left(\bar{k}, \bar{h}_i\left(\bar{k}\right)\right)}{\partial h} = \gamma'\left(\frac{\bar{h}_i\left(\bar{k}\right)}{a_i}\right),$$

which is similar to (10.41), except that λ is absent from the left-hand side. This is because each worker takes into account only his share of output, λ, when undertaking his human capital investment decisions, while the social planner considers the entire output. Consequently, as long as $\lambda < 1$,

$$\bar{h}_i(k) > \hat{h}_i(k) \text{ for all } k > 0.$$

Similarly, the social planner would also choose a higher level of capital investment for each firm, in particular, to satisfy the equation

$$\int_0^1 \frac{\partial F\left(\bar{k}, \bar{h}_i\left(\bar{k}\right)\right)}{\partial k} di = R^*,$$

which differs from (10.40) both because now the term $1-\lambda$ is not present on the left-hand side and also because the planner takes into account the differential human capital investment behavior of workers given by $\bar{h}_i\left(\bar{k}\right)$. This discussion establishes the following result:

PROPOSITION 10.4. *In the equilibrium described in Proposition 10.3, there is underinvestment both in physical and human capital.*

More interesting than the underinvestment result is the imbalance in the physical to human capital ratio of the economy, which did not feature in the previous two environments we discussed. The following proposition summarizes this imbalance result in a sharp way:

PROPOSITION 10.5. *Consider the positive activity equilibrium described in Proposition 10.3. Output is equal to 0 if either $\lambda = 0$ or $\lambda = 1$. Moreover, there exists $\lambda^* \in (0,1)$ that maximizes output.*

PROOF. See Exercise 10.20. □

Intuitively, different levels of λ create different types of "imbalances" between physical and human capital. A high level of λ implies that workers have a strong bargaining position, and this encourages their human capital investments. But symmetrically, it discourages the physical capital investments of firms, since they will only receive a small fraction of the output. Therefore, high level of λ (as long as $\lambda < 1$) creates an imbalance with too high a level of human capital relative to physical capital. This imbalance effect becomes more extreme as $\lambda \to 1$. In this limit, workers' investment behavior converges to the first-order condition of the social planner (that is, $\hat{h}_i(k) \to \bar{h}_i(k)$ for all $k > 0$). However, simultaneously, the physical capital investment of each firm, $\hat{k}$, converges to zero, and this implies that $\hat{h}_i(k) \to 0$, and production collapses. The same happens, in reverse, when λ is too low. Now there is too high a level of physical capital relative to human capital. An intermediate value of λ^* achieves a balance, though the equilibrium continues to be inefficient as shown in Proposition 10.5.

Physical-human capital imbalances can also increase the role of human capital in cross-country income differences. In the current model, the proportional impact of a change in human capital on aggregate output (or on labor productivity) is greater than the return to human capital, since the latter is determined not by the marginal product of human capital, but by the bargaining parameter λ. The deviation from competitive factor prices, therefore, decouples the contribution of human capital to productivity from market prices.

At the root of the inefficiencies and of the imbalance effect in this model are *pecuniary externalities*. Pecuniary externalities refer to external effects that work through prices (not through direct technological spillovers). By investing more, workers (and symmetrically firms) increase the return to capital (symmetrically wages), and there is underinvestment because

they do not take these external effects into consideration. Pecuniary external effects are also present in competitive markets (since, for example, supply affects price), but these are typically "second order," because prices are such that they are equal to both the marginal benefit of buyers (marginal product of firms in the case of factors of production) and to the marginal cost of suppliers. The presence of labor market frictions causes a departure from this type of marginal pricing and is the reason why pecuniary externalities are not second order.

Perhaps even more interesting is the fact that pecuniary externalities in this model take the form of *human capital externalities*, meaning that greater human capital investments by a group of workers increase other workers' wages. Notice that in competitive markets (without externalities) this does not happen. For example, in the economy analyzed in the last section, if a group of workers increase their human capital investments, this would depress the physical to human capital ratio in the economy, reducing wages per unit of human capital and thus the earnings of the rest of the workers. Interestingly, the opposite may happen in the presence of labor market imperfections. To illustrate this point, let us suppose that there are two types of workers, a fraction of workers χ with ability a_1 and $1-\chi$ with ability $a_2 < a_1$. Using this specific structure, the first-order condition of firms, (10.40), can be written as

$$(1-\lambda)\left[\chi\frac{\partial F\left(\hat{k},\hat{h}_1\left(\hat{k}\right)\right)}{\partial k}+(1-\chi)\frac{\partial F\left(\hat{k},\hat{h}_2\left(\hat{k}\right)\right)}{\partial k}\right]=R^*, \tag{10.42}$$

while the first-order conditions for human capital investments for the two types of workers take the form

$$\lambda a_j\frac{\partial F\left(\hat{k},\hat{h}_j\left(\hat{k}\right)\right)}{\partial h}=\gamma'\left(\frac{\hat{h}_j\left(\hat{k}\right)}{a_j}\right) \text{ for } j=1,2. \tag{10.43}$$

Clearly, $\hat{h}_1(k) > \hat{h}_2(k)$ since $a_1 > a_2$. Now imagine an increase in χ, which corresponds to an increase in the fraction of high-ability workers in the population. Holding $\hat{h}_1\left(\hat{k}\right)$ and $\hat{h}_2\left(\hat{k}\right)$ constant, (10.42) implies that $\hat{k}$ should increase, since the left-hand side has increased (in view of the fact that $\hat{h}_1\left(\hat{k}\right) > \hat{h}_2\left(\hat{k}\right)$ and $\partial^2 F(k,h)/\partial k\partial h > 0$). Therefore, capital-skill complementarity combined with the pecuniary externalities implies that an improvement in the pool of workers that firms face leads to greater investments by firms.

Intuitively, each firm expects the average worker that it will be matched with to have higher human capital and since physical and human capital are complements, this makes it more profitable for each firm to increase their physical capital investment. Greater investments by firms, in turn, raise $F\left(\hat{k},h\right)$ for each h, in particular for $\hat{h}_2\left(\hat{k}\right)$. Since the earnings of type 2 workers is equal to $\lambda F\left(\hat{k},\hat{h}_2\left(\hat{k}\right)\right)$, their earnings will also increase as a result of the response of firms to the change in the composition of the workforce. This is therefore an example of human capital externalities, since greater human capital investments by one

group of workers have increased the earnings of the remaining workers. In fact, human capital externalities, in this economy, are even stronger, because the increase in $\hat{k}$ also raises $\partial F\left(\hat{k}, \hat{h}_2\left(\hat{k}\right)\right)/\partial h$ and thus encourages further investments by type 2 workers. These feedback effects nonetheless do not lead to divergence or multiple equilibria, since we know from Proposition 10.3 that there exists a unique equilibrium with positive activity. This discussion is summarized in the following result:

Proposition 10.6. *The positive activity equilibrium described in Proposition 10.3 exhibits human capital externalities in the sense that an increase in the human capital investments of a group of workers raises the earnings of the remaining workers.*

10.7. Human Capital Externalities

The previous section illustrated how a natural form of human capital externalities can emerge in the presence of capital-skill complementarities combined with labor market imperfections. This is not the only channel through which human capital externalities may arise. Many economists believe that the human capital stock of the workforce creates a direct non-pecuniary (technological) spillover on the productivity of each worker. In *The Economy of Cities,* Jane Jacobs, for example, argued for the importance of human capital externalities, and suggested that the concentration of economic activity in cities is partly a result of these externalities and also acts as an engine of economic growth because it facilitates the exchange of ideas among workers and entrepreneurs. In the growth literature, a number of well-known papers, including Robert Lucas' (1988) paper and Azariadis and Drazen (1990), suggest that such technological externalities are important and play a major role in the process of economic growth. Human capital externalities are interesting in their own right, since if such external effects are present, the competitive price system may be inefficient (since it will fail to internalize these externalities, particularly if they take place across firm boundaries). Human capital externalities are also important for our understanding of the sources of income differences across countries. The discussion of the contribution of physical and human capital to cross-country income differences in Chapter 3 showed that differences in human capital are unlikely to account for a large fraction of cross-country income differences, unless external effects are important.

At this point, it is therefore useful to briefly review the empirical evidence on the extent of human capital externalities. Early work in the area, in particular, the paper by James Rauch (1993) tried to measure the extent of human capital externalities by estimating quasi-Mincerian wage regressions, with the major difference that average human capital of workers in the local labor market is also included on the right-hand side. More specifically, Rauch estimated models of the following form:

$$\ln W_{j,m} = \mathbf{X}_{\mathbf{j,m}}^{\mathbf{T}}\boldsymbol{\beta} + \gamma_p S_{j,m} + \gamma_e S_m,$$

where $\mathbf{X}_{j,m}$ is a vector of controls, $S_{j,m}$ is the years of schooling of individual j living/working in labor market m, and S_m is the average years of schooling of workers in labor market m.

Without this last term, this equation would be similar to the standard Mincerian wage regressions discussed above, and we would expect an estimate of the *private return* to schooling γ_p between 6 and 10%. When the average years of schooling, S_m, is also included in the regression, its coefficient γ_e measures the *external return* to schooling in the same units. For example, if γ_e is estimated to be of the same magnitude as γ_p, we would conclude that external returns to schooling are as important as private returns (which would correspond to very large externalities).

Rauch estimated significant external returns, with the magnitude of the external returns often exceeding the private returns. External returns of this magnitude would imply that human capital differences could play a much more important role as a proximate source of cross-country differences in income per capita than implied by the computations in Chapter 3. However, Rauch's regressions exploited differences in average schooling levels across cities, which could reflect many factors that also directly affect wages. For example, wages are much higher in New York City than Ames, Iowa, but this is not only the result of the higher average education of New Yorkers. A more convincing estimate of external returns necessitates a source of exogenous variation in average schooling.

Acemoglu and Angrist (2000) exploited differences in average schooling levels across states and cohorts resulting from changes in compulsory schooling and child labor laws. These laws appear to have had a large effect on schooling, especially at the high school margin. Exploiting changes in average schooling in state labor markets driven by these law changes, Acemoglu and Angrist estimate external returns to schooling that are typically around 1 or 2 percent and statistically insignificant (as compared to private returns of about 10%). These results suggest that there are relatively small human capital externalities in local labor markets. This result is confirmed by a study by Duflo (2004) using Indonesian data and by Ciccone and Perri (2006). Moretti (2002) also estimates human capital externalities, and he finds larger effects. This may be because he focuses on college graduation, but also partly reflects the fact that the source of variation that he exploits, changes in age composition and the presence of land-grant colleges, may have other effects on average earnings in area. Overall, the evidence appears to suggest that local human capital externalities are not very large, and calibration exercises as those in Chapter 3 that ignore these externalities are unlikely to lead to significant downward bias in the contribution of human capital to cross-country income differences.

The qualification "local" in the above discussion has to be emphasized, however. The estimates discussed above focus on local externalities originally emphasized by Jacobs. Nevertheless, if a few very talented scientists and engineers, or other very skilled workers, generate ideas that are then used in other parts of the country or even in the world economy, there may exist significant global human capital externalities. Such global external effects would not be captured by the currently available empirical strategies. Whether such global human capital externalities are important is an interesting area for future research.

10.8. The Nelson-Phelps Model of Human Capital

The discussion in this chapter so far has focused on the productivity-enhancing role of human capital. This is arguably the most important role of human capital, emphasized by Becker and Mincer's seminal analyses. However, an alternative perspective on human capital is provided by Richard Nelson and Edmund Phelps in their short and influential paper, Nelson and Phelps (1966), and also by Ted Schultz (1965). According to this perspective, the major role of human capital is not to increase productivity in existing tasks, but to enable workers to cope with change, disruptions and especially new technologies. The Nelson-Phelps view of human capital has played an important role in a variety of different literatures and features in a number of growth models. Here I provide a simple presentation of the main ideas along the lines of Nelson and Phelps' original model and a discussion of how this new dimension of human capital may enrich our view of its role in economic growth and development. This model will also act as a steppingstone towards our study of technology adoption later in the book.

Consider the following continuous-time model to illustrate the basic ideas. Suppose that output in the economy in question is given by

$$Y(t) = A(t) L, \tag{10.44}$$

where L is the constant labor force, supplying its labor inelastically, and $A(t)$ is the technology level of the economy. There is no capital (and thus no capital accumulation decision) and also no labor supply margin. The only variable that changes over time is technology $A(t)$.

Suppose that the world technological frontier is given by $A_F(t)$. This could correspond to the technology in some other country or perhaps to the technological know-how of scientists that has not yet been applied to production processes. Suppose that $A_F(t)$ evolves exogenously according to the differential equation

$$\frac{\dot{A}_F(t)}{A_F(t)} = g_F,$$

with initial condition $A_F(0) > 0$.

Let the human capital of the workforce be denoted by h. Notice that this human capital does not feature in the production function, (10.44). This is an extreme case in which human capital does not play any productivity-enhancing role. Instead, the role of human capital in the current model will be to facilitate the implementation and use of frontier technology in the production process. In particular, the evolution of the technology in use, $A(t)$, is governed by the differential equation

$$\dot{A}(t) = gA(t) + \phi(h) A_F(t),$$

with initial condition $A(0) \in (0, A_F(0))$. The parameter g is strictly less than g_F and measures the growth rate of technology $A(t)$, resulting from learning-by-doing or other sources of productivity growth. But this is only one source of improvements in technology. The other

one comes from the second term, and can be interpreted as improvements in technology because of implementation and adoption of frontier technologies. The extent of this second source of improvement is determined by the average human capital of the workforce, h. This captures the above-mentioned role of human capital in the context of technology adoption and more generally facilitation of adaptation to technological change. In particular, suppose that $\phi(\cdot)$ is increasing, with

$$\phi(0) = 0 \text{ and } \phi(h) = g_F - g > 0 \text{ for all } h \geq \bar{h},$$

where $\bar{h} > 0$. This specification implies that the human capital of the workforce regulates the ability of the economy to cope with new developments embedded in the frontier technologies; if the workforce has no human capital, there will be no adoption or implementation of frontier technologies and $A(t)$ will grow at the rate g. If, in contrast, $h \geq \bar{h}$, there will be very quick adaptation to the frontier technologies.

Since $A_F(t) = \exp(g_F t) A_F(0)$, the differential equation for $A(t)$ can be written as

$$\dot{A}(t) = gA(t) + \phi(h) A_F(0) \exp(g_F t).$$

The solution to this differential equation is (recall again Section B.4 in Appendix Chapter B):

$$A(t) = \left[\left(\frac{A(0)}{g} - \frac{\phi(h) A_F(0)}{g_F - g}\right) \exp(gt) + \frac{\phi(h) A_F(0)}{g_F - g} \exp(g_F t)\right],$$

which shows that the growth rate of $A(t)$ is faster when $\phi(h)$ is higher. Moreover, it can be verified that

$$A(t) \to \frac{\phi(h)}{g_F - g} A_F(t),$$

so that the ratio of the technology in use to the frontier technology is also determined by human capital.

The role of human capital emphasized by Nelson and Phelps is undoubtedly important in a number of situations. For example, a range of empirical evidence shows that more educated farmers are more likely to adopt new technologies and seeds (e.g., Foster and Rosenzweig, 1995). The Nelson and Phelps' conception of human capital has also been emphasized in the growth literature in connection with the empirical evidence already discussed in Chapter 1, which shows that there is a stronger correlation between economic growth and levels of human capital than between economic growth and changes in human capital. A number of authors, for example, Benhabib and Spiegel (1994), suggest that this may be precisely because the most important role of human capital is not to increase the productive capacity with existing tasks, but to facilitate technology adoption. One might then conjecture that if the role of human capital emphasized by Nelson and Phelps is important in practice, human capital could be playing a more major role in economic growth and development than the discussion so far has suggested. While this is an interesting hypothesis, it is not entirely convincing. If the role of human capital in facilitating technology adoption is taking place within the firm's boundaries, then this will be reflected in the marginal product and earnings of more skilled workers. Workers that contribute to faster and more effective technology adoption would be

compensated in line with the increase in the net present value of the firm. Then, the returns to schooling and human capital used in the calculations in Chapter 3 should have already taken into account the contribution of human capital to aggregate output (thus to economic growth). If, on the other hand, human capital facilitates technology adoption not at the level of the firm, but at the level of the labor market, this would be a form of local human capital externalities and it should have shown up in the estimates on local external effects of human capital. It therefore would appear that, unless this particular role of human capital is also external and these external effects work at a global level, the calibration exercises in Chapter 3, which use labor market returns to human capital, should not be seriously underestimating the contribution of human capital to cross-country differences in income per capita.

10.9. Taking Stock

Human capital differences are a major proximate cause of cross-country differences in economic performance. In addition, human capital accumulation may play an important role in the process of economic growth and economic development. These considerations justify a detailed analysis of human capital. This chapter has presented a number of models of human capital investments that have emphasized how human capital investments respond to future rewards and how they evolve over time (with schooling as well as on-the-job training).

Four sets of related but distinct issues arise in connection with the role of human capital in economic growth. First, if some part of the earnings of labor we observe are rewards to accumulated human capital, then the effect of policies (and perhaps technology) on income per capita could be larger, because these would affect not only physical capital accumulation but also human accumulation. The neoclassical economy with physical and human capital studied in Section 10.4 models and quantifies this effect. It also provides a tractable framework in which physical and human capital investments can be studied simultaneously. Nevertheless, any effect of human capital differences resulting from differences in distortions or policies across countries should have shown up in the measurements in Chapter 3. The findings there suggest that human capital differences, though important, can only explain a small fraction of cross-country income differences (unless there is a significant mismeasurement of the impact of human capital on productivity).

The second important issue connected to the role of human capital relates to the measurement of the contribution of education and skills to productivity. A possible source of mismeasurement of these effects is the presence of human capital externalities. There are many compelling reasons why there might exist significant pecuniary or technological human capital externalities. Section 10.6 illustrated how capital-skill complementarities in imperfect labor markets can lead to pecuniary externalities. Nevertheless, existing evidence suggests that the extent of human capital externalities is rather limited—with the important caveat that there might be global externalities that remain unmeasured. A particular channel through which global externalities may arise is R&D and technological progress, which are the topics of the next part of the book. An alternative source of mismeasurement of the contribution

of human capital is differences in human capital quality. There are significant differences in school and teacher quality even within a narrow geographical area, so we may expect much larger differences across countries. In addition, most available empirical approaches measure human capital differences across countries by using differences in formal schooling. And yet, the Ben-Porath model, analyzed in Section 10.3, suggests that human capital continues to be accumulated even after individuals complete their formal schooling. When human capital is highly rewarded, we expect both higher levels of formal schooling and greater levels of on-the-job investments. Consequently, the Ben-Porath model suggests that there might be higher quality of human capital (or greater amount of unmeasured human capital) in economies where the levels of formal schooling are higher. If this is the case, the empirical measurements reported in Chapter 3 may understate the contribution of human capital to productivity. Whether or not this is so is an interesting area for future research.

The third set of novel issues raised by the modeling of human capital is the possibility of an imbalance between physical and human capital. Empirical evidence suggests that physical and human capital are complementary. This implies that productivity will be high when the correct balance is achieved between physical and human capital. Could equilibrium incentives lead to an imbalance, whereby too much or too little physical capital is accumulated relative to human capital? We saw that such imbalances are unlikely or rather short lived in models with competitive labor markets. However, the analysis in Section 10.6 shows that they become a distinct possibility when factor prices do not necessarily reflect marginal products, as in labor markets with frictions. The presence of such imbalances might increase the impact of human capital on aggregate productivity.

The final issue relates to the role of human capital in technological change and technology adoption. Section 10.8 presented the Nelson-Phelps view of human capital, which emphasizes the role of skills in facilitating the adoption and implementation of new technologies. While this perspective is likely to be important in a range of situations, it seems that, in the absence of significant external effects, this particular role of human capital should not lead to a major mismeasurement of the contribution of human capital to aggregate productivity either, especially, in the types of exercises reported in Chapter 3.

This chapter has also contributed to our quest towards understanding the sources of economic growth and cross-country income differences. We have now arrived to a relatively simple and useful framework for understanding both physical and human capital accumulation decisions. Our next task is to develop models for the other major proximate source of economic growth and income differences; technology.

10.10. References and Literature

The concept of human capital is due to Ted Shultz (1965), Gary Becker (1965), and Jacob Mincer (1974). The standard models of human capital, used extensively in labor economics and in other areas economics, have been developed by Becker (1965), Mincer (1974) and Yoram Ben-Porath (1967). These models have been the basis of the first three sections of this

chapter. Recently there has been a renewed interest in the Ben-Porath model among macroeconomists. Recent contributions include Heckman, Lochner and Taber (1998), Manuelli and Seshadri (2005) and Guvenen and Kuruscu (2006). These models make parametric assumptions (Cobb-Douglas functional forms) and try to gauge the quantitative implications of the Ben-Porath model for cross-country income differences and for the evolution on wage inequality, respectively. Manuelli and Seshadri (2005) also emphasize how differences in on-the-job training investments will create systematic differences in unmeasured human capital across countries and argue that once these "quality" differences are taken into account, human capital differences could explain a very large fraction of cross-country income differences. Caselli (2006), on the other hand, argues that quality differences are unlikely to increase the contribution of human capital to aggregate productivity.

There is a large literature on returns to schooling. As noted in the text and also in Chapter 3, this literature typically finds that one more year of schooling increases earnings by about 6 to 10% (see, for example, the survey in Card, 1999).

There is also a large literature on capital-skill complementarity. The idea was first put forward and empirically supported in Griliches (1969). Katz and Autor (1999) summarize more recent evidence on as capital-skill complementarities.

Technological human capital externalities are emphasized in Jacobs (1965), Lucas (1988), Azariadis and Drazen (1990), while pecuniary human capital externalities were first discussed by Marshall (1961), who argued that increasing the geographic concentration of specialized inputs increases productivity since the matching between factor inputs and industries is improved. Models of pecuniary human capital externalities are constructed in Acemoglu (1996, 1997a). The model with capital-skill complementarity and labor market imperfections is based on Acemoglu (1996). In that paper, I provided a more detailed and microfounded model leading to similar results to those presented in Section 10.6 and derived the results on pecuniary externalities and human capital externalities.

The empirical literature on human capital externalities includes Rauch (1993), Acemoglu and Angrist (2000), Duflo (2004), Moretti (2002) and Ciccone and Perri (2006).

The role of human capital in adapting to change and implementing new technologies was first suggested by Schultz (1965) in the context of agricultural technologies (he emphasized the role of ability rather than human capital and stressed the importance of "disequilibrium" situations). Nelson and Phelps (1966) formulated the same ideas and presented a simple model, essentially identical to that presented in Section 10.8 above. Foster and Rosenzweig (1995) provide evidence consistent with this role of human capital. Benhabib and Spiegel (1994) and Aghion and Howitt (1999) also include extensive discussions of the Nelson-Phelps view of human capital. Recent macroeconomic models that feature this role of human capital include Galor and Tsiddon (1997), Greenwood and Yorukoglu (1997), Caselli (1999), Galor and Moav (2001), and Aghion, Howitt and Violante (2004).

10.11. Exercises

Exercise 10.1. Formulate, state and prove the Separation Theorem, Theorem 10.1, in an economy in discrete time.

Exercise 10.2. (1) Consider the environment discussed in Section 10.1. Write the flow budget constraint of the individual as

$$\dot{a}(t) = ra(t) - c(t) + W(t),$$

and suppose that there are credit market imperfections so that $a(t) \geq 0$. Construct an example in which Theorem 10.1 does not apply. Can you generalize this to the case in which the individual can save at the rate r, but can only borrow at the rate $r' > r$?

(2) Now modify the environment so that the instantaneous utility function of the individual is

$$u(c(t), 1 - l(t)),$$

where $l(t)$ denotes total hours of work, labor supply at the market is equal to $l(t) - s(t)$, so that the individual has a non-trivial leisure choice. Construct an example in which Theorem 10.1 does not apply.

Exercise 10.3. Derive eq. (10.9) from (10.8).

Exercise 10.4. Consider the model presented in Section 10.2 and suppose that the discount rate r varies across individuals (for example, because of credit market imperfections). Show that individuals facing a higher r would choose lower levels of schooling. What would happen if you estimate the wage regression similar to (10.12) in a world in which the source of difference in schooling is differences in discount rates across individuals?

Exercise 10.5. Verify that Theorems 7.13 and 7.14 from Chapter 7 can be applied both to the Ben-Porath then lead to (10.14 as necessary and sufficient conditions for an optimal path of human capital investments. [Hint: use a similar argument that in Section 7.7 in Chapter 7].

Exercise 10.6. Consider the following variant of the Ben-Porath model, where the human capital accumulation equation is given by

$$\dot{h}(t) = s(t)\phi(h(t)) - \delta_h h(t),$$

where ϕ is strictly increasing, differentiable and strictly concave, with $s(t) \in [0, 1]$. Assume that individuals are potentially infinitely lived and face a Poisson death rate of $\nu > 0$. Show that the optimal path of human capital investments involves $s(t) = 1$ for some interval $[0, T]$ and then $s(t) = s^*$ for $t \geq T$.

Exercise 10.7. Modify the Ben-Porath model studied in Section 10.3 as follows. First, assume that the horizon is finite. Second, suppose that $\phi'(0) < \infty$. Finally, suppose that

$$\phi'(h(0)) > \delta_h / (1 - \exp(-\delta_h T)),$$

where recall that δ_h is the rate of depreciation of human capital.

(1) Provide the necessary conditions for an interior solution. Highlight how these necessary conditions should be modified to allow for corner solutions where $s(t)$ might take the value of 0 or 1.
(2) Show that under these conditions the optimal path of human capital accumulation will involve an interval $[0, t')$ of full-time schooling with $s(t) = 1$ for all $t \in [0, t')$, where $t' > 0$, followed by another interval of on-the-job investment $s(t) \in (0, 1)$, and finally an interval of no human capital investment, i.e., $s(t) = 0$ for all $t \in (t'', T]$, where $t'' < T$. [Hint: suppose that the first part of the claim is not true, and show that in this case the necessary conditions must hold as equality. Combining the two necessary conditions, derive a first-order linear non-autonomous differential equation for the costate variable $\lambda(t)$, and solve this differential equation with the boundary condition $\lambda(T) = 0$. Then, show that given the implied value for $\lambda(0)$ and the inequality above, the necessary conditions at $t = 0$ cannot be satisfied. Next use the assumption that $\phi' < \infty$ together with the fact that the costate variable $\lambda(t)$ is continuous and must satisfy $\lambda(T) = 0$ to prove that $s(t)$ must be equal to zero for some interval $[T - \xi, T]$. Finally, using these intermediate steps, conclude that $s(t)$ must take intermediate values before this final interval].
(3) How do the earnings of the individual evolve over the life cycle?
(4) How would you test the implications of this model?Second, suppose that $\phi'(0) < \infty$. Finally, suppose that $\lim_{x \to h(0)} \phi'(x) > 0$. Show that under these conditions the optimal path of human capital accumulation will involve an interval of full-time schooling with $s(t) = 1$, followed by another interval of on-the-job investment $s(t) \in (0, 1)$, and finally an interval of no human capital investment, $s(t) = 0$. How do the earnings of the individual evolve over the life cycle?

Exercise 10.8. Prove that as long as $Y(t) = F(K(t), H(t))$ satisfies Assumptions 1 and 2, the inequality in (10.30) holds.

Exercise 10.9. Show that equilibrium dynamics in Section 10.5 remain unchanged if $\delta < 1$.

Exercise 10.10. Prove that the current-value Hamiltonian in (10.24) is jointly concave in $(k(t), h(t), i_k(t), i_h(t))$.

Exercise 10.11. Prove that (10.25) implies the existence of a relationship between physical and human capital of the form $h = \xi(k)$, where $\xi(\cdot)$ is uniquely defined, strictly increasing and differentiable.

Exercise 10.12. Prove 10.1. Show that the differential equation for consumption growth could have alternatively been written as

$$\frac{\dot{c}(t)}{c(t)} = \frac{1}{\varepsilon_u(c(t))} \left[f_h(k(t), \xi(k(t))) - \delta_h - \rho \right].$$

Exercise 10.13. Consider the neoclassical growth model with physical and human capital discussed in Section 10.4.

(1) Specify the consumer maximization problem in this economy.

(2) Define a competitive equilibrium (specifying firm optimization and market clearing conditions).
(3) Characterize the competitive equilibrium and show that it coincides with the solution to the optimal growth problem.

Exercise 10.14. Introduce labor-augmenting technological progress at the rate g into the neoclassical growth model with physical and human capital discussed in Section 10.4.

(1) Define a competitive equilibrium.
(2) Determine transformed variables that will remain constant in a steady state allocation.
(3) Characterize the steady state equilibrium and the transitional dynamics.
(4) Why does faster technological progress lead to more rapid accumulation of human capital?

Exercise 10.15. * Characterize the optimal growth path of the economy in Section 10.4 subject to the additional constraints that $i_k(t) \geq 0$ and $i_h(t) \geq 0$.

Exercise 10.16. Derive eq. (10.26).

Exercise 10.17. Derive eq.'s (10.33) and (10.34).

Exercise 10.18. Provide conditions on $f(\cdot)$ and $\gamma(\cdot)$ such that the unique steady-state equilibrium in the model of Section 10.5 is locally stable.

Exercise 10.19. Analyze the economy in Section 10.6 under the closed economy assumption. Show that an increase in a_1 for group 1 will now create a dynamic externality, in the sense that current output will increase and this will lead to greater physical and human capital investments next periods.

Exercise 10.20. Prove Proposition 10.5.

CHAPTER 11

First-Generation Models of Endogenous Growth

The models presented so far focused on physical and human capital accumulation, and sustained economic growth has been generated by exogenous technological progress. While such models are useful in thinking about sources of income differences among countries that have (free) access to the same set of technologies, they do not generate sustained long-run growth (of the country or of the world economy) and have relatively little to say about sources of technology differences. A full analysis of both cross-country income differences and the process of world economic growth requires models in which technology choices and technological progress are endogenized. This will be the topic of the next part of the book. While models in which technology evolves as a result of firms' and workers' decisions are most attractive in this regard, sustained economic growth is possible in the neoclassical model as well. I end this part of the book by investigating sustained endogenous economic growth in neoclassical or quasi-neoclassical models.

We have already encountered the AK model in Chapter 2. This model relaxed one of the key assumptions on the aggregate production function of the economy (Assumption 2) and prevented diminishing returns to capital. Consequently, continuous capital accumulation could act as the engine of sustained economic growth. This chapter starts with a "neoclassical" version of the AK model, which not only shows the possibility of endogenous growth in the neoclassical growth model, but also provides us with a very tractable model that has multiple applications in diverse areas. This model is not without shortcomings, however. The most major one is that capital is the only (or essentially the only) factor of production, and asymptotically, the share of national income accruing to capital tends to 1. I then present two different two-sector endogenous growth models, which behave very similarly to the baseline AK model, but avoid this counterfactual prediction. The first of these incorporates physical and human capital accumulation, and is thus a close cousin of the neoclassical growth model with physical and human capital studied in Section 10.4 in Chapter 10. The second, which builds on the work by Rebelo (1991), is a substantially richer model and is also interesting since it allows investment and consumption goods sectors to have different capital intensities.

I conclude this section with a presentation of Paul Romer's (1986) path breaking article. Romer's paper started the endogenous growth literature and rejuvenated the interest in economic growth among economists. While Romer's objective was to model "technological change," he achieved this by introducing technological spillovers—similar to those we encountered in Chapter 10. Consequently, while the competitive equilibrium of Romer's model is

not Pareto optimal and the engine of economic growth can be interpreted as a form "knowledge accumulation," in many ways the model is still neoclassical in nature. In particular, we will see that in reduced-form it is very similar to the baseline AK model (except its welfare implications).

11.1. The AK Model Revisited

Let us start with the simplest neoclassical model of sustained growth, which we already encountered in the context of the Solow growth model, in particular, Proposition 2.10 in subsection 2.6. This is the so-called AK model, where the production technology is linear in capital. The analysis below will show that what matters is not the linearity on the production technology, but that of the accumulation technology. But for now it makes sense to start with the simpler case of the AK economy.

11.1.1. Demographics, Preferences and Technology. Our focus in this chapter and the next part of the book is on economic growth, and as a first pass, we will focus on balanced economic growth, defined as a growth path consistent with the Kaldor facts (recall Chapter 2). As demonstrated in Chapter 8, balanced growth forces us to adopt the standard CRRA preferences as in the canonical neoclassical growth model (to ensure a constant intertemporal elasticity of substitution).

Throughout this chapter, I assume that the economy admits an infinitely-lived representative household, with household size growing at the exponential rate n. The preferences of the representative household at time $t = 0$ are given by

$$U = \int_0^\infty \exp\left(-\left(\rho - n\right)t\right)\left[\frac{c\left(t\right)^{1-\theta}-1}{1-\theta}\right]dt. \tag{11.1}$$

Labor is supplied inelastically. The flow budget constraint facing the household can be written as

$$\dot{a}\left(t\right) = \left(r\left(t\right) - n\right)a\left(t\right) + w\left(t\right) - c\left(t\right), \tag{11.2}$$

where $a\left(t\right)$ denotes assets per capita at time t, $r\left(t\right)$ is the interest rate, $w\left(t\right)$ is the wage rate per capita, and n is the growth rate of population. As usual, let us impose the no-Ponzi game constraintalso

$$\lim_{t\to\infty}\left\{a(t)\exp\left[-\int_0^t\left[r(s)-n\right]ds\right]\right\} \geq 0. \tag{11.3}$$

The Euler equation for the representative household is the same as before and implies the following rate of consumption growth per capita:

$$\frac{\dot{c}\left(t\right)}{c\left(t\right)} = \frac{1}{\theta}(r\left(t\right) - \rho). \tag{11.4}$$

The other necessary condition for optimality of the consumer's plans is the transversality condition,

$$\lim_{t\to\infty}\left\{a(t)\exp\left[-\int_0^t\left[r(s)-n\right]ds\right]\right\} = 0. \tag{11.5}$$

As before, the problem of the consumer is concave, thus any solution to these necessary conditions is in fact an optimal plan.

The final good sector is similar to before, except that Assumptions 1 and 2 are *not* satisfied. More specifically, suppose that the aggregate production function is given by

$$Y(t) = AK(t),$$

with $A > 0$. Notice that this production function does not depend on labor, thus wage earnings, $w(t)$, in (11.2) will be equal to zero. This is one of the unattractive features of the baseline AK model, but will be relaxed below (and it is also relaxed in Exercises 11.3 and 11.4). Dividing both sides of this equation by $L(t)$, and as usual, defining $k(t) \equiv K(t)/L(t)$ as the capital-labor ratio, we obtain per capita output as

$$\begin{aligned} y(t) &\equiv \frac{Y(t)}{L(t)} \\ &= Ak(t). \end{aligned} \tag{11.6}$$

Equation (11.6) has a number of notable differences from our standard production function satisfying Assumptions 1 and 2. First, output is only a function of capital, and there are no diminishing returns (that is, it is no longer the case that $f''(\cdot) < 0$). This feature is only for simplicity and introducing diminishing returns to capital does not affect the main results in this section (see Exercise 11.4). The more important assumption is that the Inada conditions embedded in Assumption 2 are no longer satisfied. In particular,

$$\lim_{k \to \infty} f'(k) = A > 0.$$

This feature is essential for sustained growth.

The conditions for profit-maximization are similar to before, and require that the marginal product of capital be equal to the rental price of capital, $R(t) = r(t) + \delta$. Since, as is obvious from eq. (11.6), the marginal product of capital is constant and equal to A, thus $R(t) = A$ for all t, which implies that the net rate of return on the savings is constant and equal to:

$$r(t) = r = A - \delta, \text{ for all } t. \tag{11.7}$$

Since the marginal product of labor is zero, labor earnings are also equal to zero as noted above.

11.1.2. Equilibrium. A *competitive equilibrium* of this economy consists of paths of per capita consumption, capital-labor ratio, wage rates and rental rates of capital, $[c(t), k(t), w(t), R(t)]_{t=0}^{\infty}$, such that the representative household maximizes (11.1) subject to (11.2) and (11.3) given initial capital-labor ratio $k(0)$ and factor prices $[w(t), r(t)]_{t=0}^{\infty}$ such that $w(t) = 0$ for all t, and $r(t)$ is given by (11.7).

To characterize the equilibrium, note again that $a(t) = k(t)$. Next using the fact that $r = A - \delta$ and $w = 0$, eq.'s (11.2), (11.4), and (11.5) imply

$$\dot{k}(t) = (A - \delta - n)k(t) - c(t) \tag{11.8}$$

$$\frac{\dot{c}(t)}{c(t)} = \frac{1}{\theta}(A - \delta - \rho), \tag{11.9}$$

$$\lim_{t \to \infty} k(t) \exp\left(-(A - \delta - n)t\right) = 0. \tag{11.10}$$

The important result immediately follows from eq. (11.9). Since the right-hand side of this equation is constant, there must be a constant rate of consumption growth (as long as $A - \delta - \rho > 0$). The rate of growth of consumption is therefore independent of the level of capital stock per person, $k(t)$. This will also imply that there are no transitional dynamics in this model. Starting from any $k(0) > 0$, consumption per capita (and as we will see, the capital-labor ratio) will immediately start growing at a constant rate. To develop this point, let us integrate eq. (11.9) starting from some initial level of consumption $c(0)$, which as usual is still to be determined later (from the lifetime budget constraint). This gives

$$c(t) = c(0) \exp\left(\frac{1}{\theta}(A - \delta - \rho)t\right). \tag{11.11}$$

Since there is growth in this economy, we have to ensure that the transversality condition is satisfied (i.e., that lifetime utility is bounded away from infinity), and also we want to ensure positive growth (the condition $A - \delta - \rho > 0$ mentioned above). Let us therefore impose

$$A > \rho + \delta > (1 - \theta)(A - \delta) + \theta n + \delta. \tag{11.12}$$

The first part of this condition ensures that there will be positive consumption growth, while the second part is the analog to the condition that $\rho + \theta g > g + n$ in the neoclassical growth model with technological progress, which was imposed to ensure bounded utility (and thus was used in proving that the transversality condition was satisfied).

11.1.3. Equilibrium Characterization. I first establish that there are no transitional dynamics in this economy. In particular, not only the growth rate of consumption, but the growth rates of capital and output are also constant at all points in time, and equal the growth rate of consumption given in eq. (11.9).

To do this, let us substitute for $c(t)$ from eq. (11.11) into eq. (11.8), which yields

$$\dot{k}(t) = (A - \delta - n)k(t) - c(0) \exp\left(\frac{1}{\theta}(A - \delta - \rho)t\right), \tag{11.13}$$

which is a first-order, non-autonomous linear differential equation in $k(t)$. This type of equation can be solved easily. In particular,once again recalling solutions to non-autonomous differential equations (from Section B.4 in Appendix Chapter B) the solution to eq. (11.13) is

$$k(t) = \left\{\kappa \exp((A - \delta - n)t) + \left[(A - \delta)(\theta - 1)\theta^{-1} + \rho\theta^{-1} - n\right]^{-1} \left[c(0) \exp\left(\theta^{-1}(A - \delta - \rho)t\right)\right]\right\}, \tag{11.14}$$

where κ is a constant to be determined. Assumption (11.12) ensures that

$$(A - \delta)(\theta - 1)\theta^{-1} + \rho\theta^{-1} - n > 0.$$

From (11.14), it may look like capital is not growing at a constant rate, since it is the sum of two components growing at different rates. However, this is where the transversality condition becomes useful. Let us substitute from (11.14) into the transversality condition, (11.10), which yields

$$\lim_{t\to\infty}[\kappa+\left[(A-\delta)(\theta-1)\theta^{-1}+\rho\theta^{-1}-n\right]^{-1}c(0)\exp\left(-\left(A-\delta)(\theta-1)\theta^{-1}+\rho\theta^{-1}-n\right)t\right)]=0.$$

Since $(A-\delta)(\theta-1)\theta^{-1}+\rho\theta^{-1}-n>0$, the second term in this expression converges to zero as $t\to\infty$. But the first term is a constant. Thus the transversality condition can only be satisfied if $\kappa=0$. Therefore, (11.14) implies that:

$$\begin{aligned}(11.15)\qquad k(t) &= \left[(A-\delta)(\theta-1)\theta^{-1}+\rho\theta^{-1}-n\right]^{-1}\left[c(0)\exp\left(\theta^{-1}(A-\delta-\rho)t\right)\right]\\ &= k\left(0\right)\exp\left(\theta^{-1}(A-\delta-\rho)t\right),\end{aligned}$$

where the second line immediately follows from the fact that capital is equal to $k\left(0\right)$ at $t=0$. This equation naturally implies that capital and output grow at the same rate as consumption.

It also pins down the initial level of consumption as

$$(11.16)\qquad c\left(0\right)=\left[(A-\delta)(\theta-1)\theta^{-1}+\rho\theta^{-1}-n\right]k\left(0\right).$$

Note also that in this simple AK model, growth is not only sustained, but it is also endogenous in the sense of being affected by underlying parameters. For example, consider an increase in the rate of discount, ρ. Recall that in the Ramsey model, this only influenced the level of income per capita—it could have no effect on the growth rate, which was determined by the exogenous labor-augmenting rate of technological progress. Here, it is straightforward to verify that an increase in the discount rate, ρ, will reduce the growth rate, because it will make consumers less patient and will therefore reduce the rate of capital accumulation. Since capital accumulation is the engine of growth, the equilibrium rate of growth will decline. Similarly, changes in A and θ affect the levels and growth rates of consumption, capital and output.

Finally, we can calculate the saving rate in this economy. It is defined as total investment (which is equal to increase in capital plus replacement investment) divided by output. Consequently, the saving rate is constant and given by

$$\begin{aligned}s &= \frac{\dot{K}\left(t\right)+\delta K\left(t\right)}{Y\left(t\right)}\\ &= \frac{\dot{k}\left(t\right)/k\left(t\right)+n+\delta}{A}\\ &= \frac{A-\rho+\theta n+(\theta-1)\delta}{\theta A},\end{aligned}\qquad(11.17)$$

where the last equality exploited the fact that $\dot{k}\left(t\right)/k\left(t\right)=(A-\delta-\rho)/\theta$. This equation implies that the saving rate, which was taken as constant and exogenous in the basic Solow model, is again constant over time, but now depends on preference and technology. Summarizing:

Proposition 11.1. *Consider the above-described AK economy, with a representative household with preferences given by (11.1), and the production technology given by (11.6). Suppose that condition (11.12) holds. Then, there exists a unique equilibrium path in which consumption, capital and output all grow at the same rate* $g^* \equiv (A - \delta - \rho)/\theta > 0$ *starting from any initial positive capital stock per worker* $k(0)$*, and the saving rate is endogenously determined by (11.17).*

One important implication of the AK model is that since all markets are competitive, there is a representative household, and there are no externalities, the competitive equilibrium will be Pareto optimal. This can be proved either using First Welfare Theorem type reasoning, or by directly constructing the optimal growth solution.

Proposition 11.2. *Consider the above-described AK economy, with a representative household with preferences given by (11.1), and the production technology given by (11.6). Suppose that condition (11.12) holds. Then, the unique competitive equilibrium is Pareto optimal.*

Proof. See Exercise 11.2 □

11.1.4. The Role of Policy. It is straightforward to incorporate policy differences in to this framework and investigate their implications on the equilibrium growth rate. The simplest and arguably one of the most relevant classes of policies are those affecting the rate of return to accumulation, which have already been discussed in the context of the standard neoclassical growth model in Chapter 8. In particular, suppose that there is an effective tax rate of τ on the rate of return from capital income, so that the flow budget constraint of the representative household becomes:

$$\dot{a}(t) = ((1-\tau) r(t) - n) a(t) + w(t) - c(t). \tag{11.18}$$

Repeating the analysis above immediately implies that this tax will adversely affect the growth rate of the economy, which will now become (see Exercise 11.5):

$$g = \frac{(1-\tau)(A-\delta) - \rho}{\theta}. \tag{11.19}$$

Moreover, it can be calculated that the saving rate will now be

$$s = \frac{(1-\tau) A - \rho + \theta n - (1 - \tau - \theta)\delta}{\theta A}, \tag{11.20}$$

which is a decreasing function of τ if $A - \delta > 0$. Therefore, in this model, the equilibrium saving rate is constant as in the basic Solow model, but in contrast to that model, it responds endogenously to policy. In addition, the fact that the saving rate is constant implies that differences in policies will lead to permanent differences in the rate of capital accumulation. This observation has a very important implication. While in the baseline neoclassical growth model, even reasonably large differences in distortions (for example, eightfold differences in τ) could only have limited effects on differences in income per capita, here even small differences in τ can have very large effects. In particular, consider two economies, with respective

(constant) tax rates on capital income τ and $\tau' > \tau$, and exactly the same technology and preferences otherwise. It is straightforward to verify that for any $\tau' > \tau$,

$$\lim_{t \to \infty} \frac{Y(\tau', t)}{Y(\tau, t)} = 0,$$

where $Y(\tau, t)$ denotes aggregate output in the economy with tax τ at time t. Therefore, even small policy differences can have very large effects in the long run. So why does the literature focus on the inability of the standard neoclassical growth model to generate large differences rather than the possibility that the AK model can generate arbitrarily large differences? The reason is twofold: first, as noted above, the AK model, with no diminishing returns and the share of capital in national income asymptoting to 1, is not viewed as a good approximation to reality. Second, and related to our discussion in Chapter 1, most economists believe that the relative stability of the world income distribution in the post-war era makes it more attractive to focus on models in which there is a stationary world income distribution, rather than models in which small policy differences can lead to permanent growth differences. Whether this last belief is justified is, in part, an empirical question.

11.2. The AK Model with Physical and Human Capital

As pointed out in the previous section, a major shortcoming of the baseline AK model is that the share of capital accruing to national income is equal to 1 (or limits to 1 as in the variant of the AK model studied in Exercises 11.3 and 11.4). One way of enriching the AK model and avoiding these problems is to include both physical and human capital. I now briefly discuss this extension. Suppose the economy admits a representative household with preferences given by (11.1). The production side of the economy is represented by the aggregate production function

$$Y(t) = F(K(t), H(t)), \tag{11.21}$$

where $H(t)$ denotes efficiency units of labor (or human capital), which will be accumulated in the same way as physical capital. The production function $F(\cdot, \cdot)$ now satisfies Assumptions 1 and 2. Suppose, to simplify the analysis, that there is no population growth, thus $n = 0$.

The budget constraint of the representative household is given by

$$\dot{a}(t) = r(t) a(t) + w(t) h(t) - c(t) - i_h(t), \tag{11.22}$$

where $h(t)$ denotes the effective units of labor (human capital) of the representative household, $w(t)$ is wage rate per unit of human capital, and $i_h(t)$ is investment in human capital. The human capital of the representative household evolves according to the differential equation:

$$\dot{h}(t) = i_h(t) - \delta_h h(t), \tag{11.23}$$

where δ_h is the depreciation rate of human capital. The evolution of the capital stock is again given from the observation that $k(t) = a(t)$, and the depreciation rate of physical capital is now denoted by δ_k to avoid confusion with δ_h. In this model, the representative household

maximizes its utility by choosing the paths of consumption, human capital investments and asset holdings. Competitive factor markets imply that

$$R(t) = f'(k(t)) \text{ and } w(t) = f(k(t)) - k(t) f'(k(t)), \tag{11.24}$$

where, now, the effective capital-labor ratio is given by dividing the capital stock by the stock of human capital in the economy,

$$k(t) \equiv \frac{K(t)}{H(t)}.$$

A competitive equilibrium of this economy consists of paths of per capita consumption, capital-labor ratio, wage rates and rental rates of capital, $[c(t), k(t), w(t), R(t)]_{t=0}^{\infty}$, such that the representative household maximizes (11.1) subject to (11.3), (11.22) and (11.23) given initial effective capital-labor ratio $k(0)$ and factor prices $[w(t), R(t)]_{t=0}^{\infty}$ that satisfy (11.24).

To characterize the competitive equilibrium, let us first set up at the current-value Hamiltonian for the representative household with costate variables μ_a and μ_h:

$$\begin{aligned}\mathcal{H}(a, h, c, i_h, \mu_a, \mu_k) &= \frac{c(t)^{1-\theta} - 1}{1-\theta} + \mu_a(t)\left[(r(t) a(t) + w(t) h(t) - c(t) - i_h(t)\right] \\ &\quad + \mu_h(t)\left[i_h(t) - \delta_h h(t)\right].\end{aligned}$$

Now the necessary conditions of this optimization problem imply the following (see Exercise 11.8):

$$\begin{aligned}\mu_a(t) &= \mu_h(t) = \mu(t) \text{ for all } t \\ w(t) - \delta_h &= r(t) \qquad \text{for all } t \\ \frac{\dot{c}(t)}{c(t)} &= \frac{1}{\theta}(r(t) - \rho) \text{ for all } t.\end{aligned} \tag{11.25}$$

Intuitively, there are no constraints on human and physical capital investments, thus the shadow values of these two different types of investments have to be equal at all points in time as stated in the first condition in (11.25). This in turn yields the second condition in (11.25), equating the rates of return on human and physical capital. The third condition is the standard Euler equation. It can be verified that the current-value Hamiltonian is concave and and satisfies the sufficiency conditions in Theorem 7.14 in Chapter 7. Therefore, a solution to the conditions in (11.25) necessarily solves the representative household's maximization problem.

Combining these with (11.24),

$$f'(k(t)) - \delta_k = f(k(t)) - k(t) f'(k(t)) - \delta_h \text{ for all } t.$$

Since the left-hand side is decreasing in $k(t)$, while the right-hand side is increasing, this implies that the effective capital-labor ratio must satisfy

$$k(t) = k^* \text{ for all } t.$$

This implies:

Proposition 11.3. *Consider the above-described AK economy with physical and human capital, with a representative household with preferences given by (11.1), and the production technology given by (11.21). Let k^* be given by*

$$f'(k^*) - \delta_k = f(k^*) - k^* f'(k^*) - \delta_h. \tag{11.26}$$

Suppose that $f'(k^*) > \rho + \delta_k > (1-\theta)(f'(k^*) - \delta) + \delta_k$. Then, in this economy there exists a unique equilibrium path in which consumption, human capital, physical capital and output all grow at the same rate $g^* \equiv (f'(k^*) - \delta_k - \rho)/\theta > 0$ starting from any initial conditions, where k^* is given by (11.26).The share of capital in national income is constant and less than 1 at all times.

Proof. See Exercise 11.9 □

The advantage of the economy studied here, especially as compared to the baseline AK model is that, it generates a stable factor distribution of income, with a significant fraction of national income accruing to labor (as rewards to human capital). Consequently, the current model cannot be criticized on the basis of generating counter-factual results on the capital share of GDP. A similar analysis to that in the previous section also shows that the current model generates long-run growth rate differences from small policy differences. Therefore, it can account for arbitrarily large differences in income per capita across countries. Nevertheless, it would do so partly by generating large human capital differences across countries. As such, the empirical mechanism through which these large cross-country income differences are generated may again not fit with the empirical patterns discussed in Chapter 3. Moreover, given substantial differences in policies across economies in the postwar period, like the baseline AK economy, the current model would suggest significant changes in the world income distribution, whereas the evidence in Chapter 1 points to a relatively stable postwar world income distribution.

11.3. The Two-Sector AK Model

The models studied in the previous two sections are attractive in many respects; they generate sustained growth, and the equilibrium growth rate responds to policy, to underlying preferences and to technology. Moreover, these are very close cousins of the neoclassical model. In fact, as argued there, the endogenous growth equilibrium is Pareto optimal.

One unattractive feature of the baseline AK model is that all of national income accrues to capital. Essentially, it is a one-sector model with only capital as the factor of production. This makes it difficult to apply this model to real world situations. The model in the previous section avoids this problem, but at some level it does so by creating another factor of production that accumulates linearly, so that the equilibrium structure is again equivalent to the one-sector AK economy. Therefore, in some deep sense, the economies of both sections are one-sector models. More important than this one-sector property, these models potentially blur key underlying characteristic driving growth in these environments. What is important is not that the production technology is AK, but the related feature that the *accumulation*

technology is linear. In this section, I discuss a richer two-sector model of neoclassical endogenous growth, based on Rebelo's (1991) work. This model will generate constant factor shares in national income without introducing human capital accumulation. Perhaps more importantly, it will illustrate the role of differences in the capital intensity of the production functions of consumption and investment.

The preference and demographics are the same as in the model of the previous section, in particular, eq.'s (11.1)-(11.5) apply as before (but with a slightly different interpretation for the interest rate in (11.4) as will be discussed below). Moreover, to simplify the analysis, suppose that there is no population growth, that is, $n = 0$, and that the total amount of labor in the economy, L, is supplied inelastically.

The main difference is in the production technology. Rather than a single good used for consumption and investment, let us now envisage an economy with two sectors. Sector 1 produces consumption goods with the following technology

$$C(t) = B(K_C(t))^{\alpha} L_C(t)^{1-\alpha}, \tag{11.27}$$

where the subscript "C" denotes that these are capital and labor used in the consumption sector, which has a Cobb-Douglas technology. In fact, the Cobb-Douglas assumption here is quite important in ensuring that the share of capital in national income is constant (see Exercise 11.12). The capital accumulation equation is given by:

$$\dot{K}(t) = I(t) - \delta K(t),$$

where $I(t)$ denotes investment. Investment goods are produced with a different technology than (11.27), however. In particular,

$$I(t) = AK_I(t). \tag{11.28}$$

The distinctive feature of the technology for the investment goods sector, (11.28), is that it is *linear* in the capital stock and does not feature labor. This is an extreme version of an assumption often made in two-sector models, that the investment-good sector is more capital-intensive than the consumption-good sector. In the data, there seems to be some support for this, though the capital intensities of many sectors have been changing over time as the nature of consumption and investment goods has changed.

Market clearing implies:

$$K_C(t) + K_I(t) \leq K(t),$$

for capital, and

$$L_C(t) \leq L,$$

for labor (since labor is only used in the consumption sector).

An equilibrium in this economy is defined similarly to that in the neoclassical economy, but also features an allocation decision of capital between the two sectors. Moreover, since the two sectors are producing two different goods, consumption and investment goods, there will be a relative price between the two sectors that will adjust endogenously.

Since both market clearing conditions will hold as equalities (the marginal product of both factors is always positive), we can simplify notation by letting $\kappa(t)$ denote the share of capital used in the investment sector

$$K_C(t) = (1-\kappa(t))K(t) \text{ and } K_I(t) = \kappa(t)K(t).$$

From profit maximization, the rate of return to capital has to be the same when it is employed in the two sectors. Let the price of the investment good be denoted by $p_I(t)$ and that of the consumption good by $p_C(t)$, then, eq.'s (11.27) and (11.28) imply

$$p_I(t)A = p_C(t)\alpha B\left(\frac{L}{(1-\kappa(t))K(t)}\right)^{1-\alpha}. \tag{11.29}$$

Define a steady-state (a BGP) as an equilibrium path in which $\kappa(t)$ is constant and equal to some $\kappa^* \in [0,1]$. Moreover, let us choose the consumption good as the numeraire, so that $p_C(t) = 1$ for all t. Then, differentiating (11.29) implies that at the steady state:

$$\frac{\dot{p}_I(t)}{p_I(t)} = -(1-\alpha)g_K, \tag{11.30}$$

where g_K is the steady-state (BGP) growth rate of capital.

As noted above, the Euler equation for consumers, (11.4), still holds, but the relevant interest rate has to be for *consumption-denominated loans*, denoted by $r_C(t)$. This is the interest rate that measures how many units of consumption good an individual will receive tomorrow by giving up one unit of consumption today. Since the relative price of consumption goods and investment goods is changing over time, the proper calculation goes as follows. By giving up one unit of consumption, the individual will buy $1/p_I(t)$ units of capital goods. This will have an instantaneous return of $r_I(t)$. In addition, the individual will get back the one unit of capital, which has now experienced a change in its price of $\dot{p}_I(t)/p_I(t)$, and finally, he will have to buy consumption goods, whose prices changed by $\dot{p}_C(t)/p_C(t)$. Therefore, the general formula of the rate of return denominated in consumption goods in terms of the rate of return denominated in investment goods is

$$r_C(t) = \frac{r_I(t)}{p_I(t)} + \frac{\dot{p}_I(t)}{p_I(t)} - \frac{\dot{p}_C(t)}{p_C(t)}. \tag{11.31}$$

Given the choice of numeraire, $\dot{p}_C(t)/p_C(t) = 0$. Moreover, $\dot{p}_I(t)/p_I(t)$ is given by (11.30). Finally,

$$\frac{r_I(t)}{p_I(t)} = A - \delta$$

given the linear technology in (11.28). Therefore,

$$r_C(t) = A - \delta + \frac{\dot{p}_I(t)}{p_I(t)}.$$

and in steady state, from (11.30), the steady-state consumption-denominated rate of return is:

$$r_C = A - \delta - (1-\alpha)g_K.$$

From (11.4), this implies a consumption growth rate of

$$g_C \equiv \frac{\dot{C}(t)}{C(t)} = \frac{1}{\theta}\left(A - \delta - (1-\alpha) g_K - \rho\right). \tag{11.32}$$

Finally, differentiate (11.27) and use the fact that labor is always constant to obtain

$$\frac{\dot{C}(t)}{C(t)} = \alpha \frac{\dot{K}_C(t)}{K_C(t)},$$

which, from the constancy of $\kappa(t)$ in steady state, implies the following steady-state relationship:

$$g_C = \alpha g_K.$$

Substituting this into (11.32),

$$g_K^* = \frac{A - \delta - \rho}{1 - \alpha(1-\theta)} \tag{11.33}$$

and

$$g_C^* = \alpha \frac{A - \delta - \rho}{1 - \alpha(1-\theta)}. \tag{11.34}$$

What about wages? Because labor is being used in the consumption good sector, there will be positive wages. Since labor markets are competitive, the wage rate at time t is given by

$$w(t) = (1-\alpha) p_C(t) B \left(\frac{(1-\kappa(t)) K(t)}{L}\right)^{\alpha}.$$

Therefore, in the BGP,

$$\begin{aligned} \frac{\dot{w}(t)}{w(t)} &= \frac{\dot{p}_C(t)}{p_C(t)} + \alpha \frac{\dot{K}(t)}{K(t)} \\ &= \alpha g_K^*, \end{aligned}$$

which implies that wages also grow at the same rate as consumption.

Moreover, with exactly the same arguments as in the previous section, it can be established that there are no transitional dynamics in this economy. This establishes the following result:

PROPOSITION 11.4. *In the above-described two-sector neoclassical economy, starting from any $K(0) > 0$, consumption and labor income grow at the constant rate given by (11.34), while the capital stock grows at the constant rate (11.33).*

It is straightforward to conduct policy analysis in this model, and as in the basic AK model, taxes on investment income or other policies that discourage investment will depress growth. Similarly, a lower discount rate will increase the equilibrium growth rate of the economy

One important implication of this model, different from the neoclassical growth model, is that there is continuous *capital deepening.* Capital grows at a faster rate than consumption and output. Whether this is a realistic feature is debatable. The Kaldor facts, discussed above, include constant capital-output ratio as one of the requirements of balanced growth.

The "balanced growth" here does not have this feature. For much of the 20th century, capital-output ratio has been constant, but it has been increasing steadily over the past 30 years. Part of the reason why it has been increasing recently, but not before, is because of relative price adjustments. New capital goods are of higher quality, and this needs to be incorporated in calculating the capital-output ratio. These calculations have only been performed in the recent past, which may explain why capital-output ratio has been constant in the earlier part of the century, but not recently. In view of this, it is not clear whether a constant or an increasing capital-output ratio is a better approximation to reality.

11.4. Growth with Externalities

The model that started much of endogenous growth theory and revived economists' interest in economic growth was Paul Romer's (1986) paper. Romer's objective was to model the process of "knowledge accumulation". He realized that this would be difficult in the context of a competitive economy. His initial solution (later updated and improved in his and others' work during the 1990s) was to consider knowledge accumulation to be a *byproduct* of capital accumulation. In other words, Romer introduced technological spillovers, similar to those discussed in the context of human capital in Chapter 10. While arguably crude, this captures an important dimension of knowledge, that knowledge is a largely *non-rival* good—once a particular technology has been discovered, many firms can make use of this technology without preventing others using the same knowledge. Non-rivalry does not imply knowledge is also non-excludable (which would have made it a pure public good). A firm that discovers a new technology may use patents or trade secrecy to prevent others from using it, for example, in order to gain a competitive advantage. These issues will be discussed in the next part of the book. For now, it suffices to note that some of the important characteristics of "knowledge" and its role in the production process can be captured in a reduced-form way by introducing technological spillovers. I next discuss a version of the model in Romer's (1986) paper, which introduces such technological spillovers as the engine of economic growth. While the type of technological spillovers used in this model are unlikely to be important in practice, this model is a good starting point for our analysis of endogenous technological progress, since its similarity to the baseline AK economy makes it a very tractable model of knowledge accumulation.

11.4.1. Preferences and Technology. Consider an economy without any population growth (we will see why this is important) and a production function with labor-augmenting knowledge (technology) that satisfies the standard assumptions, Assumptions 1 and 2. For reasons that will become clear, instead of working with the aggregate production function, let us assume that the production side of the economy consists of a set $[0,1]$ of firms. The production function facing each firm $i \in [0,1]$ is

$$Y_i(t) = F(K_i(t), A(t) L_i(t)), \tag{11.35}$$

where $K_i(t)$ and $L_i(t)$ are capital and labor rented by a firm i. Notice that $A(t)$ is not indexed by i, since it is technology common to all firms. Let us normalize the measure of final good producers to 1, so that the following market clearing conditions apply:

$$\int_0^1 K_i(t)\,di = K(t)$$

and

$$\int_0^1 L_i(t)\,di = L,$$

where L is the constant level of labor (supplied inelastically) in this economy. Firms are competitive in all markets, which implies that they will all hire the same capital to effective labor ratio, and moreover, factor prices will be given by their marginal products, thus

$$\begin{aligned} w(t) &= \frac{\partial F(K(t), A(t)L)}{\partial L} \\ R(t) &= \frac{\partial F(K(t), A(t)L)}{\partial K(t)}. \end{aligned}$$

The key assumption of Romer (1986a) is that although firms take $A(t)$ as given, this stock of technology (knowledge) advances endogenously for the economy as a whole. In particular, Romer assumes that this takes place because of spillovers across firms, and attributes spillovers to physical capital. Lucas (1988) develops a similar model in which the structure is identical, but spillovers work through human capital (while Romer has physical capital externalities, Lucas has human capital externalities).

The idea of externalities is not uncommon to economists, but both Romer and Lucas make an extreme assumption of sufficiently strong externalities such that $A(t)$ can grow continuously at the economy level. In particular, Romer assumes

$$A(t) = BK(t), \tag{11.36}$$

so that the knowledge stock of the economy is proportional to the capital stock of the economy. This can be motivated by "learning-by-doing" whereby, greater investments in certain sectors increases the experience (of firms, workers, managers) in the production process, making the production process itself more productive. Alternatively, the knowledge stock of the economy could be a function of the cumulative output that the economy has produced up to now, thus giving it more of a flavor of "learning-by-doing". Note also that (11.35) and (11.36) imply that the aggregate production function of this economy exhibits *increasing returns to scale*. As discussed in detail in the next part of the book, this is a very common feature of models of endogenous growth.

In any case, substituting for (11.36) into (11.35) and using the fact that all firms are functioning at the same capital-effective labor ratio, the production function of the representative firm is

$$Y(t) = F(K(t), BK(t)L).$$

Using the fact that $F(\cdot,\cdot)$ is homogeneous of degree 1,

$$\begin{aligned}\frac{Y(t)}{K(t)} &= F(1,BL)\\ &= \tilde{f}(L).\end{aligned}$$

Output per capita can therefore be written as:

$$\begin{aligned}y(t) &\equiv \frac{Y(t)}{L}\\ &= \frac{Y(t)}{K(t)}\frac{K(t)}{L}\\ &= k(t)\tilde{f}(L),\end{aligned}$$

where again $k(t) \equiv K(t)/L$ is the capital-labor ratio in the economy.

As in the standard growth model, marginal products and factor prices can be expressed in terms of the normalized production function, now $\tilde{f}(L)$. In particular,

$$w(t) = K(t)\tilde{f}'(L) \tag{11.37}$$

and

$$R(t) = R = \tilde{f}(L) - L\tilde{f}'(L), \tag{11.38}$$

which is constant.

11.4.2. Equilibrium. A competitive equilibrium is defined similarly to the neoclassical growth model, as a path of consumption and capital stock for the economy, $[C(t), K(t)]_{t=0}^{\infty}$ that maximize the utility of the representative household and wage and rental rates $[w(t), R(t)]_{t=0}^{\infty}$ that clear markets. The important feature is that because the knowledge spillovers, as specified in (11.36), are external to each firm, factor prices are given by (11.37) and (11.38)—that is, they do not price the role of the capital stock in increasing future productivity.

Since the market rate of return is $r(t) = R(t) - \delta$, it is also constant. The usual consumer Euler equation (e.g., (11.4) above) then implies that consumption must grow at the constant rate,

$$g_C^* = \frac{1}{\theta}\left(\tilde{f}(L) - L\tilde{f}'(L) - \delta - \rho\right). \tag{11.39}$$

It is also clear that capital grows exactly at the same rate as consumption, so the rate of capital, output and consumption growth are all given by g_C^* as given by (11.39)—see Exercise 11.16.

Let us assume that

$$\tilde{f}(L) - L\tilde{f}'(L) - \delta - \rho > 0, \tag{11.40}$$

so that there is positive growth, but also that growth is not fast enough to violate the transversality condition, in particular,

$$(1-\theta)\left(\tilde{f}(L) - L\tilde{f}'(L) - \delta\right) < \rho. \tag{11.41}$$

Proposition 11.5. *Consider the above-described Romer model with physical capital externalities. Suppose that conditions (11.40) and (11.41) are satisfied. Then, there exists a unique equilibrium path where starting with any level of capital stock $K(0) > 0$, capital, output and consumption grow at the constant rate (11.39).*

Proof. Much of this proposition is proved in the preceding discussion. You are asked to verify the transversality conditions and show that there are no transitional dynamics in Exercise 11.17. □

This proposition (and the accompanying analysis) therefore provides us with the first example of endogenous technological change. The technology of the economy, $A(t)$ as given in (11.36), evolves endogenously as a result of the investment decisions of firms. Consequently, the growth rate of the economy is endogenous, even though none of the firms purposefully invest in research or in acquiring new technologies.

Population must be constant in this model because of the *scale effect*. Since $\tilde{f}(L) - L\tilde{f}'(L)$ is always increasing in L (by Assumption 1), a higher population (labor force) L leads to a higher growth rate. The scale effect refers to this relationship between population and the equilibrium rate of economic growth. Now if population is growing, the economy will not admit a steady state and the growth rate of the economy will increase over time (output reaching infinity in finite time and violating the transversality condition). The implications of positive population growth are discussed further in Exercise 11.18. Scale effects and how they can be removed will be discussed in detail in Chapter 13.

11.4.3. Pareto Optimal Allocations. Given the presence of externalities, it is not surprising that the decentralized equilibrium characterized in Proposition 11.5 is not Pareto optimal. To characterize the allocation that maximizes the utility of the representative household, let us again set up the current-value Hamiltonian and look for a candidate path that satisfies the conditions in Theorem 7.13 (see Exercise 11.15). The per capita accumulation equation for this economy can be written as

$$\dot{k}(t) = \tilde{f}(L) k(t) - c(t) - \delta k(t).$$

The current-value Hamiltonian is

$$\hat{H}(k, c, \mu) = \frac{c(t)^{1-\theta} - 1}{1-\theta} + \mu(t) \left[\tilde{f}(L) k(t) - c(t) - \delta k(t) \right],$$

and has the necessary conditions:

$$\begin{aligned} \hat{H}_c(k, c, \mu) &= c(t)^{-\theta} - \mu(t) = 0 \\ \hat{H}_k(k, c, \mu) &= \mu(t) \left[\tilde{f}(L) - \delta \right] = -\dot{\mu}(t) + \rho \mu(t), \\ \lim_{t \to \infty} \left[\exp(-\rho t) \mu(t) k(t) \right] &= 0. \end{aligned}$$

Using standard arguments (recall Section 7.7 in Chapter 7), it is straightforward to verify that the current-value Hamiltonian satisfies the conditions in Theorem 7.14, so that these conditions are sufficient for a Pareto optimum (see again Exercise 11.15).

Combining these equations immediately yields that the social planner's allocation will also have a constant growth rate for consumption (and output) given by

$$g_C^S = \frac{1}{\theta}\left(\tilde{f}(L) - \delta - \rho\right),$$

which is always greater than g_C^* as given by (11.39)—since $\tilde{f}(L) > \tilde{f}(L) - L\tilde{f}'(L)$. Essentially, the social planner takes into account that by accumulating more capital, she is improving productivity in the future. Since this effect is external to the firms, the decentralized economy fails to internalize this spillover. Therefore:

Proposition 11.6. *In the above-described Romer model with physical capital externalities, the decentralized equilibrium is Pareto suboptimal and grows at a slower rate than the allocation that would maximize the utility of the representative household.*

Exercise 11.19 asks you to characterize various different types of policies that can close the gap between the equilibrium and Pareto optimal allocations.

11.5. Taking Stock

This chapter ends our investigation of neoclassical growth models. It also opens the way for the analysis of endogenous technological progress in the next part of the book. The models presented in this chapter are, in many ways, more tractable and easier than those in earlier chapters. This is a feature of the linearity of the models (most clearly visible in the AK model). This type of linearity removes transitional dynamics and leads to a more tractable mathematical structure. Linearity, of course, is an essential feature of any model that will exhibit sustained economic growth. If strong concavity sets in (especially concavity consistent with the Inada conditions as in Assumption 2), sustained growth will not be possible. Therefore, (asymptotic) linearity is an essential ingredient of any model that will lead to sustain growth. The baseline AK model and its cousins make this linear structure quite explicit. While this type of linearity will be not as apparent (and often will be derived rather than assumed), it will also be a feature of the endogenous technology models studied in the next part of the book. Consequently, many of these endogenous technology models will be relatively tractable as well. Nevertheless, we will see that the linearity will often result from much more interesting economic interactions than being imposed in the aggregate production function of the economy. There is another sense in which the material in this chapter does not do justice to issues of sustained growth. As the discussion in Chapter 3 showed, modern economic growth is largely the result of technological progress. Except for the Romer model of Section 11.4, in the models studied in this chapter do not feature technological progress. This does not imply that they are necessarily inconsistent with the data. As already noted in Chapter 3, there is a lively debate about whether the observed total factor productivity growth is partly a result of mismeasurement of inputs. If this is the case, it could be that much of what we measure as technological progress is in fact capital deepening, which is the bread-and-butter of economic growth in the AK model and

its variants. Consequently, the debate about the measurement of total factor productivity has important implications for what types of models we should use for thinking about world economic growth and cross-country income differences.

The discussion in this chapter has also revealed another important tension. Chapters 3 and 8 demonstrated that the neoclassical growth model (or the simpler Solow growth model) have difficulty in generating the very large income differences across countries that we observe in the data. Even if we choose quite large differences in cross-country distortions (for example, eightfold differences in effective tax rates), the implied steady-state differences in income per capita are relatively modest. As noted before, this observation has generated a large literature that seeks reasonable extensions of the neoclassical growth model in order to derive more elastic responses to policy distortions or other differences across countries. The models presented in this chapter, like those that we will encounter in the next part of the book, suffer from the opposite problem. They imply that even small differences in policies, technological opportunities or other characteristics of societies will lead to permanent differences in long-run growth rates. Consequently, these models can explain very large differences in living standards from small policy, institutional or technological differences. But this is both a blessing and a curse. Though capable of explaining large cross-country differences, these models also predict an ever-expanding world distribution, since countries with different characteristics should grow at permanently different rates. The relative stability of the world income distribution in the postwar era is then a challenge to the baseline endogenous growth models. However, as emphasized in Chapter 1, the world income distribution is not exactly stationary. While economists more sympathetic to the exogenous growth version of the neoclassical model emphasize the relative stability of the world income distribution, others see stratification and increased inequality. This debate can, in principle, be resolved by carefully mapping various types of endogenous growth theories to postwar data.

Nevertheless, there is more to understanding the nature of the growth process and the role of technological progress than simply looking at the postwar data. First, as illustrated in Chapter 1, the era of divergence is not the past 60 years, but the 19th century. Therefore, it is equally important to confront these models with historical data. Second, a major assumption of most endogenous growth models is that each country can be treated in isolation. This "each country as an island" approach is unlikely to be a good approximation to reality in most circumstances, and much less so when technology is endogenous. Most economies do not generate their own technology by R&D or other processes, but partly import or adopt these technologies from more advanced nations (or from the world technology frontier). Consequently, a successful mapping of the theories to data requires us to enrich these theories and abandon the "each country as an island" assumption. I discuss the number of approaches that incorporate technology flows across countries and international trade linkages later in the book. But the next part will follow the established literature and develop the models of endogenous technological progress without paying much attention to cross-country knowledge flows.

11.6. References and Literature

The AK model is a special case of Rebelo's (1991) model, which was discussed in greater detail in Section 11.3 of this chapter. Solow's (1965) book also discussed the AK model (naturally with exogenous savings), but dismissed it as uninteresting. A more complete treatment of sustained neoclassical economic growth is provided in Jones and Manuelli (1990), who show that even convex models (with production functions that satisfy Assumption 1, but naturally not Assumption 2) are consistent with sustained long-run growth. Exercise 11.4 is a version of the convex neoclassical endogenous growth model of Jones and Manuelli.

Barro and Sala-i-Martin (2004) discuss a variety of two-sector endogenous growth models with physical and human capital, similar to the model presented in Section 11.2, though the model presented here is much simpler than similar ones analyzed in the literature.

Romer (1986a) is the seminal paper of the endogenous growth literature and the model presented in Section 11.4 is based on this paper. Frankel (1962) analyzed a similar growth economy, but with an exogenous constant saving rate. The importance of Romer's paper stems not only from the model itself, but from two other features. The first is its emphasis on potential non-competitive elements in order to generate long-run economic growth (in this case knowledge spillovers). The second is its emphasis on the non-rival nature of knowledge and ideas. These issues will be discussed in greater detail in the next part of the book.

Another paper that has played a major role in the new growth literature is Lucas (1988), which constructs an endogenous growth model similar to that of Romer (1986a), but with human capital accumulation and human capital externalities. Lucas' model is also similar to the earlier contribution, this time by Uzawa (1964). Lucas's paper has played two major roles in the literature. First, it emphasized the empirical importance of sustained economic growth and thus was instrumental in generating interest in the newly emerging endogenous growth models. Second, it emphasized the importance of human capital and especially of human capital externalities. Since the role of human capital was discussed extensively in Chapter 10, which also showed that the evidence for human capital externalities is rather limited, we focused on the Romer model rather than the Lucas model. It turns out that Lucas model also generates transitional dynamics, which are slightly more difficult to characterize than the standard neoclassical transitional dynamics. A version of the Lucas model is discussed in Exercise 11.21.

11.7. Exercises

Exercise 11.1. Derive eq. (11.14).

Exercise 11.2. Prove Proposition 11.2.

Exercise 11.3. Consider the following continuous-time neoclassical growth model:

$$U(0)=\int_0^{\infty}\exp(-\rho t)\frac{(c(t))^{1-\theta}-1}{1-\theta},$$

with aggregate production function

$$Y(t) = AK(t) + BL(t),$$

where $A, B > 0$.

(1) Define a competitive equilibrium for this economy.
(2) Set up the current-value Hamiltonian for an individual and characterize the necessary conditions for consumer maximization. Combine these with equilibrium factor prices and derive the equilibrium path. Show that the equilibrium path displays non-trivial transitional dynamics.
(3) Determine the evolution of the labor share of national income over time.
(4) Analyze the impact of an unanticipated increase in B on the equilibrium path.
(5) Prove that the equilibrium is Pareto optimal.

Exercise 11.4. Consider the following continuous-time neoclassical growth model:

$$U(0) = \int_0^{\infty} \exp(-\rho t) \frac{(c(t))^{1-\theta} - 1}{1-\theta},$$

with production function

$$Y(t) = A\left[L(t)^{\frac{\sigma-1}{\sigma}} + K(t)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}.$$

(1) Define a competitive equilibrium for this economy.
(2) Set up the current-value Hamiltonian for an individual and characterize the necessary conditions for consumer maximization. Combine these with equilibrium factor prices and derive the equilibrium path.
(3) Prove that the equilibrium is Pareto optimal in this case.
(4) Show that if $\sigma \leq 1$, sustained growth is not possible.
(5) Show that if A and σ are sufficiently high, this model generates asymptotically sustained growth due to capital accumulation. Interpret this result.
(6) Characterize the transitional dynamics of the equilibrium path.
(7) What is happening to the share of capital in national income? Is this plausible? How would you modify the model to make sure that the share of capital in national income remains constant?
(8) Now assume that returns from capital are taxed at the rate τ. Determine the asymptotic growth rate of consumption and output.

Exercise 11.5. Derive eq.'s (11.19) and (11.20).

Exercise 11.6. Consider the neoclassical growth model with Cobb-Douglas technology $y(t) = Ak(t)^{\alpha}$ (expressed in per capita terms) and log preferences. Characterize the equilibrium path of this economy and show that as $\alpha \to 1$, the equilibrium path approaches that of the baseline AK economy. Interpret this result.

Exercise 11.7. Consider the baseline AK model of Section 11.1 and suppose that two otherwise-identical countries have different taxes on the rate of return on capital. Consider the following calibration of the model where $A = 0.15$, $\delta = 0.05$, $\rho = 0.02$, and $\theta = 3$.

Suppose that the first country has a capital income tax rate of $\tau = 0.2$, while the second country has a tax rate of $\tau' = 0.4$. Suppose that the two countries start with the same level of income in 1900 and experience no change in technology or policies for the next 100 years. What will be the relative income gap between the two countries in 2000? Discuss this result and explain why you do (or do not) find the implications plausible.

Exercise 11.8. (1) Verify that Theorems 7.13 and 7.14 from Chapter 7 can be applied both to the two-sector model in Section 11.2. [Hint: use a similar argument that in Section 7.7 in Chapter 7].

(2) Prove that the necessary conditions for consumer optimization in Section 11.2 lead to the conditions enumerated in (11.25).

Exercise 11.9. Prove Proposition 11.3.

Exercise 11.10. Prove that the competitive equilibrium of the economy in Section 11.2, characterized in Proposition 11.3, is Pareto optimal and coincides with the solution to the optimal growth problem.

Exercise 11.11. Show that the rate of population growth has no effect on the equilibrium growth rate of the economies studied in Sections 11.1 and 11.2. Explain why this is. Do you find this to be a plausible prediction?

Exercise 11.12. * Show that in the model of Section 11.3, if the Cobb-Douglas assumption is relaxed, there will not exist a BGP with a constant share of capital income in GDP.

Exercise 11.13. Consider the effect of an increase in α on the competitive equilibrium of the model in Section 11.3. Why does it increase the rate of capital accumulation in the economy?

Exercise 11.14. Consider a variant of the model studied in Section 11.3, where the technology in the consumption-good sector is still given by (11.27), while the technology in the investment-good sector is modified to

$$I(t) = A\left(K_I(t)\right)^{\beta}\left(L_I(t)\right)^{1-\beta},$$

where $\beta \in (\alpha, 1)$. The labor market clearing condition requires $L_C(t) + L_I(t) \leq L(t)$. The rest of the environment is unchanged.

(1) Define a competitive equilibrium.

(2) Characterize the steady-state equilibrium and show that it does not involve sustained growth.

(3) Explain why the long-run growth implications of this model differ from those of Section 11.3.

(4) Analyze the steady-state income differences between two economies taxing capital at the rates τ and τ'. What are the roles of the parameters α and β in determining these relative differences? Why do the implied magnitudes differ from those in the one-sector neoclassical growth model?

Exercise 11.15. In the Romer model presented in Section 11.4, verify that Theorems 7.13 and 7.14 from Chapter 7 can be applied both to the representative household's problem and to the social planner's problem.

Exercise 11.16. In the Romer model presented in Section 11.4, let g_C^* be the growth rate of consumption and g^* the growth rate of aggregate output. Show that $g_C^* > g^*$ is not feasible, while $g_C^* < g^*$ would violate the transversality condition.

Exercise 11.17. Consider the Romer model presented in Section 11.4. Prove that the allocation in Proposition 11.5 satisfies the transversality condition. Prove also that there are no transitional dynamics in this equilibrium.

Exercise 11.18. Consider the Romer model presented in Section 11.4 and suppose that population grows at the rate $n > 0$. Characterize the labor market clearing conditions. Formulate the dynamic optimization problem of a representative household and show that any interior solution to this problem violates the transversality condition. Interpret this result.

Exercise 11.19. Consider the Romer model presented in Section 11.4. Provide two different types of tax/subsidy policies that would make the equilibrium allocation identical to the Pareto optimal allocation.

Exercise 11.20. Consider the following infinite-horizon economy in discrete time that admits a representative household with preferences at time $t = 0$ as

$$U(0) = \sum_{t=0}^{\infty} \beta^t \left[\frac{C(t)^{1-\theta} - 1}{1-\theta} \right],$$

where $C(t)$ is consumption, and $\beta \in (0,1)$. Total population is equal to L and there is no population growth and labor is supplied inelastically. The production side of the economy consists of a continuum 1 of firms, each with production function

$$Y_i(t) = F(K_i(t), A(t) L_i(t)),$$

where $L_i(t)$ is employment of firm i at time t, $K_i(t)$ is capital used by firm i at time t, and $A(t)$ is a common technology term. Market clearing implies that $\int_0^1 K_i(t)\, di = K(t)$, where $K(t)$ is the total capital stock at time t, and $\int_0^1 L_i(t)\, di = L(t)$. Assume that capital fully depreciates, so that the resource constraint of the economy is

$$K(t+1) = \int_0^1 Y_i(t)\, di - C(t).$$

Assume also that labor-augmenting productivity at time t, $A(t)$, is given by

$$A(t) = K(t). \tag{11.42}$$

(1) Explain (11.42) and why it implies a (non-pecuniary) externality.
(2) Define a competitive equilibrium (where all agents are price takers).
(3) Show that there exists a unique BGP competitive equilibrium, where the economy grows (or shrinks) at a constant rate every period. Provide a condition on F, β and θ such that this growth rate is positive, but the transversality condition is still satisfied.
(4) Explain why any equilibrium must be along the BGP characterized in part 3 at all points.

Exercise 11.21. * Consider the following endogenous growth model due to Uzawa and Lucas. The economy admits a representative household and preferences are given by

$$\int_0^\infty \exp(-\rho t)\frac{C(t)^{1-\theta}-1}{1-\theta}dt,$$

where $C(t)$ is consumption of the final good, which is produced as

$$Y(t) = AK(t)^{\alpha} H_P^{1-\alpha}(t)$$

where $K(t)$ is capital and $H(t)$ is human capital, and $H_P(t)$ denotes human capital used in production. The accumulation equations are as follows:

$$\dot{K}(t) = I(t) - \delta K(t)$$

for capital and

$$\dot{H}(t) = BH_E(t) - \delta H(t)$$

where $H_E(t)$ is human capital devoted to education (further human capital accumulation), and the depreciation of human capital is assumed to be at the same rate, δ, as physical capital for simplicity. The resource constraints of the economy are

$$I(t) + C(t) \leq Y(t)$$

and

$$H_E(t) + H_P(t) \leq H(t).$$

(1) Interpret the second resource constraint.

(2) Denote the fraction of human capital allocated to production by $\phi(t)$ (so that $\phi(t) \equiv H_P(t)/H(t)$) and calculate the growth rate of final output as a function of $\phi(t)$ and the growth rates of accumulable factors.

(3) Assume that $\phi(t)$ is constant, and characterize the BGP of the economy (with constant interest rate and constant rate of growth for capital and output). Show that in this BGP, $r^* \equiv B - \delta$ and the growth rate of consumption, capital, human capital and output are given by $g^* \equiv (B - \delta - \rho)/\theta$. Show also that there exists a unique value of $k^* \equiv K/H$ consistent with BGP.

(4) Determine the parameter restrictions to make sure that the transversality condition is satisfied.

(5) Now analyze the transitional dynamics of the economy starting with K/H different from k^* [Hint: look at dynamics in three variables, $k \equiv K/H$, $\chi \equiv C/K$ and ϕ, and consider the cases $\alpha < \theta$ and $\alpha \geq \theta$ separately].

Part 4

Endogenous Technological Change

This part of the book focuses on models of endogenous technological change. Chapter 12 discusses various different approaches to technological change and provides a brief overview of some models of technological progress from the industrial organization literature. Chapters 13 and 14 present the baseline endogenous technological progress models developed by Romer, Grossman and Helpman and Aghion and Howitt. Chapter 15 considers a richer class of models in which the direction of technological change, for example, which factors technological change will augment or complement, is also endogenous. The models presented in this part of the book are useful for two related purposes. First, they enable us to endogenize technology, thus allowing a more in-depth study of cross-country and over-time differences in technologies. This is useful since, as discussed in Chapter 3, differences in technology and in the efficiency with which factors of production are used are a major proximate cause of growth over time and of cross-country differences in economic performance. Second, endogenous technology models provide a tractable approach to modeling sustained growth, which will be later combined with other factors to investigate the determinants of country and world growth rates.

CHAPTER 12

Modeling Technological Change

We have so far investigated models of economic growth of exogenous or endogenous variety. But economic growth has not resulted from technological change. Either it has been exogenous, or it has been sustained because of a linear neoclassical technology, or it has taken place as a byproduct of knowledge spillovers. Since our purpose is to understand the process of economic growth, models in which growth results from technological progress and technological change itself is a consequence of purposeful investments by firms and individuals are much more attractive. These models not only endogenize technological progress, but they also relate the process of technological change to the underlying market structure, anti-trust and competition policy, and intellectual property rights policy. They will also enable us to discuss issues of directed technological change. In this chapter, we begin with a brief discussion of different conceptions of technological change and provide some foundations for the models that will come later.

12.1. Different Conceptions of Technology

12.1.1. Types of Technological Change. The literature on technological change often distinguishes between different types of innovations. A first common distinction is between *process* and *product* innovation. While the latter refers to the introduction of a new product (for example, the introduction of the first DVD player), the former is concerned with innovations that reduce the costs of production of existing products (for example, the introduction of new machines to produce existing goods). Although the distinction between process and product innovations is important in the literature and in practice, many models with process innovation, in particular, those in which new innovations introduce new types of inputs, are mathematically identical to models of product innovation. Nevertheless, the distinction is still useful in mapping some of these theories to data.

A somewhat different type of process innovation is perhaps most important in practice and involves the introduction of a cost-reducing technological improvement or a higher-quality version of an existing good. Both the introduction of a better DVD player, when there are already DVD players in the market and the innovation to manufacture exactly the same DVD player at lower costs would be examples of this type of innovation. The most important implication is that these innovations will typically lead to the *replacement* of older vintages of the same good or machine and also to potential competition between existing producers and the innovator. In addition, in the context of this type of innovation, one might want to distinguish between the introduction of a higher-quality DVD player and the production

of a cheaper DVD player because heterogeneous consumers may have differential willingness to pay for quality than for quantity. Issues of differential willingness to pay for quality are important in the theory of industrial organization and for constructing accurate quality-adjusted price indices. However, most growth models represent the consumer side by a representative household and also implicitly assume perfect substitution between quality and quantity. These features create a close connection between innovations that increase the quality of existing products and process innovations. The following example illustrates why, in the context of the typical growth models, quality improvements and cost reductions are essentially equivalent.

Example 12.1. Consider an economy admitting a representative household with preferences $U(qc(q), y \mid q)$, where y stands for a generic good (perhaps representing all other goods), c is a particular consumption good available in different qualities. Here $c(q)$ denotes the amount consumed of the "vintage" of this good of quality q. The utility function is also conditioned on q. This specification, with q multiplying $c(q)$, implies that quality and quantity are perfect substitutes, so that higher-quality products increase the "effective units" of consumption. This is a typical assumption in growth models, though it is clearly restrictive; the consumption (use) of five 1 GHz computers would not give the same services as the use of a single 5 GHz computer.

Let the budget constraint of the representative household be

$$p(q)\, c(q) + y \leq m,$$

where $p(q)$ is the price of a good of quality q, the price of the generic good is normalized to 1, and m denotes the resources available to the consumer. The problem of the consumer can then be equivalently written as

$$\begin{aligned} &\max_{x(q),y} U(x(q), y \mid q) \\ &\text{subject to} \\ &\frac{p(q)}{q} x(q) + y \leq m, \end{aligned}$$

where $x(q) \equiv qc(q)$ corresponds to the effective units of consumption of good c. It is straightforward to see from this problem formulation that an $s\%$ increase in quality q and an $s\%$ decline in the price $p(q)$ have exactly the same effect on the effective units of consumption and on welfare. This justifies the claim above that in many models, process innovations reducing costs of production and quality improvements have identical effects.

Another important distinction in the technological change literature is between "macro" and "micro" innovations (see Mokyr, 1990). The first refers to radical innovations, including the introduction of general-purpose technologies, such as electricity or the computer, which potentially changed the organization of production in many different product lines. In contrast, micro innovations refer to the more common innovations that introduce newer models of existing products, improve the quality of a certain product line or simply reduce costs. Most

of the innovations modeled below can be viewed as "micro innovations". Moreover, empirically, it appears that micro innovations are responsible for most of the productivity growth in practice (see the evidence and discussion in Freeman, 1982, Myers and Marquis, 1969, and Abernathy, 1980). I will discuss some of the implications of macro or general-purpose innovations below.

12.1.2. A Production Function for Technology. A potentially confusing issue in the study of technological progress is how to conceptualize the menu of technologies available to firms or individuals. Since our purpose is to develop models of endogenous technology, firms and/or individuals must have a choice over different types of technologies, with greater effort, research spending and investment leading to better technologies. At some level, this implies that there must exist a *meta production function* (a production function over production functions), which determines how new technologies are generated as a function of inputs. In what follows I referred to this meta production function as the *innovation possibilities frontier* or as the R&D production function.

While a meta production function may appear natural to some, there are various economists and social scientists who do not find this a compelling approach. Their argument against the production function approach to technology is that, by its nature, innovation includes the discovery of the "unknown"; thus how could we put that in the context of a production function where inputs go in and outputs come out in a deterministic fashion?

Although this question has some descriptive merit (in the sense that describing the discovery of new technologies with a production function obscures some important details of the innovation process), the concern is largely irrelevant. There is no reason to assume that the meta production function for technology is deterministic. For example, we can assume that when a researcher puts l hours and x units of the final good into a research project, then there will be some probability $p(l, x)$ that any innovation will be made. Conditional on an innovation, the quality of the good will have a distribution $F(q \mid l, x)$. In this particular formulation, both the success of the research project and the quality of the research output conditional on success are uncertain. Nevertheless, all this can be formulated as part of the meta production function with stochastic output. Therefore, the production function approach to technology is not particularly restrictive, as long as uncertain outcomes are allowed and we are willing to assume that individuals can make calculations about the effect of their actions on the probability of success and quality of the research project. Naturally, some may argue that such calculations are not possible. But, without such calculations we would have little hope of modeling the process of technological change (or technology adoption). Since our objective is to model purposeful innovations, assuming that individuals and firms can make such calculations is entirely natural, and the existence of individuals and firms making such calculations is equivalent to assuming the existence of a meta production function for technologies.

12.1.3. Non-Rivalry of Ideas. Another important aspect of technology is emphasized in Paul Romer's work. As already discussed in the previous chapter, Romer's first model of endogenous growth, Romer (1986a), introduced increasing returns to scale to physical capital accumulation. The justification for this was that the accumulation of knowledge could be considered a byproduct of the economic activities of firms. Later work by Romer, which will be studied in the next chapter, took a very different approach to modeling the process of economic growth, but the same key idea is present in both his early and later work: the *non-rivalry* of ideas matters.

By non-rivalry, Romer means that the use of an idea by one producer to increase efficiency does not preclude its use by others. While the same unit of labor or capital cannot be used by multiple producers, the same idea can be used by many, potentially increasing everybody's productivity. Let us consider a production function of the form

$$F(K, L, A),$$

with A denoting "technology". Romer argues that an important part of this technology is the ideas or blueprints concerning how to produce new goods, how to increase quality, or how to reduce costs. Economists are generally comfortable assuming that the production function $F(K, L, A)$ exhibits constant returns to scale in capital and labor (K and L), and I adopted this assumption throughout the first three parts of the book. For example, replication arguments could be used to justify this type of constant returns to scale; when capital and labor double, the society can always open a replica of the same production facility, and in the absence of externalities, this will (at least) double output.

Romer argues that when we endogenize A, this will naturally lead to increasing returns to scale to all three inputs, K, L and A. To understand why "non-rivalry" is important here, imagine that A is just like any other input. Then, the replication argument would require the new production facility to replicate A as well, thus we should expect constant returns to scale when we vary all three inputs, K, L and A. But, instead, assume that ideas are non-rival. The new production facility does not need to re-create or to replicate A, because it is already out there available for all firms to use. In that case, $F(K, L, A)$ will exhibits constant returns in K and L, and *increasing returns to scale* in K, L and A.

Thus the non-rivalry of ideas and increasing returns to scale to all factors of production, including technology are intimately linked. This has motivated Romer to develop different types of endogenous growth models, exhibiting different sources of increasing returns to scale, but the non-rivalry of ideas has been a central element in all of them.

Another important implication of the non-rivalry of ideas is the *market size* effect. If, once discovered, an idea can be used as many times as one wishes, then the size of its potential market will be a crucial determinant of whether or not it is profitable to implement it and whether to research it in the first place. This is well captured by a famous quote from Matthew Boulton, James Watt's business partner, who wrote to Watt:

> "It is not worth my while to manufacture your engine for three countries only, but I find it very well worth my while to make it for all the world." (quoted in Scherer, 1984, p. 13).

To see why non-rivalry is related to the market size effect, imagine another standard (rival) input that is also essential for production. A greater market size does not necessarily induce firms to use this other input more intensively, since a greater market size and thus greater sales means that more of this input has to be used. It is the fact that, once invented, non-rival ideas can be embedded in as many units desired without further costs that makes the market size effect particularly important. In the next section, we will discuss some empirical evidence on the importance of the market size effect.

Nevertheless, the non-rivalry of ideas does not make ideas or innovations *pure public goods*. Recall that pure public goods are both non-rival and non-excludable. While some discoveries may be, by their nature, non-excludable (for example, the "discovery" that providing excessively high-powered incentives to CEOs in the form of stock options will lead to counterproductive incentives and cheating), most discoveries can be made partly excludable by *patenting*. An important aspect of the models of technological progress will be whether and how discoveries are protected from rivals. For this reason, intellectual property rights protection and patent policy often play an important role in models of technological progress.

12.2. Science and Profits

Another major question for the economic analysis of technological change is whether innovation is mainly determined by scientific constraints and stimulated by scientific breakthroughs in particular fields, or whether it is driven by profit motives. Historians and economists typically give different answers to this question. Many historical accounts of technological change come down on the side of the "science-driven" view, emphasizing the autonomous progress of science, and how important breakthroughs—perhaps macro innovations discussed above—have taken place as scientists build on each other's work, with little emphasis on profit opportunities. For example, in his *History of Modern Computing*, Ceruzzi emphasizes the importance of a number of notable scientific discoveries and the role played by certain talented individuals, such as John von Neumann, J. Presper Eckert, John Maucly, John Backus, Kenneth H. Olsen, Harlan Anderson and those taking part in the Project Whirlwind at MIT, rather than profit motives and the potential market for computers. He points out, for example, how important developments took place despite the belief of many important figures in the development of the computer, such as Howard Aiken, that there would not be more than a handful of personal computers in the United States (2000, p. 13). Many economic historians, including Rosenberg (1974) and Sherer (1984), similarly argue that a key determinant of innovation in a particular field is the largely-exogenous growth of scientific and engineering knowledge in that field.

In contrast, most economists believe that profit opportunities play a much more important role and that the demand for innovation is key to understanding the process of technological

change. John Stuart Mill provides an early and clear statement of this view in his *Principles of Political Economy*, when he writes:

> "The labor of Watt in contriving the steam-engine was as essential a part of production as that of the mechanics who build or the engineers who work the instrument; and was undergone, no less than theirs, in the prospect of a renumeration from the producers." (1890, p. 68, also quoted in Schmookler, 1966, p. 210).

In fact, profits were very much in the minds of James Watt and his business partner, Matthew Bolton as the previous quote illustrates. James Watt also praised the patent system for the same reasons, arguing that: "...an engineer's life without patent was not worthwhile" (quoted in Mokyr, 1990, p. 248). The view that profit opportunities are the primary determinant of innovation and invention is articulated by Griliches and Schmookler (1963), and most forcefully and eloquently by Schmookler's seminal study, *Invention and Economic Growth*. Schmookler writes:

> "...invention is largely an economic activity which, like other economic activities, is pursued for gain." (1966, p. 206)

Moreover, Schmookler argues against the importance of major breakthroughs in science on economic innovation. He concludes his analysis of innovations in petroleum refining, papermaking, railroading, and farming by stating that there is no evidence that past breakthroughs have been the major factor in new innovations. In particular, he argues: "Instead, in hundreds of cases the stimulus was the recognition of a costly problem to be solved or a potentially profitable opportunity to be seized..." (1966, p. 199).

If potential profits are a main driver of technological change, then *the market size* that will be commanded by new technologies or products will be a key determinant of innovations. A greater market size increases profits and makes innovation and invention more desirable. To emphasize this point, Schmookler called two of his chapters "*The amount of invention is governed by the extent of the market.*" Schmokler's argument is most clearly illustrated by the example of the horseshoe. He documented that there was a very high rate of innovation throughout the late nineteenth and early twentieth centuries in the ancient technology of horseshoe making, and no tendency for inventors to run out of additional improvements. On the contrary, inventions and patents increased because demand for horseshoes was high. Innovations came to an end only when "the steam traction engine and, later, internal combustion engine began to displace the horse..." (1966, p. 93). The classic study by Griliches (1957) on the spread of hybrid seed corn in the US agriculture also provides support for the view that technological change and technology adoption are closely linked to profitability and market size.

A variety of more recent papers also reach similar conclusions. An interesting paper by Newell, Jaffee and Stavins (1999) shows that between 1960 and 1980, the typical air-conditioner sold at Sears became significantly cheaper, but not much more energy-efficient.

On the other hand, between 1980 and 1990, there was little change in costs, but air-conditioners became much more energy-efficient, which, they argue, was a response to higher energy prices. This seems to be a clear example of the pace and the type of innovation responding to profit incentives. In a related study, Popp (2002) finds a strong positive correlation between patents for energy-saving technologies and energy prices, and thus confirms the overall picture resulting from the Newell, Jaffee and Stavins study.

Evidence from the pharmaceutical industry also illustrates the importance of profit incentives and especially of the market size on the rate of innovation. Finkelstein (2003) exploits three different policy changes affecting the profitability of developing new vaccines against 6 infectious diseases: the 1991 Center for Disease Control recommendation that all infants be vaccinated against hepatitis B, the 1993 decision of Medicare to cover the costs of influenza vaccinations, and the 1986 introduction of funds to insure vaccine manufactures against product liability lawsuits for certain kinds of vaccines. She finds that increases in vaccine profitability resulting from these policy changes are associated with a significant increase in the number of clinical trials to develop new vaccines against the relevant diseases. Acemoglu and Linn (2004) look at demographic-driven exogenous changes in the market size for drugs of different types and find a significant response in the rate of innovation to these changes in market sizes.

Overall, the evidence suggests that the market size is a major determinant of innovation incentives and the amount and type of technological change. This evidence motivates the types of models presented below, where technological change will be an economic activity and will respond to profit incentives rather than simply being driven by exogenous scientific processes.

12.3. The Value of Innovation in Partial Equilibrium

Let us now turn to the analysis of the value of innovation and R&D to a firm. The equilibrium value of innovation and the difference between this private value and the social value (defined as the value to a social planner internalizing externalities) will play a central role in our analysis below. All of the growth models presented so far are dynamic general equilibrium models. In fact, as emphasized at the beginning, economic growth is a process we can only understand in the context of general equilibrium analysis. Nevertheless, it is useful to start our investigation of the value of innovation in partial equilibrium, where much of the industrial organization literature starts.

Throughout this section, I consider a single industry. Firms in this industry have access to an existing technology that enables firms to produce one unit of the product at the marginal cost $\psi > 0$. The demand side of the industry is modeled with a demand curve

$$Q = D(p),$$

where p is the price of the product and Q is the demand at this price. Throughout I assume that $D(p)$ is strictly decreasing, differentiable and satisfies the following conditions:

$$D(\psi) > 0 \text{ and } \varepsilon_D(p) \equiv -\frac{pD'(p)}{D(p)} \in (1, \infty).$$

The first ensures that there is positive demand when price is equal to marginal cost, and the second ensures that the elasticity of demand, $\varepsilon_D(p)$, is always greater than 1, so that when we consider monopoly pricing, there will exist a well-defined profit-maximizing price. Moreover, this elasticity is less than infinity, so that the monopoly price will be above marginal cost.

Throughout this chapter and whenever there are economies with monopolistic competition, oligopolies or potential monopolies, equilibrium refers to Nash equilibrium or subgame perfect Nash equilibrium (when the game in question is dynamic).

12.3.1. No Innovation with Pure Competition. Suppose first that there is a large number of firms, say N firms, with access to the existing technology. Now imagine that one of these firms, say firm 1, also has access to a research technology for generating a process innovation. In particular, let us simplify the discussion and suppose that there is no uncertainty in research, and if the firm incurs a cost $\mu > 0$, it can innovate and reduce the marginal cost of production to ψ/λ, where $\lambda > 1$. Let us suppose that this innovation is non-rival and is also non-excludable, either because it is not patentable or because the patent system does not exist.

Let us now analyze the incentives of this firm in undertaking this innovation. First, in the equilibrium without the innovation, the presence of a large number of N firms, all with the same technology with marginal cost ψ, implies that the equilibrium price will be $p^N = \psi$, where the superscript N denotes "no innovation". The total quantity demanded will be $D(\psi) > 0$ and can be distributed among the N firms in any arbitrary fashion. Since price is equal to marginal cost, the profits of firm 1 in this equilibrium will be

$$\begin{aligned} \pi_1^N &= \left(p^N - \psi\right) q_1^N \\ &= 0, \end{aligned}$$

where q_1^N denotes the amount supplied by this firm.

Now imagine that firm 1 innovates, but because of non-excludability, the innovation can be used by all the other firms in the industry. The same reasoning implies that the equilibrium price will be $p^I = \lambda^{-1}\psi$, and total quantity supplied by all the firms will equal $D\left(\lambda^{-1}\psi\right) > D(\psi)$. Then, the net profits of firm 1 following innovation will be

$$\begin{aligned} \pi_1^I &= \left(p^I - \lambda^{-1}\psi\right) q_1^I - \mu \\ &= -\mu < 0. \end{aligned}$$

Therefore, if it undertakes the innovation, firm 1 will lose money. The reason for this is simple. The firm incurs the cost of innovation, μ, but because the knowledge generated by the innovation is non-excludable, it is unable to *appropriate* any of the gains from innovation.

This simple example underlies a claim dating back to Schumpeter that pure competition will not generate innovation.

Clearly, this outcome is potentially very inefficient. For example, μ could be arbitrarily small (but still positive), while λ, the gain from innovation, can be arbitrarily large. Nevertheless, the equilibrium would still involve no innovation. For future reference, let us calculate the social value of innovation, which is the additional gain resulting from innovation. A natural measure of social value is in the sum of the consumer and producer surpluses generated from the innovation. Presuming that after innovation, the good will be priced at marginal cost, this social value is

$$\begin{aligned} \mathcal{S}^I &= \int_{\lambda^{-1}\psi}^{\psi} D(p)\,dp - \mu \\ &= \int_{\lambda^{-1}\psi}^{\psi} [D(p) - D(\psi)]\,dp + D(\psi)\,\lambda^{-1}(\lambda-1)\,\psi - \mu. \end{aligned} \tag{12.1}$$

The first term in the second line is the increase in consumer surplus because of the expansion of output as the price falls from ψ to $\lambda^{-1}\psi$ (recall that price is equal to marginal cost in this social planner's allocation). The second term is the savings in costs for already produced units; in particular, there is a saving of $\lambda^{-1}(\lambda-1)\,\psi$ on $D(\psi)$ units. Finally, the last term is the cost of innovation. Depending on the shape of the function $D(p)$, the values of λ and μ, this social value of innovation can be quite large.

12.3.2. Some Caveats. The above example illustrates the problem of innovation under pure competition in a very sharp way. The main problem is the inability of the innovator to exclude others from using this innovation. One way of ensuring such excludability is via the protection of intellectual property rights or a patent system, which will create *ex post monopoly power* for the innovator. This type of intellectual property right protection is present in most countries and will play an important role in many of the models we study below.

Before embarking on an analysis of the implications of ex post monopoly power of innovators, there are a number of caveats. First, even without patents, "trade secrecy" may be sufficient to provide some incentives for innovation. Second, firms may engage in innovations that are only appropriate for their own firm, making their innovations de facto excludable. For example, imagine that at the same cost, the firm can develop a new technology that reduces the marginal cost of production by only $\lambda' < \lambda$. But this technology is *specific* to the needs and competencies of the current firm and cannot be used by any other (or alternatively, λ/λ' is the "proportional" cost of making the innovation "excludable"). The adoption of this technology may be profitable for the firm, since the specificity of the innovation firm acts exactly like patent protection (see next subsection and also Exercise 12.5). Therefore, some types of innovations, in particular those protected by trade secrecy, can be undertaken under pure competition.

Finally, a number of authors have recently argued that innovations in competitive markets are possible. One strand of the literature shows that competitive growth may originate because firms are able to replicate new technologies (for example, copy software or compact discs) and sell them to competitors during a certain interval of time before being imitated by others (e.g., Walde, 2002, and Boldrin and Levine, 2003). Another strand incorporates diminishing returns at the level, which creates potential profits for firms so that they can cover the fixed costs of investment in new technologies (e.g., Hellwig and Irmen, 2001). This recent work on competitive growth constitutes a very promising direction, though existing models generate innovations and sustained growth in competitive equilibria only under somewhat special assumptions (such as fully replicable technologies or strong externalities from past innovations).

12.3.3. Innovation and Ex Post Monopoly. Let us now return to the same environment as above, and suppose that if firm 1 undertakes a successful innovation it can obtain a fully-enforced patent. Once this happens, firm 1 will have better technology than the rest of the firms, and will possess *ex post* monopoly power. This monopoly power will enable the firm to earn profits from the innovation, potentially encouraging its research activity in the first place. This is the basis of the claim by Schumpeter, Arrow, Romer and others that there is an intimate link between ex post monopoly power and innovation.

Let us now analyze this situation in a little more detail. It is useful to separate two cases:

(1) *Drastic innovation*: a drastic innovation corresponds to a sufficiently high value of λ such that firm 1 becomes an effective monopolist after the innovation. To determine which values of λ will lead to a situation of this sort, let us first suppose that firm 1 does indeed act like a monopolist. This implies that it will choose its price to maximize

$$\pi_1^I = D(p)\left(p - \lambda^{-1}\psi\right).$$

The first-order condition of this maximization is

$$D'(p)\left(p - \lambda^{-1}\psi\right) + D(p) = 0.$$

Clearly the solution to this equation gives the standard monopoly pricing formula (see Exercise 12.1):

$$p^M \equiv \frac{\lambda^{-1}\psi}{1 - \varepsilon_D(p)^{-1}}. \tag{12.2}$$

We say that the innovation is *drastic* if $p^M \leq \psi$. It is clear that this will be the case when

$$\lambda \geq \lambda^* \equiv \frac{1}{1 - \varepsilon_D(p)^{-1}}.$$

When the innovation is drastic, firm 1 can set its unconstrained monopoly price, p^M, and capture the entire market.

(2) *Limit pricing*: when the innovation is not drastic, so that $p^M > \psi$ or alternatively, when $\lambda < \lambda^*$, the equilibrium will involve limit pricing, where firm 1 sets the price

$$p_1 = \psi,$$

so as to make sure that it still captures the entire market (since in this case if it were to set $p_1 = p^M$, other firms can profitably undercut firm 1). This type of limit pricing arises in many situations. In this case, limit pricing results from process innovations by some firms that now have access to a better technology than their rivals. Alternatively, it can also arise when *a fringe* of potential entrants can imitate the technology of a firm (either at some cost or with lower efficiency) and the firm may be forced to set a limit price in order to prevent the fringe from stealing its customers.

The next proposition summarizes this discussion. Recall that the difference between the drastic innovation and limit pricing cases corresponds to different values of λ—i.e., it depends on whether or not λ is greater than λ^*.

PROPOSITION 12.1. *Consider the above-described industry. Suppose that firm 1 undertakes an innovation reducing marginal cost of production from ψ to $\lambda^{-1}\psi$. If $p^M \leq \psi$ (or if $\lambda \geq \lambda^*$), then it sets the unconstrained monopoly price $p_1 = p^M$ and makes profits*

$$\hat{\pi}_1^I = D\left(p^M\right)\left(p^M - \lambda^{-1}\psi\right) - \mu. \tag{12.3}$$

If $p^M > \psi$ (if $\lambda < \lambda^*$), then firm 1 sets the limit price $p_1 = \psi$ and makes profits

$$\pi_1^I = D\left(\psi\right)\lambda^{-1}\left(\lambda - 1\right)\psi - \mu < \hat{\pi}_1^I. \tag{12.4}$$

PROOF. The proof of this proposition involves solving for the equilibrium of an asymmetric cost Bertrand competition game. While this is standard, it is useful to repeat it, especially to see why in equilibrium, all demand must be met by the low cost firm. Exercise 12.2 asks you to work through the steps of the proof. □

The fact that $\hat{\pi}_1^I > \pi_1^I$ is intuitive, since the former refers to the case where λ is greater than λ^*, whereas in the latter, firm 1 has a sufficiently low λ that it is forced to charge a price lower than the profit-maximizing monopoly price because of the competition by the remaining firms (still producing at marginal cost ψ).

It can also be easily verified that both $\hat{\pi}_1^I$ and π_1^I can be strictly positive, so that with ex post monopoly innovation becomes possible. This corresponds to a situation in which we start with pure competition, but one of the firms undertakes an innovation in order to *escape competition* and gains ex post monopoly power. The fact that the ex post monopoly power is important for providing incentives to undertake innovations is consistent with Schumpeter's emphasis on the role of monopoly in generating innovations.

Now returning to the discussion in the previous subsection, we can also see that trade secrecy or innovations that are specific only for the needs of the firm in question will act in the same way as ex post patent protection in encouraging innovation (see Exercise 12.5).

Note also that the expressions for $\hat{\pi}_1^I$ and π_1^I in this proposition also give the value of innovation to firm 1, since without innovation, it would make zero profits. Given this observation, let us contrast the value of innovation for firm 1 in these two regimes with the social value of innovation, which is still given by (12.1). Moreover, we can also compare social values in the equilibrium in which innovation is undertaken by firm 1 (who will charge the profit-maximizing price) to the full social value of innovation in (12.1), which applies when the product is priced at marginal cost. The equilibrium social surplus in the regimes with monopoly and limit pricing (again corresponding to the cases in which λ is greater than or less than λ^*) can be computed as

$$\begin{aligned} \widehat{\mathcal{S}}_1^I &= D\left(p^M\right)\left(p^M - \lambda^{-1}\psi\right) + \int_{p^M}^{\psi} D\left(p\right) dp - \mu, \text{ and} \\ \mathcal{S}_1^I &= D\left(\psi\right)\lambda^{-1}\left(\lambda - 1\right)\psi - \mu. \end{aligned} \tag{12.5}$$

We then have the following result:

PROPOSITION 12.2. *We have that*

$$\begin{aligned} \pi_1^I &< \hat{\pi}_1^I < \mathcal{S}^I. \\ \mathcal{S}_1^I &< \widehat{\mathcal{S}}_1^I < \mathcal{S}^I. \end{aligned}$$

PROOF. See Exercise 12.3. □

This proposition states that the social value of innovation is always greater than the private value in two senses. First, the first line states that a social planner interested in maximizing consumer and producer surplus will always be more willing to adopt an innovation, because of *an appropriability effect*; the firm, even if it has ex post monopoly rights, will be able to appropriate only a portion of the gain in consumer surplus created by the better technology. Second, the second line implies that even conditional on innovation, the gain in social surplus is always less in the equilibrium supported by ex post monopoly than the gain that the social planner could have achieved (by also controlling prices). Therefore, even though ex post monopoly power (for example, generated by patents) can induce innovation, the incentives for innovation and the equilibrium allocations following an innovation are still inefficient. Note also that $\widehat{\mathcal{S}}_1^I$ might be negative, so that a potentially productivity-enhancing process innovation can reduce social surplus because of the cost of innovation, μ. This would then correspond to *excessive innovation*. However, it can be shown that if $\hat{\pi}_1^I > 0$, then $\widehat{\mathcal{S}}_1^I > 0$, which implies that excessive innovation is not possible in this competitive environment (see Exercise 12.4). This will contrast with the results in the next subsection.

12.3.4. The Value of Innovation to a Monopolist: The Replacement Effect. Let us now analyze the same industry as in the previous subsection, but assuming that firm 1 is already an unconstrained monopolist with the existing technology. Then, with the existing technology, this firm would set the monopoly price of

$$\hat{p}^M \equiv \frac{\psi}{1 - \varepsilon_D\left(p\right)^{-1}}$$

and make profits equal to

$$\hat{\pi}_1^N = D\left(\hat{p}^M\right)\left(\hat{p}^M - \psi\right). \tag{12.6}$$

If it undertakes the innovation, it will reduce its marginal cost to $\lambda^{-1}\psi$ and still remain an unconstrained monopolist. Therefore, its profits will be given by $\hat{\pi}_1^I$ as in (12.3), with the monopoly price p^M given by (12.2). Now the value of innovation to the monopolist is

$$\begin{aligned}\Delta\hat{\pi}_1^I &= \hat{\pi}_1^I - \hat{\pi}_1^N \\ &= D\left(p^M\right)\left(p^M - \lambda^{-1}\psi\right) - D\left(\hat{p}^M\right)\left(\hat{p}^M - \psi\right) - \mu,\end{aligned}$$

where $\hat{\pi}_1^I$ is given by (12.3) and $\hat{\pi}_1^N$ by (12.6).

Proposition 12.3. *We have that*

$$\Delta\hat{\pi}_1^I < \pi_1^I < \hat{\pi}_1^I,$$

so that a monopolist always has lower incentives to undertake innovation than a competitive firm.

Proof. See Exercise 12.6. □

This result, which was first pointed out in Arrow's (1962) seminal paper, is referred to as the *replacement effect*. The terminology reflects the intuition for the result; the monopolist has lower incentives to undertake innovation than the firm in a competitive industry because the innovation will replace its own already existing profits. In contrast, a competitive firm would be making zero profits and thus would have no profits to replace.

An immediate and perhaps more useful corollary of this proposition is the following:

Corollary 12.1. *A potential entrant will have stronger incentives to undertake an innovation than an incumbent monopolist.*

The potential entrant would make zero profits without the innovation. If it undertakes the innovation it becomes the ex post monopolist and makes profits equal to π_1^I or $\hat{\pi}_1^I$ (depending on whether there will be limited pricing) Both of these are greater than the additional profits that the incumbent would make by innovating, $\Delta\hat{\pi}_1^I$. This is a direct consequence of the replacement effect; while the incumbent would be replacing its own profit-making technology, the entrant would be replacing the incumbent. The replacement effect and this corollary imply that in many models entrants have stronger incentives to invest in R&D than incumbents.

The observation that entrants will often be the engine of process innovations takes us to the realm of Schumpeterian models. Joseph Schumpeter characterized the process of economic growth as one of *creative destruction*, meaning a process in which economic progress goes hand-in-hand with the destruction of some existing productive units. Put differently, innovation is driven by the prospect of monopoly profits. Because of the replacement effect, it will be entrants, not incumbents, that undertake greater R&D towards inventing and implementing process innovations. Consequently, innovations will displace incumbents and destroy their rents. According to Schumpeter, this process of creative destruction is the

essence of the capitalist economic system. Chapter 14 will show that the process of creative destruction can be the engine of economic growth as well.

In addition to providing an interesting description of the process of economic growth and highlighting the importance of the market structure, the process of creative destruction is important because it also brings political economy interactions to the fore of the question of economic growth. If economic growth takes place via creative destruction, it will create losers, in particular, the incumbents who are currently enjoying profits and rents. Since incumbents may be politically powerful as well, this implies that many economic systems will create powerful barriers against the process of economic growth. Political economy of growth is partly about understanding the opposition of certain firms and individuals to technological progress and studying whether this opposition will be successful.

There is another, perhaps more surprising, implication of the analysis in this subsection. This relates to *the business stealing effect*, which is closely related to the replacement effect. The entrant, by replacing the incumbent, is also stealing the business—profits—of the incumbent. The above discussion suggests that this business stealing effect helps closing the gap between the private and the social values of innovation. It is also possible, however, for the business stealing effect to lead to *excessive innovation* by the entrant. To see the possibility of excessive innovation, let us first look at the total surplus gain from an innovation starting with the monopolist. Suppose to simplify the discussion that the innovation in question is drastic, so that if the entrant undertakes this innovation, it can set the unconstrained monopoly price p^M as given by (12.2) above. Therefore, the social value of innovation is $\widehat{\mathcal{S}}_1^I$ as given by (12.5).

Proposition 12.4. *It is possible that*

$$\widehat{\mathcal{S}}_1^I < \hat{\pi}_1^I,$$

so that the entrant has excessive incentives to innovate.

Proof. See Exercise 12.8. □

Intuitively, the social planner values the profits made by the monopolist, since these are part of the "producer surplus". In contrast, the entrant only values the profits that it will make if it undertakes the innovation. This is the essence of the business stealing effect and creates the possibility of excessive innovations. This result is important because it points out that, in general, it is not clear whether the equilibrium will involve too little or too much innovation. Whether or not it does so depends on how strong the business stealing effect is relative to the appropriability effect discussed above.

12.4. The Dixit-Stiglitz Model and "Aggregate Demand Externalities"

The analysis in the previous section focused on the private and the social values of innovations in a partial equilibrium setting. A large part of growth theory is about general equilibrium models of innovation. This requires a tractable model of industry equilibrium, which

can then be embedded in a general equilibrium framework. The most widely-used model of industry equilibrium is the model developed by Dixit and Stiglitz (1977) and Spence (1976), which captures many of the key features of Chamberlin's (1933) discussion of monopolistic competition. Chamberlin (1933) suggested that a good approximation to the market structure of many industries is one in which each firm faces a downward sloping demand curve (thus has some degree of monopoly power), but there is also free entry into the industry, so that each firm (or at the very least, the marginal firm) makes zero profits.

The distinguishing feature of the Dixit-Stiglitz-Spence model (or Dixit-Stiglitz model for short) is that it allows us to specify a structure of preferences that leads to constant monopoly markups. This turns out to be a very convenient feature in many growth models, though it also implies that this model may not be particularly well suited to situations in which market structure and competition affect monopoly markups. In this section, I present a number of variants of the Dixit Stiglitz model and emphasize its advantages and shortcomings.

12.4.1. The Dixit-Stiglitz Model with a Finite Number of Products. Consider a static economy that admits a representative household with preferences given by

$$U(c_1, ..., c_N, y) = u(C, y), \tag{12.7}$$

where

$$C \equiv \left(\sum_{i=1}^{N} c_i^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}} \tag{12.8}$$

is a consumption index of N differentiated "varieties," $c_1, ..., c_N$, of a particular good and y stands for a generic good, representing all other consumption. The function $u(\cdot, \cdot)$ is strictly increasing and differentiable in both of its arguments and is jointly strictly concave. The parameter ε in (12.8) represents the *elasticity of substitution* between the differentiated varieties and we assume that $\varepsilon > 1$. The key feature of (12.8) is that it features *love-for-variety*, meaning that the greater is the number of differentiated varieties that the individual consumes, the higher is his utility. The aggregator over the different consumption varieties in (12.8) will appear in many different models of technological change and economic growth in the remainder of the book. I refer to it as a *Dixit-Stiglitz aggregator* or more often, as a *CES aggregator* (where CES stands for constant elasticity of substitution).

More specifically, consider the case in which

$$c_1 = ... = c_N = \frac{\bar{C}}{N},$$

so that the individual purchases a total of $\bar{C}$ units of differentiated varieties, distributed equally across all N varieties. Substituting this into (12.7) and (12.8),

$$U\left(\frac{\bar{C}}{N}, ..., \frac{\bar{C}}{N}, y\right) = u\left(N^{\frac{1}{\varepsilon-1}}\bar{C}, y\right),$$

which is strictly increasing in N (since $\varepsilon > 1$) and implies that for a fixed total $\bar{C}$ units of differentiated commodities, the larger is the number of varieties over which this total number

of units are distributed, the higher is the utility of the individual. This is the essence of the love-for-variety utility function. What makes this utility function convenient is not only this feature, but also the fact that individual demands take a very simple isoelastic form. To derive the demand for individual varieties, let us normalize the price of the y good to 1 and denote the price of variety i by p_i and the total money income of the individual by m. Then, the budget constraint of the individual takes the form

$$\sum_{i=1}^{N} p_i c_i + y \leq m. \tag{12.9}$$

The maximization of (12.7) subject to (12.9) implies the following first-order condition between varieties:

$$\left(\frac{c_i}{c_{i'}}\right)^{-\frac{1}{\varepsilon}} = \frac{p_i}{p_{i'}} \text{ for any } i,i'.$$

To write this first-order condition in a more convenient form, let P denote the price index corresponding to the consumption index C. Then, combining the first-order conditions,

$$\left(\frac{c_i}{C}\right)^{-\frac{1}{\varepsilon}} = \frac{p_i}{P} \text{ for } i = 1, ..., N. \tag{12.10}$$

This first-order condition for the consumption index immediately implies that (see Exercise 12.10):

$$P \equiv \left(\sum_{i=1}^{N} p_i^{1-\varepsilon}\right)^{\frac{1}{1-\varepsilon}}. \tag{12.11}$$

Since P is the price index corresponding to the consumption index C, it is typically referred to as the *ideal price index*. In many circumstances, it will be convenient to choose this ideal price index as the numeraire. Note, however, that we cannot set this as the price index in this particular instance, since the budget constraint is already in terms of money income, m, and the price of good y is normalized to 1. The choice between C and y is straightforward in this case and boils down to the maximization of the utility function $u(C, y)$ subject to the budget constraint

$$PC + y \leq m,$$

which combines (12.10) and (12.11) with (12.9), in order to obtain a budget constraint expressed in terms of C and y. Now this maximization yields the following intuitive first-order condition:

$$\frac{\partial u(C, y)/\partial y}{\partial u(C, y)/\partial C} = \frac{1}{P},$$

which assumes that the solution is interior, an assumption I maintain throughout this section to simplify the discussion. The strict joint concavity of u, combined with the budget constraint, implies that this first-order condition can be expressed as

$$\begin{aligned} y &= g(P, m) \\ C &= \frac{m - g(P, m)}{P}, \end{aligned} \tag{12.12}$$

for some function $g(\cdot,\cdot)$.

Next, let us consider the production of the varieties. Suppose that each variety can only be produced by a single firm, who is thus an effective monopolist for this particular commodity. Also assume that all monopolists are owned by the representative household and maximize profits.

Recall that the marginal cost of producing each of these varieties is constant and equal to ψ. Let us first write down the profit maximization problem of one of these monopolists:

$$\max_{p_i} \left(\left(\frac{p_i}{P} \right)^{-\varepsilon} C \right) (p_i - \psi), \tag{12.13}$$

where the term in the first parentheses is c_i (recall (12.10)) and the second is the difference between price and marginal cost. The complication in this problem comes from the fact that P and C are potentially functions of p_i. However, for N sufficiently large, the effect of p_i on these can be ignored and the solution to this maximization problem becomes very simple (see Exercise 12.11). This enables us to derive the optimal price in the form of a constant markup over marginal cost:

$$p_i = p = \frac{\varepsilon}{\varepsilon - 1}\psi \text{ for each } i = 1, ..., N. \tag{12.14}$$

This result follows because when the effect of firm i's price choice on P and C is ignored, the demand function facing the firm, (12.10), is isoelastic with an elasticity $\varepsilon > 1$. Since each firm charges the same price, the ideal price index P can be computed as

$$P = N^{-\frac{1}{\varepsilon-1}} \frac{\varepsilon}{\varepsilon - 1}\psi. \tag{12.15}$$

Using this expression the profits for each firm are obtained as

$$\pi_i = \pi = N^{-\frac{\varepsilon}{\varepsilon-1}} C \frac{1}{\varepsilon - 1}\psi \text{ for each } i = 1, ..., N.$$

Profits are decreasing in the price elasticity for the usual reasons. In addition, profits are increasing in C because this is the total amount of expenditure on these differentiated goods, and they are decreasing in N, since for given C a larger number of varieties means less spending on each variety.

Despite this last effect, the total impact of N on profits can be positive. To see this, let us substitute for P from (12.15) to obtain

$$C = N^{\frac{1}{\varepsilon-1}} \frac{\varepsilon - 1}{\varepsilon\psi} \left(m - g\left(N^{-\frac{1}{\varepsilon-1}} \frac{\varepsilon}{\varepsilon - 1}\psi, m \right) \right)$$

and

$$\pi = \frac{1}{\varepsilon N} \left(m - g\left(N^{-\frac{1}{\varepsilon-1}} \frac{\varepsilon}{\varepsilon - 1}\psi, m \right) \right).$$

It can be verified that depending on the form of the $g(\cdot)$ function, which in turn depends on the shape of the utility function u in (12.7), profits can be increasing in the number of varieties (see Exercise 12.12). This may at first appear somewhat surprising: typically, we expect a greater number of competitors to reduce profits. But the love-for-variety effect embedded in the Dixit-Stiglitz preferences creates a countervailing effect, which is often referred to as

an *aggregate demand externality* in the macroeconomics literature. The basic idea is that a higher N raises the utility from consuming each of the varieties because of the love-for-variety effect. The impact of the entry of a particular variety (or the impact of the increase in the production of a particular variety) on the demand for other varieties is a pecuniary externality. This pecuniary externality will play an important role in many of the models of endogenous technological change and we will encounter it again in models of poverty traps in Chapter 21.

12.4.2. The Dixit-Stiglitz Model with a Continuum of Products. As discussed in the last subsection and analyzed further in Exercise 12.12, when N is finite, the equilibrium in which each firm charges the price given by (12.14) may be viewed as an approximation (where each firm only has a small effect on the ideal price index and thus ignores this effect). An alternative modeling assumption would be to assume that there is a continuum of varieties. When there is a continuum of varieties, (12.14) is no longer an approximation. Moreover, such a model will be more tractable because the number of firms, N, need not be a natural number. For this reason, the version of the Dixit-Stiglitz model with a continuum of products is often used in the literature and will also be used in the rest of this book.

This version of the model is very similar to the one discussed in the previous subsection, except that the utility function of the representative household now takes the form

$$U\left(\left[c_i\right]_{i=0}^{N}, y\right) = u\left(C, y\right),$$

where

$$C \equiv \left(\int_0^N c_i^{\frac{\varepsilon-1}{\varepsilon}} di\right)^{\frac{\varepsilon}{\varepsilon-1}}$$

and N denotes the measure of varieties. The budget constraint facing the representative household is

$$\int_0^N p_i c_i di + y \leq m.$$

An identical analysis leads to utility maximizing decisions given by (12.10) and to the ideal price index

$$P = \left(\int_0^N p_i^{1-\varepsilon} di\right)^{\frac{1}{1-\varepsilon}}.$$

Using the definition of the ideal price index and (12.10), the budget constraint is

$$PC + y \leq m.$$

Equation (12.12) then determines y and C. Since the supplier of each variety is infinitesimal, their price has no effect on P and C. Consequently, the profit-maximizing pricing decision in (12.14) applies exactly, and each firm has profits given by

$$\pi = \frac{1}{\varepsilon N}\psi\left(m - g\left(N^{-\frac{1}{\varepsilon-1}}\frac{\varepsilon}{\varepsilon-1}\psi, m\right)\right),$$

where $g\left(\cdot\right)$ is defined as (12.12) in the previous subsection.

Now using this expression, the entry margin can also be endogenized. Imagine, for example, that there is an infinite number of potential different varieties, and a particular firm can adopt one of these varieties at some fixed cost $\mu > 0$ and enter the market. Consequently, as in the Chamberlin's (1933) model of monopolistic competition, in equilibrium all varieties will make zero profits because of free entry. This implies that the following zero-profit condition has to hold for all entrants and thus for all varieties:

$$\frac{1}{\varepsilon N}\psi\left(m - g\left(N^{-\frac{1}{\varepsilon-1}}\frac{\varepsilon}{\varepsilon-1}\psi, m\right)\right) = \mu. \tag{12.16}$$

As discussed in the next chapter, there is an intimate link between entry by new products (firms) and technological change. Leaving a detailed discussion of this connection to the next chapter, here we can ask a simpler question: do the aggregate demand externalities imply that there is too little entry in a model of this sort? The answer is not necessarily. While the aggregate demand externalities imply that firms do not take into account the positive benefits their entry creates on other firms, the business stealing effect identified above is still present and implies that entry may also reduce the demand for existing products. Thus, in general, whether there is too little or too much entry in models of product differentiation depends on the details of the model and the values of the parameters (see Exercise 12.13).

12.4.3. Objectives of Monopolistic Firms. It is also useful to briefly discuss the objectives of monopolistically competitive firms in the Dixit-Stiglitz model (and related models). Throughout, I follow the industrial organization and the growth literatures and assume that all firms maximize profits, even when they are owned by a "representative household". One may object to this assumption, noting that, since this is a monopolistically competitive economy where the First Welfare Theorem (e.g., Theorem 5.5) does not hold, the representative household might be made better-off if firms pursued a non-profit maximizing strategy. However, profit maximization is still the right objective function for firms. This is because an allocation in which firms do not maximize profits (instead act in the way that a social planner would like them to act) cannot be an equilibrium. To see this, note that the representative household is itself a price taker—for example, it represents a large number of identical price-taking households. If some firms did not maximize profits, then the households would refuse to hold the stocks of these firms in their portfolios and there would be entry by other profit-maximizing firms instead. Thus, as long as the representative household or the set of households on the consumer side act as price-takers (as has been assumed to be the case throughout), profit maximization is the only consistent strategy for the monopolistically competitive firms.

The only caveat to this arises from a different type of deviation on the production side. In particular, a single firm may buy all of the monopolistically competitive firms and act as the single producer in the economy. This firm might then ensure an allocation that makes consumers better-off relative to the equilibrium allocation considered here. Nevertheless, I ignore this type of deviation for two reasons. First, as usual we are taking the market structure as given, and the market structure here is monopolistic competition not pure monopoly. A

single firm owning all production units would correspond to an entirely different market structure, with much less realism and relevance to the issues studied here. Second, in a related model Acemoglu and Zilibotti (1997) show that a single firm owning all production units would not be an equilibrium either because the possibility of free entry would encourage the entry of profit-maximizing firms at the margin, disrupting the equilibrium with the single producer. This is discussed further at the end of Chapter 17. Given these considerations, throughout the book I assume that firms are profit maximizing.

12.4.4. Limit Prices in the Dixit-Stiglitz Model. We have already encountered how limit prices can arise in the previous section, when process innovations are non-drastic relative to the existing technology. Another reason why limit prices can arise is because of the presence of a "competitive" fringe of firms that can imitate the technology of monopolists. This type of competitive pressure from the fringe of firms is straightforward to incorporate into the Dixit-Stiglitz model and will be useful in later chapters as a way of parameterizing competitive pressures.

Let us assume that there is a large number of fringe firms that can imitate the technology of the incumbent monopolists. Let us assume that this imitation is equivalent to the production of a similar good and is not protected by patents. It may be reasonable to assume that the imitating firms will be less efficient than those who have invented the variety and produced it for a while. A simple way of capturing this would be to assume that while the monopolist creates a new variety by paying the fixed cost μ and then having access to a technology with the marginal cost of production of ψ, the fringe of firms do not pay any fixed costs, but can only produce with a marginal cost of $\gamma\psi$, where $\gamma > 1$.

Similar to the analysis in the previous section, if $\gamma \geq \varepsilon/(\varepsilon-1)$, then the fringe is sufficiently unproductive that they cannot profitably produce even when the monopolists charge the unconstrained monopoly price given in (12.14). Instead, when $\gamma < \varepsilon/(\varepsilon-1)$, the monopolists will be forced to charge a limit price. The same arguments as in the previous section establish that this limit price must take the form

$$p = \gamma\psi < \frac{\varepsilon}{\varepsilon-1}\psi.$$

It is then straightforward to see that the entry condition that determines the number of varieties in the market will change to

$$N^{-\frac{\varepsilon}{\varepsilon-1}}\gamma\psi\left(m - g\left(N^{-\frac{1}{\varepsilon-1}}\frac{\varepsilon}{\varepsilon-1}\psi\right)\right) = \mu.$$

12.4.5. Limitations. The most important limitation of the Dixit-Stiglitz model is the feature that makes it tractable: the constancy of markups as in eq. (12.14). In particular, the model implies that the markup of each firm is independent of the number of varieties in the market. But this is a very special feature. Most industrial organization models imply that markups over marginal cost are declining in the number of competing products (see, for example, Exercise 12.14). While plausible, this makes endogenous growth models less tractable, because in many classes of models, endogenous technological change will correspond

to a steady increase in the number of products N. If markups decline towards zero as N increases, this would ultimately stop the process of innovation and thus prevent sustained economic growth. The alternative would be to have a model in which some other variable, perhaps capital, simultaneously increases the potential markups that firms can charge. While such models can be developed, they are more difficult than the standard Dixit-Stiglitz setup. For this reason, the literature typically focuses on Dixit-Stiglitz specifications.

12.5. Individual R&D Uncertainty and the Stock Market

The final issue I discuss in this chapter involves uncertainty in the research process. As discussed at the beginning of the chapter, it is reasonable to presume that the output of research will be uncertain. This implies that individual firms undertaking research will face a stochastic revenue stream. When individuals are risk averse, this may imply that there should be a risk premium associated with such stochastic streams of income. This is not necessarily the case, however, when the following three conditions are satisfied:

(1) there are many firms involved in research;
(2) the realization of the uncertainty across firms is independent;
(3) households and firms have access to a "stock market," where each consumer can hold a *balanced portfolio* of various research firms.

In many of the models presented in the next two chapters, firms will face uncertainty (for example, regarding whether their R&D will be successful or how long their monopoly position will last), but the three conditions outlined here will be satisfied. When this is the case, even though each firm's revenue is risky, the balanced portfolio held by the representative household will have deterministic returns. Here I illustrate this with a simple example.

Example 12.2. Suppose that the representative household has a utility function over consumption given by $u(c)$, where $u(\cdot)$ is strictly increasing, differentiable and strictly concave, so that individual is risk averse. Moreover, let us assume that $\lim_{c\to 0} u'(c) = \infty$, so that the marginal utility of consumption at zero is very high. The household starts with an endowment equal to $y > 0$. This endowment can be consumed or it can be invested in a risky R&D project. Imagine that the R&D project is successful with probability p and will have a return equal to $1 + R > 1/p$ per unit of investment. It is unsuccessful with probability $1 - p$, in which case it will have a zero return. When this is the only project available, the household would be facing consumption risk if it invests in this project. In particular, the maximization problem that determines how much it should invest will be a solution to the following expected utility maximization

$$\max_x (1-p)\, u(y-x) + pu(y+Rx).$$

The first-order condition of this problem implies that the optimal amount of investment in the risky research activity will be given by:

$$\frac{u'(y-x)}{u'(y+Rx)} = \frac{pR}{1-p}.$$

The assumption $\lim_{c\to 0} u'(c) = \infty$ implies that $x < y$, thus less than the full endowment of the individual will be invested in the research activity, even though this is a positive expected return project. Intuitively, the household requires a risk premium to bear the consumption risk associated with the risky investment.

Next imagine a situation in which many different firms can invest in similar risky research ventures. Suppose that the success or failure of each project is independent of the others. Imagine that the individual invests an amount x/N in each of N projects. The Strong Law of Large Numbers implies that as $N \to \infty$, a fraction p of these projects will be successful and the remaining fraction $1-p$ will be unsuccessful. Therefore, the household will receive (almost surely) a utility of

$$u\left(y + \left(p\left(1+R\right) - 1\right)x\right).$$

Since $1 + R > 1/p$, this is strictly increasing in x, and implies that the individual would prefer to invest all of its endowment in the risky projects, that is, $x = y$. Therefore, the ability to hold a balanced portfolio of projects with independently disputed returns allows the household to diversify the risks and act in a risk-neutral manner. A similar logic will apply in many of the models presented in the next three chapters; even though individual firms will have stochastic returns, the representative household will hold a balanced portfolio of all the firms in the economy and thus will have risk-neutral preferences in the aggregate. This observation also implies that the objective of each firm will be to maximize expected profits (without a risk premium).

12.6. Taking Stock

This chapter has reviewed a number of conceptual and modeling issues related to the economics of research and development. I introduced the distinction between process and product innovations, macro and micro innovations, and also discussed the concept of the innovation possibilities set and the importance of the non–rivalry of ideas.

We have also seen why ex post monopoly power is important to create incentives for research spending, how incentives to undertake innovations differ between competitive firms and monopolies and how these compare to the social value of innovation. In this context, I emphasized the importance of the appropriability effect, which implies that the private value of innovation often falls short of the social value of innovation, because even with ex post monopoly power an innovating firm will not be able to appropriate the entire consumer surplus created by a better product or a cheaper process. Arrow's replacement effect, in this context, implies that unless they have a cost advantage, incumbent monopolists will have weaker incentives for innovation than the entrants. Despite the appropriability and the replacement effects, the amount of innovation in equilibrium can be excessive, because of another, countervailing force, the business stealing effect, which encourages firms to undertake innovations in order to become the new monopolist and take over ("steal") the monopoly rents. Therefore, whether there is too little or too much innovation in equilibrium depends on the market structure and the parameters of the model.

This chapter has also introduced the Dixit-Stiglitz-Spence (or for short the Dixit-Stiglitz) model, which will play an important role in the analysis of the next few chapters. This model offers a very tractable formalization of Chamberlin's approach to monopolistic competition, where each firm has some monopoly power, but free entry ensures that all firms (or the marginal entrants) make zero profits. The Dixit-Stiglitz model is particularly tractable because the markup charged by monopolists is independent of the number of competing firms. This makes it an ideal model to study endogenous growth, because it will enable innovation to remain profitable even when the number of products (or the number of machines) increases continuously.

12.7. References and Literature

The literature on R&D in industrial organization is vast, and my purpose in this chapter has not been to review this literature, but to highlight the salient features that will be used in the remainder of the book. The reader who is interested in this area can start with Tirole (1990, Chapter 10), which contains an excellent discussion of the contrast between private and the social values of innovation. Tirole also provides an excellent introduction to patent races, which will be discussed in Section 14.4 in Chapter 14. A more up-to-date reference that surveys the recent developments in the economics of innovation is Scotchmer (2005).

The classic reference on the private and social values of innovation is Arrow (1962). Schumpeter (1943) was the first to emphasize the role of monopoly in R&D and innovation. The importance of monopoly power for innovation and the implications of the non-rival nature of ideas are discussed in Romer (1990, 1993) and Jones (2006). Most of the industrial organization literature also emphasizes the importance of ex post monopoly power and patent systems in providing incentives for innovation. See, for example, Scotchmer (2005). This perspective has recently been criticized by Boldrin and Levine (2003).

The idea of creative destruction was also originally developed by Schumpeter. Models of creative destruction in the industrial organization literature include Reinganum (1983, 1985). Similar models in the growth literature are developed in Aghion and Howitt (1992, 1998).

Chamberlin (1933) is the classic reference on monopolistic competition. The Dixit-Stiglitz model is developed in Dixit and Stiglitz (1977) and is also closely related to Spence (1976). This model was first used for an analysis of R&D in Dasgupta and Stiglitz (1979). An excellent exposition of the Dixit-Stiglitz model is provided in Matsuyama (1995). Tirole (1990, Chapter 7) also discusses the Dixit-Stiglitz-Spence model as well as other models of product innovation, including the Salop model, due to Salop (1979), which is presented in Exercise 12.14.

A stimulating general discussion of issues of innovation and the importance of market size and profit incentives is provided in Schmookler (1966). Recent evidence on the effect of market size and profit incentives on innovation is discussed in Popp (2002), Finkelstein (2003) and Acemoglu and Linn (2004).

Mokyr (1990) contains an excellent history of innovation. Freeman (1982) also provides a survey of the qualitative literature on innovation and discusses the different types of innovations.

The rest of this part of the book, like this chapter, will focus on monopolistic environments, where the appropriate equilibrium concept is not the competitive equilibrium, but one that incorporates game-theoretic interactions. Since all games in this book will have complete information, the appropriate notion of equilibrium is the standard Nash Equilibrium concept or when the game is multi-stage or dynamic, it is the Subgame Perfect (Nash) Equilibrium or a refinement of this concept, the Markov Perfect Equilibrium. In these situations, equilibrium always refers to a Nash Equilibrium or a Subgame Perfect Equilibrium and I typically do not add the additional "Nash" qualification. The treatment here presumes that the reader is familiar with these concepts. A quick introduction to the necessary game theory is provided in the Appendix of Tirole (1990), and a more detailed treatment can be found in Fudenberg and Tirole (1994), Myerson (1995) and Osborne and Rubinstein (1994). A brief overview of dynamic (infinite-horizon) games is provided in Appendix Chapter C.

12.8. Exercises

Exercise 12.1. Derive eq. (12.2).

Exercise 12.2. Prove Proposition 12.1. In particular:

(1) Show that even if $p^M = \psi$, the unique (Nash) equilibrium involves $q_1 = D\left(p^M\right)$ and $q_j = 0$ for all $j > 1$. Why is this?
(2) Show that when $p^M > \psi$, any price $p_1 > \psi$ or $p_1 < \psi$ cannot be profit-maximizing. Show that there cannot be an equilibrium in which $p_1 = \psi$ and $q_j > 0$ for some $j > 1$ [Hint: find a profitable deviation for firm 1].
(3) Prove that $\hat{\pi}_1^I > \pi_1^I$.

Exercise 12.3. Derive eq. (12.5). Using these relationships, prove Proposition 12.2.

Exercise 12.4. Prove that if $\hat{\pi}_1^I > 0$, then $\widehat{\mathcal{S}}_1^I > 0$ (where these terms are defined in Proposition 12.2).

Exercise 12.5. Consider the model in Section 12.3, and suppose that there is no patent protection for the innovating firm. The firm can undertake two different types of innovations at the same cost η. The first is a general technological improvement, which can be copied by all firms. It reduces the marginal cost of production to $\lambda^{-1}\psi$. The second is *specific* to the needs of the current firm and cannot be copied by others. It reduces the marginal cost of production by $\lambda' < \lambda$. Show that the firm would never adopt the λ technology, but may adopt λ' technology. Calculate the difference in the social values generated by these two technologies.

Exercise 12.6. Prove Proposition 12.3. In particular, verify that the conclusion is also true with limit pricing, i.e., $\Delta\hat{\pi}_1^I < \pi_1^I$.

Exercise 12.7. Consider the model in Section 12.3 with an incumbent monopolist and an entrant. Suppose that the cost of innovation for the incumbent is μ, while for the entrant it is $\chi\mu$, where $\chi \geq 1$.

(1) Explain why we may have $\chi > 1$.
(2) Show that there exists $\bar{\chi} > 1$ such that if $\chi < \bar{\chi}$, the entrant has greater incentives to undertake innovation, and if $\chi > \bar{\chi}$, the incumbent has greater incentives to undertake innovation.
(3) What is the effect of the elasticity of demand on the relative incentives of the incumbent and the entrant to undertake innovation? Provide in intuition for this effect.

Exercise 12.8. (1) Prove Proposition 12.4 by providing an example in which there is excessive innovation incentives.
(2) What factors make excessive innovation more likely?

Exercise 12.9. The discussion in the text presumed a particular form of patent policy, which provided ex post monopoly power to the innovator. An alternative intellectual property right policy is licensing, where firms that have made an innovation can license the rights to use this innovation to others. This exercise asks you to work through the implications of this type of licensing. Throughout, let us think of the licensing stage as follows: the innovator can make a take-it-or-leave-it-offer to one or many firms so that they can buy the rights to use the innovation (and produce as many units of the output as they like) in return for some licensing fee ν.

(1) Consider the competitive environment we started with and show that if firm 1 is allowed to license its innovation to others, this can never raise its profits and it can never increase its incentives to undertake the innovation. Provide an intuition for this result.
(2) Now modify the model, so that each firm has a strictly convex and increasing cost of producing, $\psi_1(q)$, and also has to pay a fixed cost of $\psi_0 > 0$ to be active (so that the average costs take the familiar inverse U shape). Show that licensing can be beneficial for firm 1 in this case and therefore increase incentives to undertake the innovation. Explain why the results differ between the two cases.

Exercise 12.10. Derive the expression for the ideal price index, (12.11), from (12.10) and the definition of the consumption index C.

Exercise 12.11. Consider the maximization problem in (12.13) and write down the first-order conditions taking into account the impact of p_i on P and C. Show that as $N \to \infty$, the solution to this problem converges to (12.14).

Exercise 12.12. In the Dixit-Stiglitz model, determine the conditions on the function $v(\cdot)$ such that an increase in N raises the profits of a monopolist.

Exercise 12.13. Suppose that $U(C, y) = C + v(y)$, where $v(y) = y^{1-\alpha}/(1-\alpha)$ with $\alpha \in (0, 1)$. Suppose also that new varieties can be introduced at the fixed cost μ.

(1) Consider the allocation determined by a social planner also controlling prices. Characterize the number of varieties that a social planner would choose in order to maximize the utility of the representative household in this case.
(2) Suppose that prices are given by (12.14). Characterize the number of varieties that the social planner would choose in order to maximize utility of the representative household in this case.
(3) Characterize the equilibrium number of varieties (at which all monopolistically competitive variety producers makes zero profits) and compare this with the answers to the previous two parts. Explain the sources of differences between the equilibrium and the social planner's solution in each case.

Exercise 12.14. This exercise asks you to work through the Salop (1979) model of product differentiation, which differs from the Dixit-Stiglitz model in that equilibrium markups are declining in the number of firms. Imagine that consumers are located uniformly around a circle with perimeter equal to 1. The circle indexes both the preferences of heterogeneous consumers and the types of goods. The point where the consumer is located along the circle corresponds to the type of product that he most prefers. When a consumer at point x around the circle consumes a good of type z, his utility is

$$R - t\,|z - x| - p,$$

while if he chooses not to consume, his utility is 0. Here R can be thought of as the reservation utility of the individual, while t parameterizes the "transport" costs that the individual has to pay in order to consume a good that is away from his ideal point along the circle. Suppose that each firm has a marginal cost of ψ per unit of production

(1) Imagine a consumer at point x, with the two neighboring firms at points $z_1 > x > z_2$. As long as the prices of these firms are not much higher than those further a far, the consumer will buy from one of these two firms. Denote the prices of these two firms by p_1 and p_2. Show that the price difference that would make the consumer indifferent between purchasing from the two firms satisfies

$$p_1 - p_2 = (2x - z_1 - z_2)\,t$$

with

$$t\,(z_1 - x) + p_1 \leq R.$$

(2) Suppose that p_1 and p_2 satisfy the above relationships. Then, show that all $x' \in [z_2, x)$ strictly prefer to buy from firm 2 and all $x' \in (x, z_1]$ strictly prefer to buy from firm 1.
(3) Now assume that there are three firms along the circle at locations $z_1 > z_2 > z_3$. Show that firm 2's profits are given by

$$\pi_2\,(p_1, p_2, p_3 \mid z_1, z_2, z_3) = (p_2 - \psi)\left(\frac{p_1 - p_2}{2t} + \frac{z_1 - z_2}{2} + \frac{p_3 - p_2}{2t} + \frac{z_2 - z_3}{2}\right)$$

and calculate its profit maximizing price.

(4) Now look at the location choice of firm 2. Suppose that $p_1 = p_3$. Show that it would like to locate half way between z_1 and z_3.

(5) Prove that in an equilibrium with N firms charging the same price must have the distance between any two firms equal to $1/N$.

(6) Show that when there are N equity-distant firms, it is an equilibrium for all firms to charge

$$p = \psi + \frac{t}{N}.$$

(7) Explain why the markup here is a decreasing function of the number of firms, whereas it was independent of the number of firms in the Dixit-Stiglitz model.

CHAPTER 13

Expanding Variety Models

As emphasized in the previous chapter, the key to understanding endogenous technological progress is that R&D is a purposeful activity, undertaken for profits and that the knowledge (machines, blueprints, or new technologies) that it generates increases the productivity of existing factors of production. The first endogenous technological change models were formulated by Romer (1987 and 1990). Different versions have been analyzed by Segerstrom, Anant and Dinopoulos (1990), Grossman and Helpman (1991a,b), Aghion and Howitt (1992). Some of those will be discussed in the next chapter.

The simplest models of endogenous technological change are those in which R&D expands the variety of inputs or machines used in production. This chapter focuses on models with expanding input varieties; research will lead to the creation of new varieties of inputs (machines) and a greater variety of inputs will increase the "division of labor," raising the productivity of final good firms. This can therefore be viewed as a form of *process innovation*. An alternative, formulated and studied by Grossman and Helpman (1991a,b), focuses on *product innovation*. In this model, research leads to the invention of new goods, and because individuals have "love-for-variety" preferences, they derive greater utility when they consume a greater variety of products. Consequently "real" income increases as a result of these product innovations. The comparison of Grossman-Helpman's model of product innovation with the baseline model of process innovation presented below will show that the two models are mathematically very similar (though the model of product innovation is slightly more involved and is thus treated at the end of this chapter).

In all of these models, and also in the models of quality competition in the next chapter, I use the Dixit-Stiglitz constant elasticity structure introduced in the previous chapter.

13.1. The Lab-Equipment Model of Growth with Input Varieties

Let us start with a particular version of the endogenous growth model with expanding varieties of inputs and an R&D technology that only uses output for creating new inputs. This is sometimes referred to as the *lab-equipment* model, since all that is required for research is investment in equipment or in laboratories—rather than the employment of skilled or unskilled workers or scientists.

13.1.1. Demographics, Preferences and Technology. Imagine an infinite-horizon economy in continuous time admitting a representative household with preferences

$$\int_0^{\infty} \exp\left(-\rho t\right) \frac{C\left(t\right)^{1-\theta}-1}{1-\theta} dt. \tag{13.1}$$

There is no population growth and the total population of workers, L, supplies labor inelastically throughout. Let us assume, as discussed in the previous chapter, that the representative household owns a balanced portfolio of all the firms in the economy. Alternatively, we can think of the economy as consisting of many households with the same preferences as the representative household and each household holding a balanced portfolio of all the firms.

The unique final good of the economy is produced with the following aggregate production function:

$$Y\left(t\right) = \frac{1}{1-\beta} \left[\int_0^{N(t)} x(\nu,t)^{1-\beta} d\nu \right] L^{\beta}, \tag{13.2}$$

where L is the aggregate labor input, $N\left(t\right)$ denotes the different number of varieties of inputs (machines) available to be used in the production process at time t, and $x\left(\nu,t\right)$ is the total amount of input (machine) type ν used at time t. I assume that x's depreciate fully after use, thus they can be interpreted as generic inputs, as intermediate goods, as machines, or even as capital as long as we are comfortable with the assumption that there is immediate depreciation. The assumption that the inputs or machines are "used up" in production or depreciate immediately after being used makes sure that the amounts of inputs used in the past do not become additional state variables. This simplifies the exposition of the model (though the results are identical without this assumption, see Exercise 13.22). Nevertheless, I refer to the inputs as "machines," which makes the economic interpretation of the problem easier.

The term $(1-\beta)$ in the denominator is included for notational simplicity. Notice that for given $N\left(t\right)$, which final good producers take as given, eq. (13.2) exhibits *constant returns to scale*. Therefore, final good producers are competitive and are subject to constant returns to scale, justifying our use of the aggregate production function to represent their production possibilities set.

One can also write (13.2) in the following form

$$Y\left(t\right) = \frac{1}{1-\beta} \tilde{\mathbf{X}}\left(t\right)^{1-\beta} L^{\beta},$$

where

$$\tilde{\mathbf{X}}\left(t\right) \equiv \left[\int_0^{N(t)} x(\nu,t)^{\frac{\varepsilon_\beta - 1}{\varepsilon_\beta}} d\nu \right]^{\frac{\varepsilon_\beta}{\varepsilon_\beta - 1}},$$

with $\varepsilon_\beta \equiv 1/\beta$ as the elasticity of substitution between inputs. This form emphasizes both the constant returns to scale properties of the production function and the continuity between the model here and the Dixit-Stiglitz model of the previous chapter.

The resource constraint of the economy at time t is

$$C(t) + X(t) + Z(t) \leq Y(t), \tag{13.3}$$

where $X(t)$ is investment or spending on inputs at time t and $Z(t)$ is expenditure on R&D at time t, which comes out of the total supply of the final good.

We next need to specify how quantities of machines are created and how new machines are invented. Let us first assume that once the blueprint of a particular input is invented, the research firm can create one unit of that machine at marginal cost equal to $\psi > 0$ units of the final good. The *innovation possibilities frontier*, which regulates the technology for the creation of new machines, takes the following form:

$$\dot{N}(t) = \eta Z(t), \tag{13.4}$$

where $\eta > 0$, and the economy starts with some initial technology stock $N(0) > 0$. This implies that greater spending on R&D leads to the invention of new machines. Throughout, I assume that there is free entry into research, which means that any individual or firm can spend one unit of the final good at time t in order to generate a flow rate η of the blueprints of new machines. The firm that discovers these blueprints receives a fully-enforced perpetual patent on this machine.

There is no aggregate uncertainty in the innovation process. Naturally, there will be uncertainty at the level of the individual firm, but with many different research labs undertaking such expenditure, at the aggregate level, eq. (13.4) holds deterministically.

Given the patent structure specified above, a firm that invents a new machine variety is the sole supplier of that type of machine, say machine of type ν, and sets a price of $p^x(\nu, t)$ at time t to maximize profits. Since machines depreciate after use, $p^x(\nu, t)$ can also be interpreted as a "rental price" or the user cost of this machine.

The demand for machine of type ν is obtained by maximizing net aggregate profits of the final good sector as given by (13.2) minus the cost of inputs. Since machines depreciate after use and labor is hired on the spot market for its flow services, the maximization problem of the final good sector can be considered for each point in time separately and simply requires the maximization of the instantaneous profits of a representative final good producer. These instantaneous profits can be obtained by subtracting the total inputs costs, the user costs of renting machines and labor costs, from the value of our production. Therefore, the maximization problem at time t is:

$$\max_{[x(\nu,t)]_{lv\in[0,N(t)]},L} \frac{1}{1-\beta}\left[\int_0^{N(t)} x(\nu,t)^{1-\beta} d\nu\right] L^{\beta} - \int_0^{N(t)} p^x(\nu, t)\, x(\nu,t) d\nu - w(t) L. \tag{13.5}$$

The first-order condition of this maximization problem with respect to $x(\nu, t)$ for any $\nu \in [0, N(t)]$ yields the demand for machines from the final good sector. These demands take the convenient isoelastic form:

$$x(\nu,t) = p^x(\nu,t)^{-1/\beta} L, \tag{13.6}$$

which is intuitive in view of the fact that elasticity of demand for different machine varieties is $\varepsilon_\beta \equiv 1/\beta$ (so that $x(\nu,t) = p^x(\nu,t)^{-\varepsilon_\beta} L$). This equation implies that the demand for machines only depends on the user cost of the machine and on equilibrium labor supply, but not on the interest rate, $r(t)$, the wage rate, $w(t)$, or the total measure of available machines, $N(t)$. This feature makes the model very tractable.

Next, consider the maximization problem of a monopolist owning the blueprint of a machine of type ν invented at time t. Since the representative household holds a balanced portfolio of all the firms in the economy and there is a continuum of firms, there will be no aggregate uncertainty, so each monopolist's objective is to maximize profits. Consequently, this monopolist chooses an investment plan and a sequence of capital stocks so as to maximize the present discounted value of profits starting from time t. Recalling that the interest rate at time t is $r(t)$ and that the marginal cost of producing machines (in terms of the final good) is ψ, the net present discounted value can be written as:

$$V(\nu,t) = \int_t^\infty \exp\left[-\int_t^s r\left(s'\right) ds'\right] \pi(\nu,s)\, ds \tag{13.7}$$

where

$$\pi(\nu,t) \equiv p^x(\nu,t)x(\nu,t) - \psi x(\nu,t)$$

denotes profits of the monopolist producing intermediate ν at time t, $x(\nu,t)$ and $p^x(\nu,t)$ are the profit-maximizing choices for the monopolist, and $r(t)$ is the market interest rate at time t. Alternatively, assuming that the value function is differentiable in time, this could be written in the form of Hamilton-Jacobi-Bellman equations as in Theorem 7.10 in Chapter 7:

$$r(t)\, V(\nu,t) - \dot{V}(\nu,t) = \pi(\nu,t). \tag{13.8}$$

Exercise 13.1 asks you to provide a different derivation of this equation than in Theorem 7.10.

13.1.2. Characterization of Equilibrium. An allocation in this economy is defined by the following objects: time paths of consumption levels, aggregate spending on machines, and aggregate R&D expenditure $[C(t), X(t), Z(t)]_{t=0}^\infty$, time paths of available machine types, $[N(t)]_{t=0}^\infty$, time paths of prices and quantities of each machine and the net present discounted value of profits from that machine, $[p^x(\nu,t), x(\nu,t), V(\nu,t)]_{\nu\in N(t),t=0}^\infty$, and time paths of interest rates and wage rates, $[r(t), w(t)]_{t=0}^\infty$.

An equilibrium is an allocation in which all existing research firms choose $[p^x(\nu,t), x(\nu,t)]_{\nu\in[0,N(t)],t=0}^\infty$ to maximize profits, the evolution of $[N(t)]_{t=0}^\infty$ is determined by free entry, the time paths of interest rates and wage rates, $[r(t), w(t)]_{t=0}^\infty$, are consistent with market clearing, and the time paths of $[C(t), X(t), Z(t)]_{t=0}^\infty$ are consistent with consumer optimization. Note that this is not a competitive equilibrium, since technology firms have market power.

For characterizing the equilibrium, let us start with the firm side. Since (13.6) defines isoelastic demands, the solution to the maximization problem of any monopolist $\nu \in [0, N(t)]$

involves setting the same price in every period (see Exercise 13.2):

$$p^x(\nu,t)=\frac{\psi}{1-\beta} \text{ for all } \nu \text{ and } t. \tag{13.9}$$

That is, all monopolists charge a constant rental rate, equal to a markup over the marginal cost. Without loss of generality, let us normalize the marginal cost of machine production to $\psi \equiv (1-\beta)$, so that

$$p^x(\nu,t)=p^x=1 \text{ for all } \nu \text{ and } t.$$

Profit-maximization also implies that each monopolist rents out the same quantity of machines in every period, equal to

$$x(\nu,t)=L \text{ for all } \nu \text{ and } t. \tag{13.10}$$

This gives monopoly profits as:

$$\pi(\nu,t)=\beta L \text{ for all } \nu \text{ and } t. \tag{13.11}$$

The major implication of this equation is that each monopolist sells exactly the same amount of machines, charges the same price and makes the same amount of profits at all time points. This particular feature simplifies the analysis of endogenous technological change models with expanding variety.

Substituting (13.6) and the machine prices into (13.2) yields a (derived) production function for the final good,

$$Y(t)=\frac{1}{1-\beta}N(t)L. \tag{13.12}$$

This is one of the main equations of the expanding input or product variety models. It shows that even though the aggregate production function exhibits constant returns to scale from the viewpoint of final good firms (which take $N(t)$ as given), there are *increasing returns to scale* for the entire economy. Moreover, (13.12) makes it clear that an increase in the variety of machines, $N(t)$, raises the productivity of labor and that when $N(t)$ increases at a constant rate so will output per capita.

The labor demand of the final good sector follows from the first-order condition of maximizing (13.5) with respect to L and implies the equilibrium condition

$$w(t)=\frac{\beta}{1-\beta}N(t). \tag{13.13}$$

Finally, free entry into research implies that at all points in time

$$\eta V(\nu,t)\leq 1,\ Z(\nu,t)\geq 0 \text{ and } (\eta V(\nu,t)-1)\,Z(\nu,t)=0, \text{ for all } \nu \text{ and } t, \tag{13.14}$$

where $V(\nu,t)$ is given by (13.7). To understand (13.14), recall that one unit of final good spent on R&D leads to the invention of η units of new inputs, each with a net present discounted value of profits given by (13.7). This free-entry condition is written in the complementary slackness form, since research may be sufficiently unprofitable that there may be zero R&D

effort, in which case $\eta V(\nu,t)$ could be strictly less than 1. Nevertheless, for the relevant parameter values there will be positive entry and economic growth (and technological *progress*), so we often simplify the exposition by writing the free-entry condition as

$$\eta V(\nu,t)=1.$$

Note also that since each monopolist $\nu\in[0,N(t)]$ produces machines given by (13.10), and there are a total of $N(t)$ monopolists, the total expenditure on machines is

$$X(t)=(1-\beta)N(t)L. \tag{13.15}$$

Finally, the representative household's problem is standard and implies the usual Euler equation for consumption,

$$\frac{\dot{C}(t)}{C(t)}=\frac{1}{\theta}(r(t)-\rho), \tag{13.16}$$

and the transversality condition

$$\lim_{t\to\infty}\left[\exp\left(-\int_0^t r(s)\,ds\right)N(t)V(t)\right]=0, \tag{13.17}$$

which is written in the "market value" form and requires the value of the total wealth of the representative household, which is equal to the value of corporate assets $N(t)V(t)$, not to grow faster than the discount rate (see Exercise 13.3).

In light of the previous equations, an equilibrium can be defined more formally as time paths of consumption, expenditures, R&D decisions and total number of varieties, $[C(t),X(t),Z(t),N(t)]_{t=0}^{\infty}$, such that (13.3), (13.15), (13.16), (13.17) and (13.14) are satisfied; time paths of prices and quantities of each machine and the net present discounted value of profits from that machine, $[p^x(\nu,t),x(\nu,t)]_{\nu\in N(t),t=0}^{\infty}$, that satisfy (13.9) and (13.10); and time paths of interest rate and wages, $[r(t),w(t)]_{t=0}^{\infty}$, such that (13.13) and (13.16), hold.

A *balanced growth path (BGP)* is again an equilibrium path where $C(t),X(t),Z(t)$ and $N(t)$ grow at a constant rate. Such an equilibrium can alternatively be referred to as a "steady state", since it is a steady state in transformed variables (even though the original variables grow at a constant rate). This is a feature of all the growth models and I therefore use the terms steady state and BGP interchangeably in the context of models with sustained growth.

13.1.3. Balanced Growth Path. A balanced growth path (BGP) requires that consumption grows at a constant rate, say g_C^*. This is only possible from (13.16) if the interest rate is constant. Let us therefore look for an equilibrium allocation in which

$$r(t)=r^* \text{ for all } t,$$

where "*" refers to BGP values. Since profits at each date are given by (13.11) and since the interest rate is constant, (13.8) implies that $\dot{V}(t)=0$. Substituting this in either (13.7) or (13.8),

$$V^*=\frac{\beta L}{r^*}. \tag{13.18}$$

This equation is intuitive: a monopolist makes a flow profit of βL, and along the BGP, this is discounted at the constant interest rate r^*.

Let us next suppose that the (free entry) condition (13.14) holds as an equality, in which case we also have

$$\frac{\eta\beta L}{r^*} = 1.$$

This equation pins down the steady-state interest rate, r^*, as:

$$r^* = \eta\beta L$$

The consumer Euler equation, (13.16), then implies that the rate of growth of consumption in BGP must be given by

$$g_C^* = \frac{\dot{C}(t)}{C(t)} = \frac{1}{\theta}(r^* - \rho). \tag{13.19}$$

Moreover, it can be verified that the current-value Hamiltonian for the consumer's maximization problem is concave, thus this condition, together with the transversality condition, are sufficient for characterizing the optimal consumption plans of the consumer (recall Theorem 7.14 in Chapter 7).

In a BGP, consumption cannot grow at a different rate than total output (see Exercise 13.5), thus we must also have the growth rate of output in the economy is

$$g^* = g_C^*.$$

Therefore, given the BGP interest rate, the long-run growth rate of the economy is obtained as:

$$g^* = \frac{1}{\theta}(\eta\beta L - \rho). \tag{13.20}$$

Let us next assume that

$$\eta\beta L > \rho \text{ and } (1-\theta)\eta\beta L < \rho, \tag{13.21}$$

which will ensure that $g^* > 0$ and that the transversality condition is satisfied.

PROPOSITION 13.1. *Suppose that condition (13.21) holds. Then, in the above-described lab-equipment expanding input variety model, there exists a unique BGP in which technology, output and consumption all grow at the same rate, g^*, given by (13.20).*

PROOF. The preceding discussion establishes all the claims in the proposition except that the transversality condition holds. You are asked to check this in Exercise 13.7. □

An important feature of this class of endogenous technological progress models is the presence of the *scale effect*, which was already discussed in Section 11.4 in Chapter 11 in the context of the Romer (1986a) model: the larger is L, the greater is the growth rate. The scale effect comes from a very strong form of the *market size effect* discussed in the previous chapter. As illustrated there, the increasing returns to scale nature of the technology (for example, as highlighted in eq. (13.12)) is responsible for this strong form of the market size effect and thus for the scale effect. We will see in Section 15.5 in Chapter 15 that it is possible

to have variants of the current model that feature the market size effect, but not the scale effect.

13.1.4. Transitional Dynamics. It is also straightforward to see that there are no transitional dynamics in this model. To derive this result, let us go back to the value function for each monopolist. Substituting for profits, this gives

$$r(t) V(\nu, t) - \dot{V}(\nu, t) = \beta L.$$

The key observation is that positive growth at any point implies that $\eta V(\nu, t) = 1$ for all t. In other words, if $\eta V(\nu, t') = 1$ for some t', then $\eta V(\nu, t) = 1$ for all t (see Exercise 13.4). Now differentiating $\eta V(\nu, t) = 1$ with respect to time yields $\dot{V}(\nu, t) = 0$, which is only consistent with $r(t) = r^*$ for all t, thus

$$r(t) = \eta \beta L \text{ for all } t.$$

This establishes:

Proposition 13.2. *Suppose that condition (13.21) holds. In the above-described lab-equipment expanding input-variety model, with initial technology stock $N(0) > 0$, there is a unique equilibrium path in which technology, output and consumption always grow at the rate g^* as in (13.20).*

At some level, this result is not too surprising. While the microfoundations and the economics of the expanding varieties model studied here are very different from the neoclassical AK economy, the mathematical structure of the model is very similar to the AK model (as most clearly illustrated by the derived equation for output, (13.12)). Consequently, as in the AK model, the economy always grows at a constant rate.

Even though the mathematical structure of the model is similar to the neoclassical AK economy, it is important to emphasize that the economics here is very different. The equilibrium in Proposition 13.2 exhibits *endogenous technological progress*. In particular, research firms spend resources in order to invent new inputs. They do so because, given their patents, they can profitably sell these inputs to final good producers. It is therefore profit incentives that drive R&D, and R&D drives economic growth. We have therefore arrived to our first model in which market-shaped incentives determine the rate at which the technology of the economy evolves over time.

13.1.5. Pareto Optimal Allocations. The presence of monopolistic competition implies that the competitive equilibrium is not necessarily Pareto optimal (and Theorems 5.5 and 5.6 in Chapter 5 no longer apply). In particular, the current model exhibits a version of the *aggregate demand externalities* discussed in the previous chapter and features two sources of potential inefficiencies. First, there is a markup over the marginal cost of production of inputs. Second, the number of inputs produced at any point in time may not be optimal. The first source of inefficiency is familiar from models of static monopoly, while the second emerges from the fact that in this economy the set of traded (Arrow-Debreu) commodities is endogenously determined. This latter source of potential inefficiency relates to the issue of

endogenously incomplete markets (there is no way to purchase an input that is not supplied in equilibrium) and will be discussed in greater detail in Section 17.6 in Chapter 17.

To contrast the equilibrium allocations with the Pareto optimal allocations, let us set up the problem of the social planner and derive the optimal growth rate. Notice that the social planner will also use the same quantity of all types of machines in production, but because of the absence of a markup, this quantity will be different than in the equilibrium allocation. The social planner will also take into account the effect of an increase in the variety of inputs on the overall productivity in the economy, which monopolists did not, because they did not appropriate the full surplus from inventions.

More explicitly, given $N(t)$, the social planner will choose

$$\max_{[x(\nu,t)]_{\nu\in[0,N(t)]},L} \frac{1}{1-\beta}\left[\int_0^{N(t)} x(\nu,t)^{1-\beta}d\nu\right]L^{\beta}-\int_0^{N(t)}\psi x(\nu,t)d\nu,$$

which only differs from the equilibrium profit maximization problem, (13.5), because the marginal cost of machine creation, ψ, is used as the cost of machines rather than the monopoly price, and the cost of labor is not subtracted. Recalling that $\psi\equiv 1-\beta$, the solution to this program involves

$$x^{S}(\nu,t)=(1-\beta)^{-1/\beta}L,$$

so that the social planner's output level will be

$$\begin{aligned} Y^{S}(t) &= \frac{(1-\beta)^{-(1-\beta)/\beta}}{1-\beta}N^{S}(t)L \\ &= (1-\beta)^{-1/\beta}N^{S}(t)L, \end{aligned}$$

where superscripts "S" are used to emphasize that the level of technology and thus the level of output will differ between the social planner's allocation and the equilibrium allocation. The aggregate resource constraint is still given by (13.3). Let us define net output, which subtracts the cost of machines from total output, as

$$\tilde{Y}^{S}(t)\equiv Y^{S}(t)-X^{S}(t).$$

This is relevant, since it is net output that will be distributed between R&D expenditure and consumption. Net output can be written as

$$\begin{aligned} \tilde{Y}^{S}(t) &= (1-\beta)^{-1/\beta}N^{S}(t)L-\int_0^{N^{S}(t)}\psi x^{S}(\nu,t)\,d\nu \\ &= (1-\beta)^{-1/\beta}N^{S}(t)L-(1-\beta)^{-(1-\beta)/\beta}N^{S}(t)L \\ &= (1-\beta)^{-1/\beta}\beta N^{S}(t)L. \end{aligned}$$

Given this expression and (13.4), the maximization problem of the social planner can be written as

$$\max\int_0^{\infty}\frac{C(t)^{1-\theta}-1}{1-\theta}\exp(-\rho t)\,dt$$

subject to

$$\dot{N}(t)=\eta(1-\beta)^{-1/\beta}\beta N(t)L-\eta C(t).$$

In this problem, $N(t)$ is the state variable, and $C(t)$ is the control variable. Let us set up the current-value Hamiltonian,

$$\hat{H}(N,C,\mu)=\frac{C(t)^{1-\theta}-1}{1-\theta}+\mu(t)\left[\eta(1-\beta)^{-1/\beta}\beta N(t)L-\eta C(t)\right].$$

The necessary conditions for the social planner' as problem are

$$\begin{aligned}\hat{H}_C(N,C,\mu) &= C(t)^{-\theta}-\eta\mu(t)=0\\ \hat{H}_N(N,C,\mu) &= \mu(t)\eta(1-\beta)^{-1/\beta}\beta L=\rho\mu(t)-\dot{\mu}(t)\\ \lim_{t\to\infty}\left[\exp(-\rho t)\mu(t)N(t)\right] &= 0.\end{aligned}$$

It can be verified easily that the current-value Hamiltonian of the social planner is concave and satisfy the conditions of Theorem 7.14, thus the above conditions are both necessary and sufficient for an optimal solution (see Exercise 13.9).

Combining these conditions, we obtain the following growth rate for consumption in the social planner's allocation (see again Exercise 13.9):

$$\frac{\dot{C}^S(t)}{C^S(t)}=\frac{1}{\theta}\left(\eta(1-\beta)^{-1/\beta}\beta L-\rho\right). \tag{13.22}$$

Like the equilibrium, the social planner also chooses and allocation with a constant rate of consumption growth (thus no transitional dynamics). The growth rate of consumption chosen by the social planner, (13.22), can be directly compared to the growth rate in the decentralized equilibrium, (13.20). The comparison boils down to that of

$$(1-\beta)^{-1/\beta}\beta \text{ to } \beta,$$

and it is straightforward to see that the former is always greater since $(1-\beta)^{-1/\beta}>1$ by virtue of the fact that $\beta\in(0,1)$. This implies that the socially-planned economy will always grow faster than the decentralized economy

Proposition 13.3. *In the above-described expanding input variety model, the decentralized equilibrium is always Pareto suboptimal. Starting with any $N(0)>0$, the Pareto optimal allocation involves a constant growth rate*

$$g^S=\frac{1}{\theta}\left(\eta(1-\beta)^{-1/\beta}\beta L-\rho\right),$$

which is strictly greater than the equilibrium growth rate g^* given in (13.20).

Proof. Most of the proof is provided in the preceding discussion. Exercise 13.11 asks you to show that the Pareto optimal allocation always involves a constant growth rate and no transitional dynamics. □

Intuitively, the Pareto optimal growth rate is greater than the equilibrium growth rate because the social planner values innovation more. The greater social value of innovations stems from the fact that the social planner is able to use the machines more intensively after innovation, since the monopoly markup reducing the demand for machines is absent in the social planner's allocation. This is related to the *pecuniary externality* resulting from the

monopoly markups (and is thus related to the aggregate demand externalities discussed in the previous chapter). Moreover, this pecuniary externality also affects the set of traded commodities (thus the rate of growth of inputs and technology). Other models of endogenous technological progress we will study in this chapter incorporate technological spillovers and thus generate inefficiencies both because of the pecuniary externality isolated here and because of the standard technological spillovers.

13.1.6. Policy in Models of Endogenous Technological Progress. The divergence between the decentralized equilibrium and the socially planned allocation introduces the possibility of Pareto-improving policy interventions. There are two natural alternatives to consider:

(1) *Subsidies to Research:* by subsidizing research, the government can increase the growth rate of the economy, and this can be turned into a Pareto improvement if taxation is not distortionary and there can be appropriate redistribution of resources so that all parties benefit.
(2) *Subsidies to Capital Inputs:* inefficiencies also arise from the fact that the decentralized economy is not using as many units of the machines/capital inputs (because of the monopoly markup); so subsidies to capital inputs given to final good producers would also be useful in increasing the growth rate.

Moreover, it is noteworthy that as in the first-generation endogenous growth models, a variety of different policy interventions, including taxes on investment income and subsidies of various forms, will have growth effects not just level effects in this framework (see, for example, Exercise 13.13).

Naturally, once we start thinking of policy in order to close the gap between the decentralized equilibrium and the Pareto optimal allocation, we also have to think of the objectives of policymakers and this brings us to political economy issues, which are the subject matter of Part 8. For that reason, I will not go into a detailed discussion of optimal policy (leaving some of this to you in Exercises 13.12-13.14). Nevertheless, it is useful to briefly discuss the role of competition policy in models of endogenous technological progress.

Recall that the optimal price that the monopolist charges for machines is

$$p^x = \frac{\psi}{1-\beta}.$$

Imagine, instead, that a fringe of competitive firms can copy the innovation of any monopolist, but they will not be able to produce at the same level of costs (because the inventor has more know-how). In particular, as in the previous chapter, suppose that instead of a marginal cost ψ, they will have marginal cost of $\gamma\psi$ with $\gamma > 1$. If $\gamma > 1/(1-\beta)$, then this fringe is not a threat to the monopolist, since the monopolist could set its ideal, profit maximizing, markup and the fringe would not be able to enter without making losses. However, if $\gamma < 1/(1-\beta)$, then the fringe would prevent the monopolist from setting its ideal monopoly price. In

particular, in this case the monopolist would be forced to set a "limit price",

$$p^x = \gamma\psi, \tag{13.23}$$

which has an explanation identical to the limit price in the previous chapter.[1]

When the monopolist charges this limit price, its profits per unit would be

$$\text{profits per unit } = (\gamma - 1)\,\psi = (\gamma - 1)\,(1 - \beta)\,,$$

which is less than β, the profits per unit that the monopolist made in the absence of the competitive fringe.

What is the implication of this on the rate of economic growth? It is straightforward to work out that in this case the economy would grow at a slower rate. For example, in the baseline model with the lab-equipment technology, this growth rate would be (see Exercise 13.15):

$$\hat{g} = \frac{1}{\theta}\left(\eta\gamma^{-1/\beta}\,(\gamma - 1)\,(1 - \beta)^{-(1-\beta)/\beta}\,L - \rho\right),$$

which is less than g^* given in (13.20). Therefore, in this model, greater competition, which reduces markups (and thus static distortions), also reduces long-run growth. This might at first appear counter-intuitive, since the monopoly markup may be thought to be the key source of inefficiency and greater competition (lower γ) reduces this markup. Nevertheless, as mentioned above, inefficiency results both because of monopoly markups and because the set of available inputs may not be appropriately chosen. As γ declines, monopoly markups decline, but the problem of underprovision of inputs becomes more severe. This is because profits are important in this model to encourage innovation by new research firms. If these profits are reduced, then incentives for research are also reduced. Since γ can also be interpreted as a parameter of anti-trust (competition) policy, this result implies that in the baseline endogenous technological change models more strict anti-trust policy reduces economic growth.

Welfare is not the same as growth, however, and some degree of competition reducing prices below the unconstrained monopolistic level might be useful for welfare depending on the discount rate of the representative household. Essentially, with a lower markup, households are happier in the present, but suffer slower consumption growth. The tradeoff between these two opposing effects depends on the discount rate of the representative household (see Exercise 13.15 for the details).

Similar results apply when we consider patent policy. In practice, patents are for limited durations. In the baseline model, we assumed that patents are perpetual; once a firm invents a new good, it has a fully-enforced patent forever and it becomes the monopolist for that good forever. If patents are enforced strictly, then this might rule out the competitive fringe from competing, restoring the growth rate of the economy to (13.20). Also, even in the absence of the competitive fringe, we can imagine that once the patent runs out, the firm

[1] More specifically, if the price of the monopolist were higher than $\gamma\psi$, the fringe could undercut the price of the monopolist, take over to market and make positive profits. If it were below this, then the monopolist could increase its price towards the unconstrained monopoly price and make more profits. Thus, there is a unique equilibrium price given by (13.23).

will cease making profits on its innovation. In this case, it can easily be shown that growth is maximized by having as long patents as possible. Again there is a tradeoff here between the equilibrium growth rate of the economy and the static level of welfare.

Perhaps, more important than these tradeoffs between growth and level is the fact that the models discussed in this chapter do not feature an interesting type of competition among firms. The quality competition (Schumpeterian) models introduced in the next chapter will allow a richer analysis of the effect of competition on innovation and economic growth.

13.2. Growth with Knowledge Spillovers

In the model of the previous section, growth resulted from the use of final output for R&D. This is similar, in some way, to the endogenous growth model of Rebelo (1991) discussed in Chapter 11, since the accumulation equation is linear in accumulable factors. As a result, in equilibrium output took a linear form in the stock of knowledge (new machines), thus a AN form instead of Rebelo's AK form.

An alternative is to have "scarce factors" used in R&D. In other words, instead of the lab-equipment specification, scientists are now the key creators of R&D. The lab-equipment model generated sustained economic growth by investing more and more resources in the R&D sector. This is impossible with scarce factors, since, by definition, a sustained increase in the use of these factors in the R&D sector is not possible. Consequently, with this alternative specification, there cannot be endogenous growth unless there are knowledge spillovers from past R&D, making the scarce factors used in R&D more and more productive over time. In other words, we now need current researchers to *"stand on the shoulder of past giants".* The original formulation of the endogenous technological change model by Romer (1990) relied on this type of knowledge spillovers. While these types of knowledge spillovers might be important in practice, the lab-equipment model studied in the previous section was a better starting point because it clearly delineated the role of technology accumulation and showed that growth need not be generated by technological externalities or spillovers.

Since knowledge spillovers play an important role in many models of economic growth, it is useful to see how the baseline endogenous technological progress model works in the presence of such spillovers. I now present the simplest version of the endogenous technological change model with knowledge spillovers. The environment is identical to that of the previous section, with the exception of the innovation possibilities frontier, which now takes the form

$$\dot{N}(t) = \eta N(t) L_R(t) \tag{13.24}$$

where $L_R(t)$ is labor allocated to R&D at time t. The term $N(t)$ on the right-hand side captures spillovers from the stock of existing ideas. The greater is $N(t)$, the more productive is an R&D worker. Notice that (13.24) imposes that these spillovers are proportional or *linear*. This linearity will be the source of endogenous growth in the current model. In the next section, we will see that a different kind of endogenous growth model can be formulated with less than proportional spillovers.

In (13.24), $L_R(t)$ is research employment, which comes out of the regular labor force. An alternative, which was originally used by Romer (1990), would be to suppose that only skilled workers or scientists can work in the knowledge-production (R&D) sector. Here we use the assumption that a homogeneous workforce is employed both in the R&D sector and in the final good sector. The advantage of this formulation is that competition between the production and the R&D sectors for workers ensures that the cost of workers to the research sector is given by the wage rate in final good sector. The only other change we need to make to the underlying environment is that now the total labor input *employed* in the final good sector, represented by the production function (13.2), is $L_E(t)$ rather than L, since some of the workers are working in the R&D sector. Labor market clearing then requires that

$$L_R(t) + L_E(t) \leq L.$$

The fact that not all workers are in the final good sector implies that the aggregate output of the economy (by an argument similar to before) is given by

$$Y(t) = \frac{1}{1-\beta} N(t) L_E(t), \tag{13.25}$$

and profits of monopolists from selling their machines is

$$\pi(t) = \beta L_E(t). \tag{13.26}$$

The net present discounted value of a monopolist (for a blueprint ν) is still given by $V(\nu, t)$ as in (13.7) or (13.8), with the flow profits given by (13.26). However, the free-entry condition is no longer the same as that which followed from eq. (13.4). Instead, (13.24) implies the following free-entry condition (when there is positive research):

$$\eta N(t) V(\nu, t) = w(t), \tag{13.27}$$

where $N(t)$ is on the left-hand side because it parameterizes the productivity of an R&D worker, while the flow cost of undertaking research is hiring workers for R&D, thus is equal to the wage rate $w(t)$.

The equilibrium wage rate must be the same as in the lab-equipment model of the previous section, in particular, $w(t) = \beta N(t) / (1-\beta)$ as in eq. (13.13), since the final good sector is unchanged. Moreover, balanced growth again requires that the interest rate must be constant at some level r^*. Using these observations together with the free-entry condition, BGP requires

$$\eta N(t) \frac{\beta L_E(t)}{r^*} = \frac{\beta}{1-\beta} N(t). \tag{13.28}$$

Hence the BGP equilibrium interest rate must be

$$r^* = (1-\beta)\eta L_E^*,$$

where L_E^* is the number of workers employed in production in BGP (given by $L_E^* = L - L_R^*$). The fact that the number of workers in production must be constant in BGP follows from

(13.28). Now using the Euler equation of the representative household, (13.16),

$$\frac{\dot{C}(t)}{C(t)} = \frac{1}{\theta}\left((1-\beta)\eta L_E^* - \rho\right) \equiv g^*, \tag{13.29}$$

for all t.

To complete the characterization of the BGP equilibrium, the BGP level of employment, L_E^*, needs to be determined. The innovation possibilities frontier, (13.24), implies $\dot{N}(t)/N(t) = \eta L_R^* = \eta(L - L_E^*)$. Moreover, by definition, the BGP growth rate of consumption must be equal to the rate of technological progress, thus $g^* = \dot{N}(t)/N(t)$. This implies that the BGP level of employment is uniquely pinned down as

$$L_E^* = \frac{\theta\eta L + \rho}{(1-\beta)\eta + \theta\eta}. \tag{13.30}$$

The rest of the analysis is unchanged. It can also be verified that there are no transitional dynamics in the decentralized equilibrium (see Exercise 13.17).

PROPOSITION 13.4. *Consider the above-described expanding input-variety model with knowledge spillovers and suppose that*

$$(1-\theta)(1-\beta)\eta L_E^* < \rho < (1-\beta)\eta L_E^*, \tag{13.31}$$

where L_E^* is the number of workers employed in production in BGP, given by (13.30). Then, there exists a unique BGP in which technology, output and consumption grow at the same rate, $g^* > 0$, given by (13.29) starting from any initial level of technology stock $N(0) > 0$.

PROOF. Most of the proof is given by the preceding discussion. Exercise 13.16 asks you to verify that the transversality condition is satisfied and that there are no transitional dynamics. □

Also, as in the lab-equipment model, the equilibrium allocation is Pareto suboptimal, and the Pareto optimal allocation involves a higher rate of output and consumption growth. Intuitively, while firms disregard the future increases in the productivity of R&D resulting from their own R&D spending, the social planner internalizes this effect (see Exercise 13.17). Finally, it is also useful to note that there is again a scale effect here—greater L increases the interest rate and the growth rate in the economy.

13.3. Growth without Scale Effects

The models used so far feature a scale effect in the sense that a larger population, L, translates into a higher interest rate and a higher growth rate. This is problematic for three reasons as argued in a series of papers by Chad Jones and others:

(1) Larger countries do not necessarily grow faster (though the larger market of the United States or European economies may have been an advantage during the early phases of the industrialization process; see also Chapter 21).

(2) The population of most nations has not been constant. With population growth as in the standard neoclassical growth model, e.g., $L(t) = \exp(nt) L(0)$, these models would not feature balanced growth, rather, the growth rate of the economy would steadily increase over time and output per capita would reach infinity in finite time ("explode").

(3) In the data, the total amount of resources devoted to R&D appears to increase steadily, but there is no associated increase in the aggregate growth rate.

Each one of these arguments against scale effects can be debated (for example, by arguing that countries do not provide the right level of analysis because of international trade linkages or that the growth rate of the world economy has indeed increased when we look at the past 2000 years rather than the past 100 years). Nevertheless together they do suggest that the strong form of scale effects embedded in the baseline endogenous technological change models may not provide a good approximation to reality. These observations have motivated Jones (1995) to suggest a modified version of the baseline endogenous technological progress model. While the type of modification to remove scale effect can be formulated in the lab-equipment model (see Exercise 13.21), it is conceptually simpler to do so in the context of the model with knowledge spillovers discussed in the previous section. In particular, in that model the scale effect can be removed by reducing the impact of knowledge spillovers.

More specifically, consider the model of the previous section with only two differences. First, there is population growth at the constant exponential rate n, so that $\dot{L}(t) = nL(t)$. The economy admits a representative household, which is also growing at the rate n, so that its preferences can be represented by the standard CRRA form:

$$\int_0^\infty \exp(-(\rho - n)t) \frac{C(t)^{1-\theta} - 1}{1-\theta} dt, \tag{13.32}$$

where $C(t)$ is consumption of the final good of the economy at time t, which is produced as before (with the production function (13.2)).

Second, in contrast to the knowledge-spillovers model studied in the previous section, the R&D sector only admits limited knowledge spillovers and (13.24) is replaced by

$$\dot{N}(t) = \eta N(t)^{\phi} L_R(t) \tag{13.33}$$

where $\phi < 1$ and $L_R(t)$ is labor allocated to R&D activities at time t. Labor market clearing requires

$$L_E(t) + L_R(t) = L(t), \tag{13.34}$$

where $L_E(t)$ is the level of employment in the final good sector and $L(t)$ is population at time t.

The key assumption for the model is that $\phi < 1$. The case where $\phi = 1$ is the one analyzed in the previous section, and as commented above, with population growth this would lead to an exploding path, leading to infinite utility. However, the model is well behaved when $\phi < 1$.

Aggregate output and profits are given by (13.25) and (13.26) as in the previous section. An equilibrium is also defined similarly. Let us focus on the BGP, where a constant fraction of workers are allocated to R&D, and the interest rate and the growth rate are constant. Suppose that this BGP involves positive growth, so that the free-entry condition holds as equality. Then, the BGP free-entry condition can be written as (see Exercise 13.18):

$$\eta N(t)^{\phi} \frac{\beta L_E(t)}{r^* - n} = w(t). \tag{13.35}$$

As before, the equilibrium wage is determined by the production side, (13.13), as $w(t) = \beta N(t)/(1-\beta)$. Combining this with the previous equation gives the following free-entry condition

$$\eta N(t)^{\phi-1} \frac{(1-\beta) L_E(t)}{r^*} = 1.$$

Now differentiating this condition with respect to time:

$$(\phi - 1)\frac{\dot{N}(t)}{N(t)} + \frac{\dot{L}_E(t)}{L_E(t)} = 0.$$

Since in BGP, the fraction of workers allocated to research is constant, $\dot{L}_E(t)/L_E(t) = n$. This implies that the BGP growth rate of technology is given by

$$g_N^* \equiv \frac{\dot{N}(t)}{N(t)} = \frac{n}{1-\phi}. \tag{13.36}$$

From eq. (13.12), this implies that total output grows at the rate $g_N^* + n$. But now there is population growth, so consumption per capita grows only at the rate

$$\begin{aligned} g_C^* &= g_N^* \\ &= \frac{n}{1-\phi}. \end{aligned} \tag{13.37}$$

The consumer Euler equation, equivalent of (11.4) incorporating the fact that the discount factor is $\rho - n$ instead of ρ, then determines the BGP interest rate as

$$\begin{aligned} r^* &= \theta g_N^* + \rho - n \\ &= \frac{\phi - (1-\theta)}{1-\phi} n + \rho > 0. \end{aligned}$$

The most noteworthy feature is that this model generates sustained and exponential growth in income per capita in the presence of population growth. More interestingly, in order to achieve this growth rate, it allocates more and more of labor to R&D. This is possible because of population growth. The reason for this is that the technology for creating new ideas, (13.33), only features limited spillovers, thus to maintain sustained growth, more resources need to be allocated to R&D.

Proposition 13.5. *In the above-described expanding input-variety model with limited knowledge spillovers as given by (13.33), starting from any initial level of technology stock $N(0) > 0$, there exists a unique BGP in which, technology and consumption per capita grow at the rate g_N^* as given by (13.36), and output grows at rate $g_N^* + n$.*

This analysis therefore shows that sustained equilibrium growth of per capita income is possible in an economy with growing population. Intuitively, instead of the linear (proportional) spillovers in the baseline Romer model, the current model allows only a limited amount of spillovers. Without population growth, these spillovers would affect the level of output, but would not be sufficient to sustain long-run growth. Continuous population growth, on the other hand, steadily increases the market size for new technologies and generates growth from these limited spillovers. While this pattern is referred to as "growth without scale effects," it is useful to note that there are two senses in which there are limited scale effects in these models. First, a faster rate of population growth translates into a higher equilibrium growth rate. Second, a larger population size leads to higher output per capita (see Exercise 13.20). It is not clear whether the data support these types of scale effects either. Put differently, some of the evidence suggested against the scale effects in the baseline endogenous technological change models may be inconsistent with this class of models as well. For example, there does not seem to be any evidence in the postwar data or from the historical data of the past 200 years that faster population growth leads to a higher equilibrium growth rate.

It is also worth noting that these models are sometimes referred to as "semi-endogenous growth" models, because while they exhibit sustained growth, the per capita growth rate of the economy, (13.37), is determined only by population growth and technology, and does not respond to taxes or other policies. Some papers in the literature have developed models of endogenous growth without scale effects, with equilibrium growth responding to policies, though this normally requires a combination of restrictive assumptions.

13.4. Growth with Expanding Product Varieties

Finally, I now brief the present the equivalent model in which growth is driven by *product innovations*, that is, by expanding product varieties rather than expanding varieties of inputs. The economy is in continuous time and has constant population L. It admits a representative household with preferences given by

$$\int_0^{\infty} \exp\left(-\rho t\right) \log C\left(t\right) dt, \tag{13.38}$$

where

$$C\left(t\right) \equiv \left[\int_0^{N(t)} c\left(\nu, t\right)^{\frac{\varepsilon-1}{\varepsilon}} d\nu\right]^{\frac{\varepsilon}{\varepsilon-1}} \tag{13.39}$$

is the consumption index, which is a CES aggregate of the consumption of different varieties. Here $c\left(\nu, t\right)$ denotes the consumption of product ν at time t, while $N\left(t\right)$ is the total measure of products. I assume throughout that $\varepsilon > 1$. Therefore, expanding input varieties have been replaced with expanding product varieties. The log specification in this utility function is for simplicity, and can be replaced by a CRRA utility function.

The patent to produce each product $\nu \in \left[0, N\left(t\right)\right]$ belongs to a monopolist, and the monopolist who invents the blueprints for a new product receives a fully enforced perpetual

patent on this product. Each product can be produced with the technology

$$y(\nu,t)=l(\nu,t), \tag{13.40}$$

where $l(\nu,t)$ is labor allocated to the production of this variety. Since the economy is closed, $y(\nu,t)=c(\nu,t)$.

As in model with knowledge spillovers of Section 13.2, I assume that new products can be produced with the production function

$$\dot{N}(t)=\eta N(t) L_R(t). \tag{13.41}$$

The reader will notice that there is a very close connection between the model here and the models of expanding input variety studied so far, especially the model with knowledge spillovers in Section 13.2. For instance, if $y(\nu,t)$s were interpreted as intermediate goods or inputs instead of products and if $C(t)$ in (13.39) were interpreted as the production function for the final good rather than part of the utility function of the representative household, the two models would be essentially identical. The only difference would be that, with this interpretation, labor would now be used in the production of the inputs, while in Section 13.2 it is used in the final good sector. This similarity emphasizes that the distinction between process and product innovations is fairly minor in theory, though this distinction might still be useful in mapping these models to reality.

An equilibrium and a BGP are defined similarly to before. The representative household now determines both the allocation of its expenditure on different varieties and the time path of consumption expenditures. Since the economy is closed and there is no capital, all output must be consumed. Nevertheless, the consumer Euler equation will apply to determine the equilibrium interest rate. Labor market clearing requires that

$$\int_0^{N(t)} l(\nu,t)\,d\nu + L_R(t) \leq L. \tag{13.42}$$

Let us start with expenditure decisions. Since the representative household has Dixit-Stiglitz preferences, the following consumer demands can be derived (see Exercise 13.24):

$$c(\nu,t)=\frac{p^x(\nu,t)^{-\varepsilon}}{\left(\int_0^{N(t)} p^x(\nu,t)^{1-\varepsilon}\,d\nu\right)^{-\frac{\varepsilon}{1-\varepsilon}}} C(t), \tag{13.43}$$

where $p^x(\nu,t)$ is the price of product variety ν at time t, and $C(t)$ is defined in (13.39). The term in the denominator is the ideal price index raised to the power $-\varepsilon$. It is most convenient to set this ideal price index as the numeraire, so that the price of output at every instant is normalized to 1, that is,

$$\left(\int_0^{N(t)} p^x(\nu,t)^{1-\varepsilon}\,d\nu\right)^{\frac{1}{1-\varepsilon}} = 1 \text{ for all } t. \tag{13.44}$$

With this choice of numeraire, the consumer Euler equation is (see Exercise 13.25):

$$\frac{\dot{C}(t)}{C(t)} = r(t)-\rho. \tag{13.45}$$

With similar arguments to before, the net present discounted value of the monopolist owning the patent for product ν can be written as

$$V(\nu,t)=\int_t^{\infty}\exp\left[-\int_t^s r\left(s'\right)ds'\right]\left[p^x(\nu,s)c(\nu,s)-w\left(s\right)c(\nu,s)\right]\,ds,$$

where $w\left(t\right)c(\nu,t)$ is the total expenditure of the firm to produce a total quantity of $c(\nu,t)$ (given the production function (13.40) and the wage rate at time t equal to $w\left(t\right)$), while $p^x(\nu,t)c(\nu,t)$ is its revenue, consistent with the demand function (13.43). The maximization of the net present discounted value again requires profit maximization at every instant. Since each monopolist faces the isoelastic demand curve given in (13.43), the profit-maximizing monopoly price is

$$p^x\left(\nu,t\right)=\frac{\varepsilon}{\varepsilon-1}w\left(t\right)\text{ for all }\nu\text{ and }t.$$

Since all firms charge the same price, they will all produce the same amount and employ the same amount of labor. At time t, there are $N\left(t\right)$ products, so the labor market clearing condition (13.42) implies that

$$c\left(\nu,t\right)=l\left(\nu,t\right)=\frac{L_E\left(t\right)}{N\left(t\right)}\text{ for all }\nu\text{ and }t, \tag{13.46}$$

where $L_E\left(t\right)=L-L_R\left(t\right)$. The instantaneous profits of each monopolist at time t can then be written as

$$\begin{aligned}\pi\left(\nu,t\right) &= p^x(\nu,t)c(\nu,t)-w\left(t\right)c(\nu,t)\\ &= \frac{1}{\varepsilon-1}\frac{L_E\left(t\right)}{N\left(t\right)}w\left(t\right)\text{ for all }\nu\text{ and }t.\end{aligned} \tag{13.47}$$

Since prices, sales and profits are equal for all monopolists, we can simplify notation by letting

$$V\left(t\right)=V\left(\nu,t\right)\text{ for all }\nu\text{ and }t.$$

In addition, since $c(\nu,t)=c\left(t\right)$ for all ν,

$$\begin{aligned}C\left(t\right) &= N\left(t\right)^{\frac{\varepsilon}{\varepsilon-1}}c\left(t\right).\\ &= L_E\left(t\right)N\left(t\right)^{\frac{1}{\varepsilon-1}},\end{aligned} \tag{13.48}$$

where the second equality uses (13.46).

Labor demand comes from the research sector as well as from the final good producers. Labor demand from research can again be determined using the free-entry condition. Assuming that there is positive research, so that the free-entry condition holds as an equality, this takes the form

$$\eta N\left(t\right)V\left(t\right)=w\left(t\right). \tag{13.49}$$

Combining this equation with (13.47) yields

$$\pi\left(t\right)=\frac{1}{\varepsilon-1}L_E\left(t\right)\eta V\left(t\right),$$

where we use $\pi\left(t\right)$ to denote the profits of all monopolists at time t. In BGP, where the fraction of the workforce working in research is constant, this implies that profits and the net present discounted value of monopolists must grow at the same rate. Let us denote the

BGP growth rate of the number of products, $N(t)$, by g_N, the growth rate of profit and values by g_V, and the growth rate of wages by g_w. Moreover, given the choice of numeraire, the consumption growth rate in this economy, g^*, must equal the wage growth rate, g_w. The free-entry condition, (13.49), then implies that $g^* = g_N + g_V$. Given these constant growth rates and the constant BGP interest rate r^*, (13.47) implies that in BGP:

$$V(t) = \frac{\pi(t)}{r^* - g^* + g_N}. \tag{13.50}$$

Intuitively, at time t, profits are equal to $\pi(t)$. Subsequently, because of product expansion, the number of employees per product decreases at the rate g_N, reducing profits, and wages increase at the rate g^*, increasing profits. Taking into account discounting at the rate r^* yields (13.50). Now combining (13.49) with (13.50) implies that

$$r^* = \frac{\eta}{\varepsilon - 1}(L - L_R^*) + g^* - g_N,$$

with L_R^* denoting the BGP size of the research sector. The R&D employment level of L_R^* combined with the R&D sector production function, (13.41) then implies

$$\frac{\dot{N}(t)}{N(t)} = g_N = \eta L_R^*.$$

However, we also know from the consumer Euler equation, (13.45) that in BGP $g^* = r^* - \rho$. Combining this with the previous two equations, the steady-state size of the research sector is obtained as

$$L_R^* = \frac{\eta L - (\varepsilon - 1)\rho}{\eta \varepsilon}. \tag{13.51}$$

I assume that $L_R^* > 0$, so that there is positive growth (otherwise, the free-entry condition would hold is inequality and there would be zero growth). Moreover, from (13.48) $g^* = g_N/(\varepsilon - 1)$, therefore

$$g^* = \frac{\eta L - (\varepsilon - 1)\rho}{\varepsilon(\varepsilon - 1)}. \tag{13.52}$$

Since $r^* > g^*$, the relevant transversality condition is always satisfied.

Proposition 13.6. *In the above-described expanding product variety model, provided that $\eta L > (\varepsilon - 1)\rho$, there exists a unique BGP, in which aggregate consumption expenditure, $C(t)$, and the wage rate, $w(t)$, grow at the rate g^* given by (13.52).*

A couple of features are worth noting. First, in this equilibrium, there is growth of "real income," even though the production function of each good remains unchanged. This is because, while there is no process innovation reducing costs or improving quality, the number of products available to consumers expands because of product innovations. Since the utility function of the representative household, (13.38), exhibits *love-for-variety*, the expanding variety of products increases utility. What happens to income depends on what we choose as the numeraire. The natural numeraire is the one setting the ideal price index, (13.44), equal to 1, which amounts to measuring incomes in similar units at different dates. With this

choice of numeraire, real incomes grow at the same rate as $C(t)$, at the rate g^*. Second, even though the equilibrium was characterized in a somewhat different manner than our baseline expanding input variety model, there is a close parallel between expanding product varieties and expanding input varieties. This can be seen, for example, in Exercise 13.23, which looks at an economy with expanding input varieties produced by labor. It can be verified that the structure of the equilibrium is very similar to the one studied here. Third, Exercise 13.26 will show that as in other models of endogenous technological progress in this chapter, there are no transitional dynamics and the equilibrium is again Pareto suboptimal. Moreover, log preferences now ensure that the transversality condition is always satisfied. Finally, it can be verified that there is again a scale effect here. This discussion then reveals that whether one wishes to use the expanding input variety or the expanding product model is mostly a matter of taste, and perhaps one of context. Both models lead to a similar structure of equilibria, to similar equilibrium growth rates and to similar welfare properties.

13.5. Taking Stock

In this chapter, we had our first look at models of endogenous technological progress. The distinguishing feature of these models is the fact that profit incentives shape R&D spending and investments, which in turn determines the rate at which the technology of the economy evolves over time. At some level, there are many parallels between the models studied here and the Romer (1986a) model of growth with externalities studied in Section 11.4 in Chapter 11; both have a mathematical structure similar to the neoclassical AK models (constant long-run growth rate, no transitional dynamics) and both generate externalities causing an equilibrium growth rate less than the Pareto optimal growth rate (because of physical capital externalities in the Romer (1986a) model, because of *aggregate demand externalities* in the lab-equipment model of Section 13.1 here, and because of a mixture of these in the other models studied in this chapter). The difference between the Romer (1986a) model and the endogenous technological change model should not be understated, however. While one may interpret the Romer (1986a) model as involving "knowledge accumulation," the accumulation of knowledge and technology is *not* an economic activity—it is a byproduct of other decisions (in this particular instance, individual physical capital accumulation decisions). Hence, while such a model may "endogenize" technology, it does so without explicitly specifying the costs and benefits of investing in new technologies. Since, as discussed in Chapter 3, technology differences across countries are likely to be important in accounting for their income differences, understanding the sources of technology differences is a major part of our effort to understand the mechanics of economic growth. In this respect, the models presented in this chapter constitute a major improvement over those presented so far.

The models studied in this chapter, like those of the previous chapter, emphasize the importance of profits in shaping technology choices. We have also seen the role of monopoly power and patent length on the equilibrium growth rate. In addition, the same factors that influenced the equilibrium growth rate in the neoclassical AK model also affect equilibrium

economic growth here. These include the discount rate, ρ, as well as taxes on capital income or corporate profits. Nevertheless, the effect of the market structure on equilibrium growth and innovation rates is somewhat limited in the current models because the Dixit-Stiglitz structure and expanding product or input varieties limit the extent to which firms can compete with each other. The models of quality competition in the next chapter will feature a richer interaction between market structure and equilibrium growth.

While the models in this chapter highlight certain major determinants of the rate of technological progress, another shortcoming of these models should be noted. The technology stock of a society is determined only by its own R&D. Thus technological differences will result simply from R&D differences. In our current world of relatively free knowledge-flows, many countries will not only generate technological know-how by their own R&D but will also benefit from the advances in the world technology frontier. Consequently, in practice, technology adoption decisions and the patterns of technology diffusion may be equally important as, or more important than, R&D towards the invention of new technologies (see Chapter 18 below). Therefore, the major contribution of the approaches studied in this chapter to our knowledge may be not in pinpointing the exact source of technology differences across countries, but in their emphasis on the endogenous nature of technology and in the perspective they provide for modeling technological investments.

In addition, even if technology adoption and imitation may be more important than innovation for the growth of some countries, models of endogenous technological change are essential for understanding world economic growth, since presumably the world technology frontier does largely advance because of R&D. Therefore, for our purpose of understanding that world economic growth, the new perspectives on the determinants of technological progress are important. Nevertheless, the AK structure of these models implies that they may have relatively little to say about why R&D investments and rapid technological progress has been a feature of the past 200 years and why the stock of knowledge and output per capita did not exhibit steady growth before the 19th century. Some of these questions will be addressed later in the book.

13.6. References and Literature

Models of endogenous technological progress were introduced in Romer (1987 and 1990), and then subsequently analyzed by, among others, Segerstrom, Anant and Dinopoulos (1990), Grossman and Helpman (1991a,b), Aghion and Howitt (1992). The lab-equipment model presented in Section 13.1 appears in Rivera-Batiz and Romer (1991). The model in Romer (1990) is similar to that presented in Section 13.2, but with skilled workers working in R&D. Gancia and Zilibotti (2005) provide an excellent survey of many of the models discussed in this chapter. Matsuyama (1995) gives a very lucid and informative discussion of the sources of inefficiency in Dixit-Stiglitz type models, which is related to the sources of inefficiency in the lab-equipment model presented in Section 13.1.

The critique of endogenous growth models because of scale effect is contained in Backus, Kehoe and Kehoe (1992) and in Jones (1995). The first of these papers pointed out that countries with larger sizes (either without adjustment or adjusted for international trade) do not grow faster in the postwar era. Jones (1995), on the other hand, focused on time-series patterns and pointed out the substantial increase in R&D inputs, for example, the total number of workers involved in research, with no corresponding increase in the equilibrium growth rate. Others argued that looking at the 20th century data may not be sufficient to reach a conclusion on whether there is a scale effect or not. Kremer (1993) argues, on the basis of estimates of world population, that there must have been an increase in economic growth over the past one million years. Laincz and Perreto (1996) argue that R&D resources allocated to specific product lines have not increased.

The model in Section 13.3 is similar to that presented in Jones (1995) and Jones (1999). As pointed out there, these models generate sustained growth of per capita income, but the growth rate of the economy does not respond to policies or preferences (given the rate of population growth). A number of authors have developed models of endogenous growth without scale effect, where policy might have an effect on the equilibrium growth rate. See, among others, Dinopoulos and Thompson (1998), Segerstrom (1998), Howitt (1999) and Young (1998). Aghion and Howitt (1998) and Ha and Howitt (2005) argue that semi-endogenous growth models along these lines also faced difficulties when confronted with the time-series evidence.

The model of expanding product varieties was first suggested by Judd (1985), but in the context of a model of exogenous growth. The endogenous growth models with expanding product variety is presented in Grossman and Helpman (1991a,b). The treatment here is somewhat different from that in Grossman and Helpman, especially because we used the ideal price index rather than Grossman and Helpman's choice of total expenditure as the numeraire.

13.7. Exercises

Exercise 13.1. This exercise asks you to derive (13.8) from (13.7).

(1) Rewrite (13.7) at time t as:

$$\begin{aligned} V(\nu,t) &= \int_t^{t+\Delta t} \exp\left[-\int_t^s r(\tau)\,d\tau\right](p^x(\nu,s)-\psi)\,x(\nu,s)ds \\ &\quad + \int_{t+\Delta t}^{\infty} \exp\left[-\int_{t+\Delta t}^s r(\tau)\,d\tau\right][p^x(\nu,s)x(\nu,s)-\psi x(\nu,s)]\,ds \end{aligned}$$

which is just an identity for any Δt. Interpret this equation and relate this to the *Principle of Optimality.*

(2) Show that for small Δt, this can be written as

$$V(\nu,t) = \Delta t\,(p^x(\nu,t)-\psi)\,x(\nu,t) + \exp(r(t)\,\Delta t)\,V(\nu,t+\Delta t) + o(\Delta t),$$

and thus derive the equation

$$\Delta t\,(p^x(\nu,t)-\psi)\,x(\nu,t) + \exp(r(t)\,\Delta t)\,V(\nu,t+\Delta t) - \exp(r(t)\times 0)\,V(\nu,t) + o(\Delta t) = 0,$$

where, recall that, $\exp(r(t) \times 0) = 1$. Interpret this equation and the significance of the term $o(\Delta t)$.

(3) Now divide both sides by Δt and take the limit $\Delta t \to 0$, to obtain

$$(p^x(\nu,t) - \psi)\, x(\nu,t) + \lim_{\Delta t \to 0} \frac{\exp(r(t)\,\Delta t)\, V(\nu, t+\Delta t) - \exp(r(t) \times 0)\, V(\nu,t)}{\Delta t} = 0.$$

(4) When the value function is differentiable in its time argument, the previous equations is equivalent to

$$(p^x(\nu,t) - \psi)\, x(\nu,t) + \left.\frac{\partial(\exp(r(t)\,\Delta t)\, V(\nu, t+\Delta t))}{\partial t}\right|_{\Delta t=0} = 0.$$

Now derive (13.8).

(5) Provide an economic intuition for the eq. (13.8).

Exercise 13.2. Derive (13.9) and (13.10) from the profit maximization problem of a monopolist.

Exercise 13.3. Formulate the consumer optimization problem in terms of the current-value Hamiltonian and derive the necessary conditions. Show that these are equivalent to (13.16) and (13.17).

Exercise 13.4. Prove that in the model of Section 13.1, if $\eta V(\nu, t') = 1$ and there is entry at some t', then $\eta V(\nu, t) = 1$ for all t and conversely that if $\eta V(\nu, t') < 1$ for some t', then $\eta V(\nu, t) \leq 1$ for all t and there is no entry at any t.

Exercise 13.5. Consider the expanding variety model of Section 13.1 and denote the BGP growth rates of consumption and total output by g_C^* and g^*.

(1) Show that $g_C^* > g^*$ is not feasible.

(2) Show that $g_C^* < g^*$ violates the transversality condition.

Exercise 13.6. This exercise asks you to construct and analyze the equivalent of the lab-equipment expanding variety model of Section 13.1 in discrete time. Suppose that the economy admits a representative household with preferences at time 0 given by

$$\sum_{t=0}^{\infty} \beta^t \frac{C(t)^{1-\theta} - 1}{1-\theta},$$

with $\beta \in (0,1)$ and $\theta \geq 0$. Production technology is the same as in the text and the innovation possibilities frontier of the economy is given by

$$N(t+1) - N(t) = \eta Z(t).$$

(1) Define the equilibrium in BGP allocations.

(2) Characterize the BGP and compare the structure of the equilibrium to that in Section 13.1.

(3) Show that there are no transitional dynamics, so that starting with any $N(0) > 0$, the economy grows at a constant rate.

Exercise 13.7. Complete the proof of Proposition 13.1, in particular, showing that condition (13.21) is sufficient for the transversality condition to be satisfied.

Exercise 13.8. Consider a world economy consisting of $j = 1, ..., M$ economies. Suppose that each of those are closed and have access to the same production and R&D technology as described in Section 13.1. The only differences across countries are in the size of labor force, L_j, productivity of R&D, η_j, and discount rate ρ_j. Also assume that one unit of R&D expenditure costs ζ_j units of final good in country j, and this varies across countries. There are no technological exchanges across countries.

(1) Define the "world equilibrium" in which each country is in equilibrium in analogy with the equilibrium path of the one country economy in Section 13.1.
(2) Characterize the world equilibrium. Show that in the world equilibrium, each country will grow at a constant rate starting at $t = 0$. Provide explicit solutionse for these growth rates.
(3) Show that except in "knife-edge" cases, output in each country will grow at a different long-run rate.
(4) Now return to the discussion in Chapters 3 and 8 regarding the effect of policy and taxes on long-run income per capita differences. Show that, in the model discussed in this exercise, arbitrarily small differences in policy or discount factors across countries will lead to "infinitely large" differences in long-run income per capita. Does this resolve the empirical challenges discussed in those chapters? Does the environment in this exercise provide a satisfactory model for the study of long-run income per capita differences across countries? If yes, please elaborate how such a model can be mapped to reality. If not, explain which features of the model appear unsatisfactory to you and how you would want to change them.

Exercise 13.9. (1) Verify that Theorem 7.14 from Chapter 7 can be applied to the social planner's problem in Section 13.1.
(2) Derive the consumption growth rate in the socially-planned economy, (13.22).

Exercise 13.10. Consider the expanding input variety model of Section 13.1. Show that it is possible for the equilibrium allocation to satisfy the transversality condition, while the social planner's solution may violate it. Interpret this result. Does it imply that the social planner's allocation is less compelling?

Exercise 13.11. Complete the proof of Proposition 13.3, in particular showing that the Pareto optimal allocation always involves a constant growth rate and no transitional dynamics.

Exercise 13.12. Consider the expanding input variety model of Section 13.1.

(1) Suppose that a benevolent government has access only to research subsidies, which can be financed by lump-sum taxes. Can these subsidies be chosen so as to ensure that the equilibrium growth rate is the same as the Pareto optimal growth rate? Can they be used to replicate the Pareto optimal equilibrium path? Would it be desirable for the government to use subsidies so as to achieve the Pareto optimal growth rate (from the viewpoint of maximizing social welfare at time $t = 0$)?

(2) Suppose that the government now has only access to subsidies to machines, which can again be financed by lump-sum taxes. Can these be chosen to induce the Pareto optimal growth rate? Can they be used to replicate the Pareto optimal equilibrium path?
(3) Will the combination of subsidies to machines and subsidies to research be better than either of these two policies by themselves?

Exercise 13.13. Consider the expanding input variety model of Section 13.1 and assume that corporate profits are taxed at the rate τ.

(1) Characterize the equilibrium allocation.
(2) Consider two economies with identical technologies and identical initial conditions, but with different corporate tax rates, τ and τ'. Determine the relative income of these two economies (possibly as a function of time).

Exercise 13.14. * Consider the expanding input variety model of Section 13.1, with one difference. A firm that invents a new machine receives a patent, which expires at the Poisson rate ι. Once the patent expires, that machine is produced competitively and is supplied to final good producers at marginal cost.

(1) Characterize the equilibrium in this case and show how the equilibrium growth rate depends on ι. [Hint: notice that there will be two different types of machines, supplied at different prices].
(2) What is the value of ι that maximizes the equilibrium rate of economic growth?
(3) Show that a policy of $\iota = 0$ does not necessarily maximize social welfare at time $t = 0$.

Exercise 13.15. Consider the formulation of competition policy in subsection 13.1.6.

(1) Characterize the equilibrium fully.
(2) Write down the welfare of the representative household at time $t = 0$ in this equilibrium.
(3) Maximize this welfare function by choosing a value of γ.
(4) Why is the optimal value of γ not equal to some $\gamma^* \geq 1/(1-\beta)$? Provide an interpretation in terms of the tradeoff between level and growth effects.
(5) What is the relationship between the optimal value of γ and ρ. Interpret.

Exercise 13.16. Complete the proof of Proposition 13.4. In particular, show that the equilibrium path involves no transitional dynamics and that under (13.31), the transversality condition is satisfied.

Exercise 13.17. Characterize the Pareto optimal allocation in the economy of Section 13.2. Show that it involves a constant growth rate greater than the equilibrium growth rate in Proposition 13.4 and no transitional dynamics.

Exercise 13.18. Derive eq. (13.35) and explain why the denominator is equal to $r^* - n$.

Exercise 13.19. Consider the model of endogenous technological progress with limited knowledge spillover as discussed in Section 13.3.

(1) Characterize the transitional dynamics of the economy starting from an arbitrary $N(0) > 0$.
(2) Characterize the Pareto optimal allocation and compare it to the equilibrium allocation in Proposition 13.5.
(3) Analyze the effect of the following two policies: first, a subsidy to research; second, a patent policy, where each patent expires at the rate $\iota > 0$. Explain why the effects of these policies on economic growth are different than their effects in the baseline endogenous growth model.

Exercise 13.20. Consider the model in Section 13.3. Suppose that there are two economies with identical preferences, technology and initial conditions, except country 1 starts with population $L_1(0)$ and country 2 starts with $L_2(0) > L_1(0)$. Show that income per capita is always higher in country 2 than in country 1.

Exercise 13.21. Consider the lab-equipment model of Section 13.1, but modify the innovation possibilities frontier to

$$\dot{N}(t) = \eta N(t)^{-\phi} Z(t),$$

where $\phi > 0$.

(1) Define an equilibrium.
(2) Characterize the market clearing factor prices and determine the free-entry condition.
(3) Show that without population growth, there will be no sustained growth in this economy.
(4) Now consider population growth at the exponential rate n, and show that this model generates sustained equilibrium growth as in the model analyzed in Section 13.3.

Exercise 13.22. Consider the baseline endogenous technological change model with expanding machine varieties in Section 13.1. Suppose that x's now denote machines that do not immediately depreciate. In contrast, once produced these machines depreciate as an exponential rate δ. Preferences and the rest of the production structure remain unchanged.

(1) Define the equilibrium in BGP allocations.
(2) Formulate the maximization problem of producers of machines. [Hint: it is easier to formulate the problem in terms of machine rentals rather than machine sales].
(3) Characterize the equilibrium in this economy and show that all the results are identical to those in Section 13.1.

Exercise 13.23. Consider the following model. Population at time t is $L(t)$ and grows at the constant rate n (i.e., $\dot{L}(t) = nL(t)$). All agents have preferences given by

$$\int_0^\infty \exp(-\rho t) \frac{C(t)^{1-\theta} - 1}{1-\theta} dt,$$

where C is consumption defined over the final good of the economy. This good is produced as

$$Y(t) = \left[\int_0^N y(\nu,t)^\beta \, d\nu\right]^{1/\beta},$$

where $y(\nu,t)$ is the amount of intermediate good ν used in production at time t. The production function of each intermediate is

$$y(\nu,t) = l(\nu,t)$$

where $l(\nu,t)$ is labor allocated to this good at time t.

New goods are produced by allocating workers to the R&D process, with the production function

$$\dot{N}(t) = \eta N^\phi(t) L_R(t)$$

where $\phi \leq 1$ and $L_R(t)$ is labor allocated to R&D at time t. So labor market clearing requires $\int_0^{N(t)} l(\nu,t)\, dv + L_R(t) = L(t)$. Risk-neutral firms hire workers for R&D. A firm who discovers a new good becomes the monopoly supplier, with a perfectly and indefinitely enforced patent.

(1) Characterize the BGP in the case where $\phi = 1$ and $n = 0$, and show that there are no transitional dynamics. Why is this? Why does the long-run growth rate depend on θ? Why does the growth rate depend on L? Do you find this plausible?
(2) Now suppose that $\phi = 1$ and $n > 0$. What happens? Interpret.
(3) Now characterize the BGP when $\phi < 1$ and $n > 0$. Does the growth rate depend on L? Does it depend on n? Why? Do you think that the configuration $\phi < 1$ and $n > 0$ is more plausible than the one with $\phi = 1$ and $n = 0$?

EXERCISE 13.24. Derive eq. (13.43). [Hint: use the first-order condition between two products ν and ν', and then substitute into the budget constraint of the representative household with total expenditure denoted by $C(t)$].

EXERCISE 13.25. Using (13.43) and the choice of numeraire in (13.44), set up the consumer optimization problem in the form of the current-value Hamiltonian (derive the consumer's budget constraint explicitly). Derive the consumer Euler equation (13.45).

EXERCISE 13.26. Consider the model analyzed in Section 13.4.

(1) Show that the allocation described in Proposition 13.6 always satisfies the transversality condition.
(2) Show that in this model there are no transitional dynamics.
(3) Characterize the Pareto optimal allocation and show that the equilibrium growth rate in Proposition 13.6 is less than the growth rate in the Pareto optimal allocation.

CHAPTER 14

Models of Schumpeterian Growth

The previous chapter presented the basic endogenous technological change models based on expanding input or product varieties. The advantage of these models is their relative tractability. While the expansion of the variety of machines used in production captures certain aspects of the economics of innovation, most process innovations in practice either increase the quality of an existing product or reduce the costs of production. Therefore, typical process innovations have a number of distinct features compared to the "horizontal innovations" of the previous chapter. For example, in the expanding machine variety model a newly-invented computer is used alongside all previous vintages of computers, though, in reality, a newly-invented computer often *replaces* existing vintages. Thus in some fundamental sense, models of expanding machine variety do not provide a good description of innovation dynamics in practice because they do not capture the *competitive* aspect of innovations. These competitive aspects bring us to the realm of Schumpeterian *creative destruction* in which economic growth is driven, at least in part, by new firms replacing incumbents, and new machines and products replacing old ones. For this reason, the models discussed in this chapter are often referred to as *Schumpeterian growth models*. My purpose in this chapter is to develop tractable models of Schumpeterian growth.

As Chapter 12 discussed, innovations that involve quality improvements or cost reductions will feature the replacement effect, which implies that entrants should be more active in the research process than incumbents. Schumpeterian growth therefore raises a number of novel and important issues. First, in contrast to the models of expanding varieties, there may be direct price competition between different producers with different vintages of quality or different costs of producing the same product. This will affect both the description of the growth process and a number of its central implications. For example, market structure and anti-trust policy can play potentially richer roles in models exhibiting this type of price competition. Second, competition between incumbents and entrants brings the business stealing effect discussed in Chapter 12 to the fore and raises the possibility of excessive innovations.

This description suggests that a number of new and perhaps richer issues arise when in the context of Schumpeterian growth models. One may then expect models of Schumpeterian models to be significantly more complicated than expanding varieties models. This is not necessarily the case, however. This chapter presents the basic models of competitive innovations, first proposed by Aghion and Howitt (1992) and then further developed by Grossman

and Helpman (1991a,b) and Aghion and Howitt (1998). The literature on models of Schumpeterian economic growth is now large and an excellent survey is presented in Aghion and Howitt (1998). My purpose here is not to provide a detailed survey, but to emphasize the most important implications of these models. I will also present these models in a way that parallels the mathematical structure of the expanding varieties models, both to emphasize the similarities and to clarify the differences. A number of distinct applications of these models are also discussed later in the chapter and in the exercises.

14.1. A Baseline Model of Schumpeterian Growth

14.1.1. Preferences and Technology. The economy is in continuous time and admits a representative household with the standard CRRA preferences, (13.1), as in the previous chapter. Population is constant at L and labor is supplied inelastically. The resource constraint at time t again takes the form

$$C(t) + X(t) + Z(t) \leq Y(t), \tag{14.1}$$

where $C(t)$ is consumption, $X(t)$ is aggregate spending on machines, and $Z(t)$ is total expenditure on R&D at time t.

There is a continuum of machines used in the production of a unique final good. Since there will be no expansion of inputs/machine varieties, the measure of inputs can be normalized to 1 without loss of any generality. Each machine line is denoted by $\nu \in [0,1]$. The engine of economic growth here will be process innovations that lead to *quality improvements*. Let us first specify how the qualities of different machine lines change over time. Let $q(\nu,t)$ be the quality of machine line ν at time t. The following "quality ladder" determines the quality of each machine type:

$$q(\nu,t) = \lambda^{n(\nu,t)} q(\nu,0) \text{ for all } \nu \text{ and } t, \tag{14.2}$$

where $\lambda > 1$ and $n(\nu,t)$ denotes the number of innovations on this machine line between time 0 and time t. This specification implies that there is a ladder of quality for each machine type, and each innovation takes the machine quality up by one rung in this ladder. These rungs are proportionally equi-distant, so that each improvement leads to a proportional increase in quality by an amount $\lambda > 1$. Growth will be the result of these quality improvements.

The production function of the final good is similar to that in the previous chapter, except that now the quality of the machines matters for productivity. The aggregate production function takes the form

$$Y(t) = \frac{1}{1-\beta}\left[\int_0^1 q(\nu,t) x(\nu,t \mid q)^{1-\beta} d\nu\right] L^{\beta}, \tag{14.3}$$

where $x(\nu,t \mid q)$ is the quantity of the machine of type ν of quality q used in the production process. As in the previous chapter, this production function can be written as

$$Y(t) = \frac{1}{1-\beta}\tilde{\mathbf{X}}(t)^{1-\beta} L^{\beta},$$

where

$$\tilde{\mathbf{X}}(t) \equiv \left[\int_0^1 q(\nu,t)x(\nu,t \mid q)^{\frac{\varepsilon_\beta-1}{\varepsilon_\beta}}d\nu\right]^{\frac{\varepsilon_\beta}{\varepsilon_\beta-1}},$$

with $\varepsilon_\beta \equiv 1/\beta$, which again emphasizes the continuity with the Dixit-Stiglitz model of Chapter 12.

An implicit assumption in (14.3) is that at any point in time only one quality of any machine is used. This is without loss of any generality, since in equilibrium only the highest-quality machines of each type will be used, and this will be the source of *creative destruction*; when a higher-quality machine is invented it will replace ("destroy") the previous vintage of of the same machine.

I next specify the technology for producing machines of different qualities and the innovation possibilities frontier of this economy. New machine vintages are invented by R&D. The R&D process is cumulative, in the sense that new R&D builds on the knowhow of existing machines. For example, consider the machine line ν that has quality $q(\nu,t)$ at time t. R&D on this machine line will attempt to improve over this quality. If a firm spends $Z(\nu,t)$ units of the final good for research on this machine line, then it generates a flow rate $\eta Z(\nu,t)/q(\nu,t)$ of innovation. The innovation advances the knowhow of the production of this machine to the new rung of the quality ladder, thus creates a machine of type ν with quality $\lambda q(\nu,t)$. Note that one unit of R&D spending is proportionately less effective when applied to a more advanced machine. This is intuitive, since we expect research on more advanced machines to be more difficult. It is also convenient from a mathematical point of view, since the benefit of research is also increasing with the quality of the machine (in particular, the quality improvements are *proportional*, with an innovation increasing quality from $q(\nu,t)$ to $\lambda q(\nu,t)$). Note also that the costs of R&D are identical for the current incumbent and new firms (see Exercise 14.7 for alternative formulations). There is again free entry into research, thus any firm or individual can undertake this type of research on any of the machine lines. As in the expanding varieties models of the previous chapter, the firm that makes an innovation has a perpetual patent on the new machine that it has invented. But the patent system does not preclude other firms undertaking research based on the machine (or product) invented by this firm. How different patenting arrangements might affect incentives will be discussed below.

Once a particular machine of quality $q(\nu,t)$ has been invented, any quantity of this machine can be produced at the marginal cost $\psi q(\nu,t)$. Once again, the fact that the marginal cost is proportional to the quality of the machine is natural, since producing higher-quality machines should be more expensive.

One noteworthy issue here concerns the identity of the firm that will undertake R&D and innovation. In the expanding varieties model, this was irrelevant, since machines could not be improved upon, so there was only R&D for new machines, and who undertook the R&D was not important. Here, in contrast, existing machines can be (and are) improved, and this is the source of economic growth. We have already seen in Chapter 12 that if the

costs of R&D are identical for incumbents and new firms, Arrow's replacement effect will imply that it will be the new entrants that undertake the R&D. The same applies in this model. The incumbent has weaker incentives to innovate, since it would be replacing its own machine, and thus destroying the profits that it is already making. In contrast, a new entrant does not have this replacement calculation in mind. As a result, with the same technology of innovation, it will always be the entrants that undertake the R&D investments in this model (see Exercise 14.2). This is an attractive implication, since it creates a real sense of creative destruction or churning. Of course in practice we observe established and leading firms undertaking innovations. This might be because the technology of innovation differs between incumbents and new potential entrants, or because there is only a limited number of new entrants as in the model studied in Section 14.4 below.

14.1.2. Equilibrium. An allocation in this economy is similar to that in the previous chapter. It consists of time paths of consumption levels, aggregate spending on machines, and aggregate R&D expenditure, $[C(t), X(t), Z(t)]_{t=0}^{\infty}$, time paths of machine qualities denoted by, $[q(\nu,t)]_{\nu\in[0,1],t=0}^{\infty}$, time paths of prices and quantities of each machine and the net present discounted value of profits from that machine, $[p^x(\nu,t \mid q), x(\nu,t), V(\nu,t \mid q)]_{\nu\in[0,1],t=0}^{\infty}$, and time paths of interest rates and wage rates, $[r(t), w(t)]_{t=0}^{\infty}$. An equilibrium is defined below.

Let us start with the aggregate production function for the final good producers. A similar analysis to that in the previous chapter implies that the demand for machines is given by

$$x(\nu,t \mid q) = \left(\frac{q(\nu,t)}{p^x(\nu,t \mid q)}\right)^{1/\beta} L \quad \text{for all } \nu \in [0,1] \text{ and all } t, \tag{14.4}$$

where $p^x(\nu,t \mid q)$ refers to the price of machine type ν of quality $q(\nu,t)$ at time t. This expression stands for $p^x(\nu,t \mid q(\nu,t))$, but there should be no confusion in this notation since it is clear that q here refers to $q(\nu,t)$, and I will use this notation for other variables as well. The price $p^x(\nu,t \mid q)$ will be determined by the profit-maximization of the monopolist holding the patent for machine of type ν of quality $q(\nu,t)$. Note that the demand from the final good sector for machines in (14.4) is isoelastic as in the previous chapter, so the unconstrained monopoly price is again a constant markup over marginal cost. However, contrary to the situation in the previous chapter, there is now competition between firms that have access to different vintages of the machine. This implies that, as in Chapter 12, we need to consider two regimes, one in which the innovation is "drastic" so that each firm can charge the unconstrained monopoly price, and the other one in which limit prices have to be used. Which regime we are in does not make any difference to the mathematical structure or to the substantive implications of the model. Nevertheless, we have to choose one of these two alternatives for consistency. Here let us assume that the quality gap between a new machine and the machine that it replaces, λ, is sufficiently large. In particular, it satisfies

$$\lambda \geq \left(\frac{1}{1-\beta}\right)^{\frac{1-\beta}{\beta}}, \tag{14.5}$$

so that we are in the drastic innovations regime (see Exercise 14.9 for the derivation of this condition and Exercise 14.10 for the structure of the equilibrium under the alternative assumption). Let us also normalize $\psi \equiv 1-\beta$ as in the previous chapter, which implies that the profit-maximizing monopoly price is

$$p^x(\nu,t \mid q) = q(\nu,t). \tag{14.6}$$

Combining this with (14.4) implies

$$x(\nu,t \mid q) = L. \tag{14.7}$$

Consequently, the flow profits of a firm with the monopoly rights on the machine of quality $q(\nu,t)$ can be computed as:

$$\pi(\nu,t \mid q) = \beta q(\nu,t) L.$$

This only differs from the flow profits in the previous chapter because of the presence of the quality term, $q(\nu,t)$. Next, substituting (14.7) into (14.3), total output is given by

$$Y(t) = \frac{1}{1-\beta} Q(t) L, \tag{14.8}$$

where

$$Q(t) = \int_0^1 q(\nu,t) d\nu \tag{14.9}$$

is the average total quality of machines. This expression closely parallels the derived production function (13.12) in the previous chapter, except that labor productivity is now determined by the average quality of the machines, $Q(t)$ (instead of the number of machine varieties, $N(t)$, in the previous chapter). This expression also clarifies the reasoning for the particular functional form assumptions above. In particular, the reader can verify that it is the linearity of the aggregate production function of the final good, (14.3), in the quality of machines that makes labor productivity depend on average qualities. With alternative assumptions, a similar expression to (14.8) would still obtain, but with a more complicated aggregator of machine qualities than the simple average (see, for example, Section 14.4). Next, aggregate spending on machines is immediately obtained as

$$X(t) = (1-\beta) Q(t) L. \tag{14.10}$$

Similar to the previous chapter, labor, which is only used in the final good sector, receives an equilibrium wage rate of

$$w(t) = \frac{\beta}{1-\beta} Q(t). \tag{14.11}$$

I next specify the value function for the monopolist of variety ν of quality $q(\nu,t)$ at time t. As before, despite the fact that each firm generates a stochastic stream of revenues, the presence of many firms with independent "draws" implies that each should maximize expected profits. The net present value of expected profits can be written in the Hamilton-Jacobi-Bellman form as follows

$$r(t) V(\nu,t \mid q) - \dot{V}(\nu,t \mid q) = \pi(\nu,t \mid q) - z(\nu,t \mid q) V(\nu,t \mid q), \tag{14.12}$$

where $z(\nu, t \mid q)$ is the rate at which new innovations occur in sector ν at time t, while $\pi(\nu, t \mid q)$ is the flow of profits. This value function is somewhat different from the ones in the previous chapter (e.g., (13.8)), because of the last term on the right-hand side, which captures the essence of Schumpeterian growth. When a new innovation occurs, the existing monopolist loses its monopoly position and is replaced by the producer of the higher-quality machine. From then on, it receives zero profits, and thus has zero value. In writing this equation, I made use of the fact that because of Arrow's replacement effect, it is an entrant that is undertaking the innovation, thus $z(\nu, t \mid q)$ corresponds to the flow rate at which the incumbent will be replaced by a new entrant.

Free entry again implies

$$\eta V(\nu, t \mid q) \leq \lambda^{-1} q(\nu, t) \text{ and } \eta V(\nu, t \mid q) = \lambda^{-1} q(\nu, t) \text{ if } Z(\nu, t \mid q) > 0. \tag{14.13}$$

In other words, the value of spending one unit of the final good should not be strictly positive. Recall that one unit of the final good spent on R&D for a machine of quality $\lambda^{-1} q$ has a flow rate of success equal to $\eta / \left(\lambda^{-1} q\right)$, and in this case, it generates a new machine of quality q, which will have a net present value gain of $V(\nu, t \mid q)$. If there is positive R&D, that is, if $Z(\nu, t \mid q) > 0$, then the free-entry condition must hold as equality.

Note also that even though the quality of individual machines, the $q(\nu, t)$'s, are stochastic (and depend on success in R&D), as long as R&D expenditures, the $Z(\nu, t \mid q)$'s, are nonstochastic, average quality $Q(t)$, and thus total output, $Y(t)$, and total spending on machines, $X(t)$, will be nonstochastic. This feature will significantly simplify the notation and the analysis.

Consumer maximization again implies the familiar Euler equation,

$$\frac{\dot{C}(t)}{C(t)} = \frac{1}{\theta}(r(t) - \rho), \tag{14.14}$$

and the transversality condition takes the form

$$\lim_{t \to \infty} \left[\exp\left(- \int_0^t r(s)\, ds \right) \int_0^1 V(\nu, t \mid q)\, d\nu \right] = 0 \tag{14.15}$$

for all q. This transversality condition is intuitive, since the total value of corporate assets is now $\int_0^1 V(\nu, t \mid q)\, d\nu$. Even though the evolution of the quality of each machine line is stochastic, the value of a machine of type ν of quality q at time t, $V(\nu, t \mid q)$, is nonstochastic. Either q is not the highest quality in this machine line, in which case $V(\nu, t \mid q)$ is equal to 0, or alternatively, it is given by (14.12).

These equations complete the description of the environment. An equilibrium can then be represented as time paths of consumption, aggregate spending on machines, and aggregate R&D, $[C(t), X(t), Z(t)]_{t=0}^{\infty}$ that satisfy (14.1), (14.10), (14.15), time paths of aggregate machine quality $[Q(t)]_{t=0}^{\infty}$ and value functions $[V(\nu, t \mid q)]_{\nu \in [0,1], t=0}^{\infty}$ consistent with (14.9), (14.12) and (14.13), time paths of prices and quantities of machines that have highest quality in their lines at that point in time and the net present discounted value of profits from those

machines, $[p^x(\nu,t \mid q), x(\nu,t)]_{\nu\in[0,1],t=0}^{\infty}$ given by (14.6) and (14.7), and time paths of interest rates and wage rates, $[r(t), w(t)]_{t=0}^{\infty}$ that are consistent with (14.11) and (14.14).

I again start with BGP, where output and consumption grow at constant rates.

14.1.3. Balanced Growth Path. In the BGP, consumption grows at the constant rate g_C^*. With familiar arguments, this must be the same rate as output growth, g^*. Moreover, from (14.14), the interest rate must be constant, that is, $r(t) = r^*$ for all t.

If there is positive growth in this BGP equilibrium, then there must be research at least in some sectors. Since both profits and R&D costs are proportional to quality, whenever the free-entry condition (14.13) holds as equality for one machine type, it will hold as equality for all of them. This, in turn, implies that

$$V(\nu,t \mid q) = \frac{q(\nu,t)}{\lambda\eta}. \tag{14.16}$$

Moreover, if this condition holds between t and $t+\Delta t$, then $\dot{V}(\nu,t \mid q) = 0$, because the right-hand side of eq. (14.16) is constant over time—$q(\nu,t)$ refers to the quality of the machine supplied by the incumbent, which does not change. In BGP, R&D expenditures, $z(\nu,t)$, must be constant and the same for all machine types, and I denote this common value by z^* (see Exercise 14.1). Then, (14.12) implies

$$V(\nu,t \mid q) = \frac{\beta q(\nu,t) L}{r^* + z^*}. \tag{14.17}$$

Notice the difference between this value function and those in the previous chapter: instead of the discount rate r^*, the denominator has the effective discount rate, $r^* + z^*$, because incumbent monopolists understand that new innovations will replace them.

Combining this equation with (14.16),

$$r^* + z^* = \lambda\eta\beta L. \tag{14.18}$$

Moreover, from the fact that $g_C^* = g^*$ and (14.14), $g^* = (r^* - \rho)/\theta$, or

$$r^* = \theta g^* + \rho. \tag{14.19}$$

To solve for the BGP equilibrium, we need a final equation relating the BGP growth rate of the economy, g^*, to z^*. From (14.8), this is given by

$$\frac{\dot{Y}(t)}{Y(t)} = \frac{\dot{Q}(t)}{Q(t)}.$$

Note that, by definition, in an interval of time Δt, there will be $z(t)\Delta t$ sectors that experience one innovation, and this will increase their productivity by λ. The measure of sectors experiencing more than one innovation within this time interval is $o(\Delta t)$—i.e., it is second-order in Δt, so that as $\Delta t \to 0$, $o(\Delta t)/\Delta t \to 0$. Therefore,

$$Q(t+\Delta t) = \lambda Q(t) z(t) \Delta t + (1 - z(t)\Delta t) Q(t) + o(\Delta t).$$

Now subtracting $Q(t)$ from both sides, dividing by Δt and taking the limit as $\Delta t \to 0$,

$$\dot{Q}(t) = (\lambda - 1) z(t) Q(t).$$

Therefore,

$$g^* = (\lambda - 1) z^*. \tag{14.20}$$

Now combining (14.18)-(14.20), the BGP growth rate of output and consumption is:

$$g^* = \frac{\lambda \eta \beta L - \rho}{\theta + (\lambda - 1)^{-1}}. \tag{14.21}$$

This establishes the following proposition.

PROPOSITION 14.1. *Consider the model of Schumpeterian growth described above. Suppose that*

$$\lambda \eta \beta L > \rho > (1 - \theta) \lambda \eta \beta L. \tag{14.22}$$

Then, there exists a unique BGP in which average quality of machines, output and consumption grow at rate g^* given by (14.21). The rate of innovation is $g^*/(\lambda - 1)$.

PROOF. Most of the proof is given in the preceding analysis. In Exercise 14.4 you are asked to check that the BGP equilibrium is unique and satisfies the transversality condition. □

The above analysis illustrates that the mathematical structure of the model is quite similar to those analyzed in the previous chapter. Nevertheless, the feature of creative destruction, the process of incumbent monopolists being replaced by new entrants, is new and provides a very different interpretation of the growth process. I return to some of the applications of creative destruction below.

Before doing this, let us briefly look at transitional dynamics in this economy. Similar arguments to those used in the previous chapter establish the following result:

PROPOSITION 14.2. *In the model of Schumpeterian growth described above, starting with any average quality of machines $Q(0) > 0$, there are no transitional dynamics and the equilibrium path always involves constant growth at the rate g^* given by (14.21).*

PROOF. See Exercise 14.5. □

A notable feature of the model, which is again related to the functional form of the aggregate production function (14.3), is that only the average quality of machines, $Q(t)$, matters for the allocation of resources. Moreover, the incentives to undertake research are identical for two machine types ν and ν', with different quality levels $q(\nu, t)$ and $q(\nu', t)$. Thus there are no differential incentives to undertake different R&D investments for more and less advanced machines. This is again a feature of the functional forms chosen here, and Exercise 14.14 shows that in different circumstances this result may not apply. Nevertheless, the specification chosen in this section is appealing, since it implies that research will be directed at a broad range of machines rather than a specific subset of the available types of machines.

14.1.4. Pareto Optimality. This equilibrium, like that of the endogenous technology model with expanding varieties, is typically Pareto suboptimal. The first reason for this is the appropriability effect, which results because monopolists are not able to capture the entire social gain created by an innovation. However, Schumpeterian growth also introduces the business stealing effect discussed in Chapter 12. Consequently, the equilibrium rate of innovation and growth can now be too high or too low. I now investigate this question.

I proceed as in the previous chapter, first deriving quantities of machines that will be used in the final good sector by the social planner. In the social planner's allocation there is no markup on machines, thus

$$\begin{aligned} x^S(\nu, t \mid q) &= \psi^{-1/\beta} L \\ &= (1-\beta)^{-1/\beta} L. \end{aligned}$$

Substituting this into (14.3),

$$Y^S(t) = (1-\beta)^{-1/\beta} Q^S(t) L,$$

where again the superscript S refers to the social planner's allocation. The net output that can be distributed between consumption and research expenditure is

$$\begin{aligned} \tilde{Y}^S(t) &\equiv Y^S(t) - X^S(t) \\ &= (1-\beta)^{-1/\beta} Q^S(t) L - \int_0^1 \psi q(\nu, t) x^S(\nu, t \mid q) d\nu \\ &= (1-\beta)^{-1/\beta} \beta Q^S(t) L. \end{aligned} \tag{14.23}$$

Moreover, given the specification of the innovation possibilities frontier above, the social planner can improve the aggregate technology as follows:

$$\dot{Q}^S(t) = \eta(\lambda - 1) Z^S(t),$$

because R&D spending of $Z^S(t)$ leads to the discovery of better vintages at the flow rate η and each new vintage increases average quality by a proportional amount $\lambda - 1$.

Now, given this equation, the maximization problem of the social planner can be written as

$$\max \int_0^\infty \exp(-\rho t) \frac{C^S(t)^{1-\theta} - 1}{1-\theta} dt$$

subject to

$$\dot{Q}^S(t) = \eta(\lambda - 1)(1-\beta)^{-1/\beta} \beta Q^S(t) L - \eta(\lambda - 1) C^S(t),$$

where the constraint equation uses net output, (14.23), and the resource constraint, (14.1). In this problem, $Q^S(t)$ is the state variable, and $C^S(t)$ is the control variable. To characterize this solution, let us set up the current-value Hamiltonian as

$$\hat{H}\left(Q^S, C^S, \mu^S\right) = \frac{C^S(t)^{1-\theta} - 1}{1-\theta} + \mu^S(t) \left[\eta(\lambda - 1)(1-\beta)^{-1/\beta} \beta Q^S(t) L - \eta(\lambda - 1) C^S(t)\right].$$

The necessary conditions for a maximum are

$$\begin{aligned}\hat{H}_C\left(Q^S, C^S, \mu^S\right) &= C^S(t)^{-\theta} - \mu^S(t)\,\eta(\lambda-1) = 0\\ \hat{H}_Q\left(Q^S, C^S, \mu^S\right) &= \mu^S(t)\,\eta(\lambda-1)(1-\beta)^{-1/\beta}\beta L = \rho\mu^S(t) - \dot{\mu}^S(t)\\ \lim_{t\to\infty}\left[\exp(-\rho t)\,\mu^S(t)\,Q^S(t)\right] &= 0.\end{aligned}$$

Moreover, it is straightforward to verify that the current-value Hamiltonian is concave in C and Q, so that the sufficiency conditions in Theorem 7.14 are satisfied and the solution to these conditions gives the unique optimal plan (see Exercise 14.8). Combining these conditions, we obtain the following growth rate for consumption in the social planner's allocation (see again Exercise 14.8):

$$\frac{\dot{C}^S(t)}{C^S(t)} = g^S \equiv \frac{1}{\theta}\left(\eta(\lambda-1)(1-\beta)^{-1/\beta}\beta L - \rho\right). \tag{14.24}$$

Clearly, total output and average quality will also grow at the rate g^S in this allocation.

Comparing g^S to g^* in (14.21), we can see that either could be greater. In particular, when λ is very large, $g^S > g^*$, and there is insufficient growth in the equilibrium. For example, as $\lambda \to \infty$, we have $g^S/g^* \to (1-\beta)^{-1/\beta} > 1$. In contrast, to obtain an example in which there is excessive growth in equilibrium, suppose that $\theta = 1$, $\beta = 0.9$, $\lambda = 1.3$, $\eta = 1$, $L = 1$ and $\rho = 0.38$. In this case, it can be verified that $g^S \approx 0$, while $g^* \approx 0.18 > g^S$.[1]

This illustrates the counteracting influences of the appropriability and business stealing effects discussed above. The following proposition summarizes this result:

PROPOSITION 14.3. *In the model of Schumpeterian growth described above, the decentralized equilibrium is generally Pareto suboptimal and may have a higher or lower rate of innovation and growth than the Pareto optimal allocation.*

It is also straightforward to verify that as in the models of the previous section, there is a scale effect, and thus population growth would lead to an exploding growth path. Exercise 14.11 asks you to construct a Schumpeterian growth model without scale effects.

14.1.5. Policy in the Model of Schumpeterian Growth. I now use the Schumpeterian growth model to analyze the effects of policy on economic growth. As in the model of the previous two chapters, anti-trust policy, patent policy and taxation will affect the equilibrium growth rate. For example, two economies that tax corporate incomes at different rates, say τ and τ', will grow at different rates.

There is a sense in which the current model is much more appropriate for conducting policy analysis than the expanding varieties models, however. In those models, there was no reason for any agent in the economy to support distortionary taxes.[2] In contrast, the fact that

[1] Notice that the combination of $\beta = 0.9$ and $\lambda = 1.3$ are consistent with (14.5), which was used in deriving the equilibrium growth rate g^*.

[2] Naturally, one can enrich these models so that tax revenues are distributed unequally across agents, for example, with taxes on capital distributed to workers. In this case, even in the basic neoclassical growth model, some groups could prefer distortionary taxes. Such models will be discussed in Part 8 of the book.

growth takes place through creative destruction here implies that there is a natural conflict of interest, and certain types of distortionary policies may have a natural constituency. To illustrate this point, which will be discussed in greater detail in Part 8 of the book, suppose that there is a tax τ imposed on R&D spending. This has no effect on the profits of existing monopolists, and only influences their net present discounted value via replacement. Since taxes on R&D will discourage R&D, there will be replacement at a slower rate, that is, BGP R&D expenditure z^* will fall. This directly increases the steady-state value of all monopolists given by (14.17). In particular, denoting the value of a monopolist with a machine of quality q by $V(q)$, we have

$$V(q) = \frac{\beta q L}{r^*(\tau) + z^*(\tau)},$$

where the equilibrium interest rate and the replacement rate have been written as functions of τ. With the tax rate on R&D, the free-entry condition, (14.13) becomes

$$V(q) = \frac{(1+\tau)}{\lambda \eta} q.$$

This equation shows that $V(q)$ is clearly increasing in the tax rate on R&D, τ. Combining the previous two equations, it can be seen that in response to a positive rate of taxation, $r^*(\tau) + z^*(\tau)$ must adjust downward, so that the value of current monopolists increases (consistent with the previous equation). Intuitively, when the costs of R&D are raised because of tax policy, the value of a successful innovation, $V(q)$, must increase to satisfy the free-entry condition. This can only happen through a decline in the effective discount rate $r^*(\tau)+z^*(\tau)$. A lower effective discount rate, in turn, is achieved by a decline in the equilibrium growth rate of the economy, which now takes the form

$$g^*(\tau) = \frac{(1+\tau)^{-1} \lambda \eta \beta L - \rho}{\theta + (\lambda - 1)^{-1}}.$$

This growth rate is indeed strictly decreasing in τ. Nevertheless, as the previous expression shows, incumbent monopolists directly benefit from an increase in τ and would be in favor of such a policy. Essentially, in this model, slowing down the process of creative destruction is beneficial for incumbents, creating a rationale for growth-retarding policies to emerge in equilibrium.

Therefore, an important advantage of models of Schumpeterian growth is that they start providing us clues about why some societies may adopt policies that reduce the growth rate. Since taxing R&D by new entrants benefits incumbent monopolists, when incumbents are politically powerful, such distortionary taxes can emerge in the political economy equilibrium, even though they are not in the interest of the society at large.

14.2. A One-Sector Schumpeterian Growth Model

The model of Schumpeterian growth presented in the previous section was designed to maximize the parallels between this class of models and those based on expanding varieties. In this section, I discuss a model more closely related to the original Aghion and Howitt

(1992) setup, which is simpler in some ways and more complicated in others. There are two major differences from the model presented in the previous section. First, there is only one sector experiencing quality improvements rather than a continuum of machine types. Second, the innovation possibilities frontier uses a scarce factor, labor, as in the model of knowledge spillovers in Section 13.2 of the previous chapter. Since there are many parallels between this model and the one in the previous section, I only provide a brief exposition.

14.2.1. The Basic Aghion-Howitt Model. The consumer side is the same as before, with the only difference that we now assume consumers are risk neutral, so that the interest rate is determined as

$$r^* = \rho$$

at all points in time. Population is again constant at a level L and labor is supplied inelastically. The aggregate production function of the unique final good is now given by

$$Y(t) = \frac{1}{1-\beta} x(t \mid q)^{1-\beta} \left(q(t) L_E(t)\right)^{\beta}, \tag{14.25}$$

where $q(t)$ is the quality of the unique machine used in production and is written in the labor-augmenting form for simplicity; $x(t \mid q)$ is the quantity of this machine used at time t; and $L_E(t)$ denotes the amount labor used in production at time t, which is less than L, since $L_R(t)$ workers will be employed in the R&D sector. Market clearing requires that

$$L_E(t) + L_R(t) \leq L.$$

Once invented, a machine of quality $q(t)$ can be produced at the constant marginal cost ψ in terms of final goods. Let us again normalize $\psi \equiv 1 - \beta$. The innovation possibilities frontier now involves labor being used for R&D. In particular, each worker employed in the R&D sector generates a flow rate η of a new machine. When the current machine used in production has quality $q(t)$, the new machine has quality $\lambda q(t)$.

Let us once again assume that (14.5) above is satisfied, so that the monopolist can charge the unconstrained monopoly price. Then, an analysis similar to that in the previous section immediately implies that the demand for the leading-edge machine of quality q is given by

$$x(t \mid q) = p^x(t)^{-1/\beta} q(t) L_E(t),$$

where again $p^x(\nu, t)$ denotes the price of the machine of quality q. The monopoly price for the highest quality machine is:[3]

$$p^x(t \mid q) = \frac{\psi}{1-\beta} = 1,$$

[3]This expression follows the original Aghion and Howitt (1992) paper and assumes that innovators ignore their effect on equilibrium wages. In the baseline model presented above, since there was a continuum of monopolistically competitive firms, each had no effect on equilibrium wages. Here, instead, the single monopolist could recognize that its price will also affect equilibrium wages and thus the cost of R&D for its competitors. In this case, it may want to set a lower price than $\psi/(1-\beta)$. Following Aghion and Howitt (1992), I ignore this issue.

for all q and t. Consequently, the demand for the machine of quality q at time t is given by

$$x(t \mid q) = q(t) L_E(t),$$

and monopoly profits are

$$\pi(t \mid q) = \beta q(t) L_E(t).$$

Aggregate output can now be written as

$$Y(t \mid q) = \frac{1}{1-\beta} q(t) L_E(t),$$

which again conditions on the quality of the machine available at the time, q. This also implies that the equilibrium wage, determined from the final good sector, is given by

$$w(t \mid q) = \frac{\beta}{1-\beta} q(t).$$

When there is no need to emphasize time dependence, I write this wage rate as a function of machine quality, that is, as $w(q)$.

Let us now focus on a "steady-state equilibrium" in which the flow rate of innovation is constant and equal to z^*. Steady state here is in quotation marks since, even though the flow rate of innovation is constant, consumption and output growth will not be constant because of the stochastic nature of innovation (and this is the reason why I do not use the term "BGP" in this context). This implies that a constant number (and thus a constant fraction) of workers, L_R^*, must be working in research. Since the interest rate is equal to $r^* = \rho$, this implies that the steady-state value of a monopolist with a machine of quality q is given by

$$V(q) = \frac{\beta q (L - L_R^*)}{\rho + z^*},$$

where I used the fact that in steady state total employment in the final good sector is equal to $L_E^* = L - L_R^*$. To simplify the notation, I also wrote V as a function of q only rather than a function of both q and time. Free entry requires that when the current machine quality is q, the wage paid to one more R&D worker, $w(q)$, must be equal to the flow benefits, $\eta V(\lambda q)$, thus

$$w(q) = \eta V(\lambda q).$$

Flow benefits from R&D are equal to $\eta V(\lambda q)$, since, when current machine quality is q, one more worker in R&D leads to the discovery of a new machine of quality λq at the flow rate η. In addition, given the R&D technology, we must have $z^* = \eta L_R^*$. Combining these four equations:

$$\frac{\lambda(1-\beta)\eta(L - L_R^*)}{\rho + \eta L_R^*} = 1,$$

which uniquely determines the steady-state number of workers in research as

$$L_R^* = \frac{\lambda(1-\beta)\eta L - \rho}{\eta + \lambda(1-\beta)\eta}, \tag{14.26}$$

as long as this expression is positive (otherwise the free-entry condition would be slack and there would be zero growth).

Contrary to the model in the previous section, however, this does not imply that output grows at a constant rate. Since there is only one sector undergoing technological change and this sector experiences growth only at finite intervals, the growth rate of the economy will have an *uneven* nature; in particular, it can be verified that the economy will have constant output for an interval of time (of average length $1/\eta L_R^*$; see Exercise 14.16) and then will have a burst of growth when a new machine is invented. This pattern of uneven growth is a consequence of having only one sector rather than the continuum of sectors in the model of the previous section. Whether it provides a better approximation to reality is open to debate. While modern capitalist economies do not grow at constant rates, they also do not have as jagged a growth performance as that implied by this model.

The results of this analysis are summarized in the next proposition.

Proposition 14.4. *Consider the one-sector Schumpeterian growth model presented in this section and suppose that*

$$0 < \lambda (1-\beta) \eta L - \rho < \frac{1+\lambda(1-\beta)\rho}{\ln \lambda}. \tag{14.27}$$

Then, there exists a unique steady-state equilibrium in which L_R^* workers work in the research sector, where L_R^* is given in eq. (14.26). The economy has an average growth rate of $g^* = \eta L_R^* \ln \lambda$. Equilibrium growth is "uneven," in the sense that the economy has constant output for a while and then grows by a discrete amount when an innovation takes place.

Proof. Much of the proof is provided by the preceding analysis. Exercise 14.17 asks you to verify that the average growth is given by $g^* = \eta L_R^* \ln \lambda$ and that (14.27) is necessary for the above-described equilibrium to exist and to satisfy the transversality condition. □

Therefore, this analysis shows that the basic insights of the one-sector Schumpeterian model, as originally developed by Aghion and Howitt (1992), are very similar to the baseline model of Schumpeterian growth presented in the previous section. The main difference is that growth has an uneven flavor in the one-sector model, because it is driven by infrequent bursts of innovation, preceded and followed by periods of no growth.

14.2.2. Uneven Growth and Endogenous Cycles*. The analysis in the previous subsection showed how the basic one-sector Schumpeterian growth leads to an uneven pattern of economic growth. This is driven by the discrete nature of innovations in continuous time. There is another source of uneven growth in this basic model, which is more closely related to the process of creative destruction. The nature of Schumpeterian growth implies that future growth reduces the value of current innovations, because it causes more rapid replacement of existing technologies. This effect did not play a role in our analysis so far, because in the model with a continuum of sectors, growth takes a smooth form and as Proposition 14.2 showed, there is a unique equilibrium path with no transitional dynamics. The one-sector growth model analyzed in this section allows these effects to manifest themselves. To show the potential for these creative destruction effects, I now construct a variant of the model

which exhibits endogenous growth cycles. Throughout, I focus on an equilibrium path with such a cycle.

The only difference is that the R&D technology now implies that L_R workers in research leads to innovation at the rate

$$\eta\left(L_R\right) L_R,$$

where $\eta\left(\cdot\right)$ is a strictly decreasing function, representing an externality in the research process. When more firms try to discover the next generation of technology, there will be more crowding-out in the research process, making it less likely for each of them to innovate. Each firm ignores its effect on the aggregate rate of innovation, thus takes $\eta\left(L_R\right)$ as given (this assumption is not important as shown by Exercise 14.22). Consequently, when the current machine quality is q, the free-entry condition takes the form

$$\eta\left(L_R\left(q\right)\right) V\left(\lambda q\right) = w\left(q\right),$$

where $L_R\left(q\right)$ is the number of workers employed in research when the current machine quality is q.

Let us now look for an equilibrium with the following cyclical property: the rate of innovation differs when the innovation in question is an odd-numbered innovation versus an even-numbered innovation (say with the number of innovations counted starting from some arbitrary date $t = 0$). This type of equilibrium is possible when all agents in the economy expect there to be such an equilibrium (i.e., it is a "self-fulfilling" equilibrium). Denote the number of workers in R&D for odd and even-numbered innovations by L_R^1 and L_R^2. Then, following the analysis in the previous subsection, in any equilibrium with a cyclical pattern the values of odd and even-numbered innovations (with a machine of quality q) can be written as (see Exercise 14.20):

$$V^2\left(\lambda q\right) = \frac{\beta q\left(L - L_R^2\right)}{\rho + \eta\left(L_R^2\right) L_R^2} \text{ and } V^1\left(\lambda q\right) = \frac{\beta q\left(L - L_R^1\right)}{\rho + \eta\left(L_R^1\right) L_R^1}, \tag{14.28}$$

and the free-entry conditions take the form

$$\eta\left(L_R^1\right) V^2\left(\lambda q\right) = w\left(q\right) \text{ and } \eta\left(L_R^2\right) V^1\left(\lambda q\right) = w\left(q\right),$$

where $w\left(q\right)$ is the equilibrium wage with technology of quality q. The reason why $\eta\left(L_R^1\right)$ multiplies the value for an even-numbered innovation is because L_R^1 researchers are employed for innovation today, when the current technology is odd-numbered, but the innovation that this research will produce will be even-numbered and thus will have value $V^2\left(\lambda q\right)$. Therefore, the following two equilibrium conditions apply:

$$\eta\left(L_R^1\right) \frac{\lambda\left(1-\beta\right) q\left(L - L_R^2\right)}{\rho + \eta\left(L_R^2\right) L_R^2} = 1 \text{ and } \eta\left(L_R^2\right) \frac{\lambda\left(1-\beta\right) q\left(L - L_R^1\right)}{\rho + \eta\left(L_R^1\right) L_R^1} = 1. \tag{14.29}$$

It can then be verified that these two equations can have solutions L_R^1 and $L_R^2 \neq L_R^1$, which would correspond to the possibility of a two-period endogenous cycle (see Exercise 14.21).

14.2.3. Labor Market Implications of Creative Destruction. Another important implication of creative destruction is related to the fact that growth destroys existing productive units. So far this only led to the destruction of the monopoly rents of incumbent producers, without any loss of employment. In more realistic economies, creative destruction may dislocate previously employed workers and these workers may experience some unemployment before finding a new job. How creative destruction may lead to unemployment is discussed in Exercise 14.19.

A final implication of creative destruction that is worth noting relates to the destruction of firm-specific skills. It may be efficient for workers to accumulate human capital that is specific to their employers. Creative destruction implies that productive units may have shorter horizons in an economy with rapid economic growth. An important consequence of this might be that in rapidly-growing economies, workers (and sometimes firms) may be less willing to make a range of specific human capital and other investments.

14.3. Innovation by Incumbents and Entrants and Sources of Productivity Growth

A key aspect of the growth process is the interplay between innovations and productivity improvements by existing firms on the one hand and entry by more productive, new firms on the other. The evidence from industry studies, which will be discussed in greater detail in Section 18.1, suggests that a large part of productivity growth at the industry level (and thus in the aggregate) comes from productivity improvements by continuing plants, though entry by new plants also makes a nontrivial contribution to industry productivity growth. The Schumpeterian models presented in this section have emphasized entry by new firms as the engine of growth. Interpreted literally, these models predict that all growth should be driven by entry, which is at odds with the facts. The expanding variety models presented in the previous chapter also do not provide a framework for the analysis of the interplay between existing firms and new entrants.[4] In this and the next section, I discuss models that feature productivity growth by continuing plants (firms). The model in this section will feature productivity growth both by incumbents and entrants. The model in the next section will be richer in many respects, but will not allow entry. The two models together provide a first glimpse of the types of models that might be useful for studying the industrial organization of innovation and productivity growth.

14.3.1. Model. The economy is again in continuous time and admits a representative household with the standard CRRA preferences, as in (13.1) in the previous chapter. Population is constant at L and labor is supplied inelastically. The resource constraint at time t takes the usual form

$$C(t) + X(t) + Z(t) \leq Y(t), \tag{14.30}$$

[4]In the expanding variety models, the identity of the firms that are undertaking the innovation does not matter, so one could assume that it is the existing producers that are inventing new varieties, though this will be essentially determining the distribution of productivity improvements across firms by assumption.

where $C(t)$ is consumption, $X(t)$ is aggregate spending on machines, and $Z(t)$ is total expenditure on R&D at time t.

The production function of the unique final good is given by

$$Y(t) = \frac{1}{1-\beta}\left[\int_0^1 q(\nu,t)^{\beta} x(\nu,t \mid q)^{1-\beta} d\nu\right] L^{\beta}, \tag{14.31}$$

where $x(\nu,t \mid q)$ is the quantity of the machine of type ν of quality q used in the production process and the measure of machines is again normalized to 1. This aggregate production function is very similar to (14.3) used above, except that the quality of machines comes in with an exponent β. This has no effects on any of the results concerning growth, but will imply that firms with different productivity levels will have different levels of sales (see Exercise 14.27). It will therefore enable us to make predictions about the size distribution of firms as well.

The engine of economic growth is again quality improvements, but these will be driven by two types of innovations:

(1) Innovation by incumbents

(2) Creative destruction by entrants.

Let $q(\nu,t)$ be the quality of machine line ν at time t. Consider the following "quality ladder" for each machine type:

$$q(\nu,t) = \lambda^{n(\nu,t)} q(\nu,s) \text{ for all } \nu \text{ and } t,$$

where $\lambda > 1$ and $n(\nu,t)$ denotes the number of *incremental* innovations on this machine line between time $s \leq t$ and time t, where time s is the date at which this particular type of technology was first invented and $q(\nu,s)$ refers to its quality at that point. The incumbent has a fully enforced patent on the machines that it has developed (though this patent does not prevent entrants leapfrogging the incumbent's machine quality). I assume that at time $t = 0$, each machine line starts with some quality $q(\nu,0) > 0$ owned by an incumbent with fully enforce patent on this initial machine quality. Incremental innovations can only be performed by the incumbent producer. So we can think of those as "tinkering" innovations that improve the quality of the machine. The assumption that incumbents have access to a technology to create incremental innovations is consistent with case study evidence (e.g., Freemen, 1982, or Sherer, 1984).

More specifically, if the current incumbent spends an amount $z(\nu,t)\,q(\nu,t)$ of the final good for this type of innovation activity on a machine of current quality $q(\nu,t)$, it has a flow rate of innovation equal to $\phi z(\nu,t)$ for $\phi > 0$ (more formally, this implies that for any interval $\Delta t > 0$, the probability of one incremental innovation is $\phi z(\nu,t)\,\Delta t$ and the probability of more than one incremental innovation is $o(\Delta t)$ (with $o(\Delta t)/\Delta t \to 0$ as $\Delta t \to 0$). Recall that such an innovation results in an improvement in quality and the resulting new machine will be of quality $\lambda q(\nu,t)$.

Alternatively, a new firm (entrant) can undertake R&D to innovate over the existing machines in machine line ν at time t. If the current quality of machine is $q(\nu,t)$, then by

spending one unit of the final good this new firm innovate at the flow rate $\eta\left(\hat{z}\left(\nu,t\right)\right)/q\left(\nu,t\right)$, where $\eta\left(\cdot\right)$ is a strictly decreasing, continuous and differentiable function and $\hat{z}\left(\nu,t\right)$ is the total amount of R&D by new entrants towards machine line ν at time t. Incumbents can also be allowed to have access to the same technology for radical innovation as the entrants. However, the Arrow replacement effect then immediately implies that incumbents would never use this technology (since entrants will be making zero profits from this technology, the profits of incumbents would be negative; see Exercise 14.23). Incumbents will still find it profitable to use the technology for incremental innovations, which is not available to entrants.

The presence of the strictly decreasing function η, which was also used in Section 14.2, captures the fact that when many firms are undertaking R&D to replace the same machine line, they are likely to try similar ideas, thus there will be some amount of "external" diminishing returns (new entrants will be "fishing out of the same pond"). Since each entrant attempting R&D on this line is potentially small, they will all take $\eta\left(\hat{z}\left(\nu,t\right)\right)$ as given. Throughout I assume that $z\eta\left(z\right)$ is strictly increasing in z so that greater aggregate R&D towards a particular machine line increases the overall probability of discovering a superior machine. I also suppose that $\eta\left(z\right)$ satisfies the following Inada-type assumptions:

$$\lim_{z\to\infty}\eta\left(z\right)=0 \text{ and } \lim_{z\to 0}\eta\left(z\right)=\infty. \tag{14.32}$$

An innovation by an entrant leads to a new machine of quality $\kappa q\left(\nu,t\right)$, where $\kappa>\lambda$. Therefore, innovation by entrants are more "radical" than those of incumbents. Existing empirical evidence from studies of innovation support the notion that innovations by new entrants are more significant or radical than those of incumbents (though it may take a while for the successful entrants to realize the full productivity gains from these innovations; I am abstracting from this aspect). Whether the entrant was a previous incumbent on this specific machine line or whether it is currently an incumbent in some other machine line do not matter for its technology of innovation.

Once a particular machine of quality $q\left(\nu,t\right)$ has been invented, any quantity of this machine can be produced at the marginal cost ψ. I again normalize $\psi\equiv 1-\beta$. This implies that the total amount of expenditure on the production of intermediate goods at time t is

$$X\left(t\right)=\int_{0}^{1}\psi x\left(\nu,t\right)d\nu, \tag{14.33}$$

where $x\left(\nu,t\right)$ is the quantity of this machine used in final good production. Similarly, the total expenditure on R&D is

$$Z\left(t\right)=\int_{0}^{1}\left[z\left(\nu,t\right)+\hat{z}\left(\nu,t\right)\right]q\left(\nu,t\right)d\nu, \tag{14.34}$$

where $q\left(\nu,t\right)$ refers to the highest quality of the machine of type ν at time t. Notice also that total R&D is the sum of R&D by incumbents and entrants ($z\left(\nu,t\right)$ and $\hat{z}\left(\nu,t\right)$ respectively).

An allocation in this economy consists of time paths of consumption levels, aggregate spending on machines, and aggregate R&D expenditure, $\left[C\left(t\right),X\left(t\right),Z\left(t\right)\right]_{t=0}^{\infty}$, time paths for R&D expenditure by incumbent and entrants, $\left[z\left(\nu,t\right),\hat{z}\left(\nu,t\right)\right]_{\nu\in[0,1],t=0}^{\infty}$, time paths of

prices and quantities of each machine and the net present discounted value of profits from that machine, $[p^x(\nu,t \mid q), x(\nu,t), V(\nu,t \mid q)]_{\nu\in[0,1],t=0}^{\infty}$, and time paths of interest rates and wage rates, $[r(t), w(t)]_{t=0}^{\infty}$. An equilibrium is given by an allocation in which R&D decisions by entrants maximize their net present discounted value, pricing, quantity and R&D decisions by incumbents maximize their net present discounted value, consumers choose the path of consumption optimally, and the labor market clears.

Let us start with the aggregate production function for the final good producers. Profit-maximization by the final good sector implies the demand for machines of highest-quality is given by a slight variant of eq. (14.4) in Section 14.1:

$$x(\nu,t \mid q) = p^x(\nu,t \mid q)^{-1/\beta} q(\nu,t) L \quad \text{for all } \nu \in [0,1] \text{ and all } t, \tag{14.35}$$

where $p^x(\nu,t \mid q)$ refers to the price of machine type ν of quality $q(\nu,t)$ at time t. Since the demand from the final good sector for machines in (14.35) is isoelastic, the unconstrained monopoly price will given by the usual formula, as a constant markup over marginal cost. In this context, I introduce the analogue of condition (14.5) above,

$$\kappa > \left(\frac{1}{1-\beta}\right)^{\frac{1-\beta}{\beta}}, \tag{14.36}$$

which ensures that new entrants can charge the unconstrained monopoly price. By implication, incumbents that make further innovations can also charge the unconstrained monopoly price.

14.3.2. Equilibrium. Since the demand for machines in (14.35) is isoelastic and $\psi \equiv 1-\beta$, the profit-maximizing monopoly price is

$$p^x(\nu,t \mid q) = 1. \tag{14.37}$$

Combining this with (14.35) implies that

$$x(\nu,t \mid q) = qL. \tag{14.38}$$

Consequently, the flow profits of a firm with the monopoly rights on the machine of quality q can be computed as:

$$\pi(\nu,t \mid q) = \beta q L. \tag{14.39}$$

Next, substituting (14.38) into (14.31), total output is given by an expression identical to (14.8) above

$$Y(t) = \frac{1}{1-\beta} Q(t) L, \tag{14.40}$$

with average quality of machines $Q(t)$ given as in (14.9) in Section 14.1. Aggregate spending on machines is

$$X(t) = (1-\beta) Q(t) L. \tag{14.41}$$

Moreover, since the labor market is competitive, the wage rate at any point in time is given by (14.11) as before.

To characterize the full equilibrium, we need to determine R&D effort levels by incumbents and entrants. To do this, let us write the net present value of a monopolist with the highest quality of machine q at time t in machine line ν. This value satisfies the standard Hamilton-Jacobi-Bellman equation

$$\begin{aligned}(14.42)\quad r(t)V(\nu,t\mid q)-\dot{V}(\nu,t\mid q)=\max_{z(\nu,t|q)\geq o}\{&\pi(\nu,t\mid q)-z(\nu,t\mid q)q(\nu,t)\\&+\phi z(\nu,t\mid q)(V(\nu,t\mid\lambda q)-V(\nu,t\mid q))-\hat{z}(\nu,t\mid q)\eta(\hat{z}(\nu,t\mid q))V(\nu,t\mid q)\},\end{aligned}$$

where $\hat{z}(\nu,t\mid q)\eta(\hat{z}(\nu,t\mid q))$ is the rate at which radical innovations by entrants occur in sector ν at time t and $z(\nu,q\mid q)$ is the rate at which the incumbent improves its technology. The first term in (14.42), $\pi(\nu,t\mid q)$, is the flow of profits given by (14.39), while the second term is the expenditure of the incumbent for improving the quality of its product. The second line includes changes in the value of the incumbent due to innovation either by itself (at the rate $\phi z(\nu,t\mid q)$, the quality of its product increases from q to λq) or by an entrant (at the rate $\hat{z}(\nu,t\mid q)\eta(\hat{z}(\nu,t\mid q))$, the incumbent is replaced and receives zero value from then on).[5] The value function is written with a maximum on the right hand side, since $z(\nu,t\mid q)$ is a choice variable for the incumbent.

Free entry by entrants implies a free-entry condition similar to (14.13) in Section 14.1,

$$\begin{aligned}(14.43)\qquad \eta(\hat{z}(\nu,t\mid q))V(\nu,t\mid\kappa q)&\leq q(\nu,t),\text{ and}\\ \eta(\hat{z}(\nu,t\mid q))V(\nu,t\mid\kappa q)&=q(\nu,t)\text{ if }\hat{z}(\nu,t\mid q)>0,\end{aligned}$$

which takes into account that by spending an amount $q(\nu,t)$, the entrant generates a flow rate of innovation of $\eta(\hat{z})$, and if this innovation is successful, the entrant will end up with a product of quality κq, thus earning the value $\eta(\hat{z}(\nu,t\mid q))V(\nu,t\mid\kappa q)$.

In addition, the incumbent's choice of R&D effort implies a similar complementary slackness condition

$$\begin{aligned}(14.44)\qquad \phi(V(\nu,t\mid\lambda q)-V(\nu,t\mid q))&\leq q(\nu,t)\text{ and}\\ \phi(V(\nu,t\mid\lambda q)-V(\nu,t\mid q))&=q(\nu,t)\text{ if }z(\nu,t\mid q)>0.\end{aligned}$$

Finally, consumer maximization implies the familiar Euler equation and the transversality condition given by (14.14) and (14.15) as before.

In light of this analysis, an equilibrium can be more succinctly defined as time path of $[C(t),X(t),Z(t)]_{t=0}^{\infty}$ that satisfy (14.30), (14.34), (14.41) and (14.15); time paths for R&D expenditure by incumbent and entrants $[z(\nu,t),\hat{z}(\nu,t)]_{\nu\in[0,1],t=0}^{\infty}$ that satisfy (14.43) and (14.44); time paths of prices and quantities of each machine and the net present discounted value of profits from that machine, $[p^x(\nu,t\mid q),x(\nu,t),V(\nu,t\mid q)]_{\nu\in[0,1],t=0}^{\infty}$ given by (14.37), (14.7) and (14.42); and time paths of wage and interest rates, $[w(t),r(t)]_{t=0}^{\infty}$ that satisfy (14.11) and (14.14).

[5]The fact that the incumbent receives a zero value from then on follows from the assumption that a previous incumbent has no advantage relative to other entrants in competing for another round of innovations.

As usual, I define a BGP as an equilibrium path in which innovation, output and consumption growth a constant rate. Notice that in BGP, aggregates grow at the constant rate, but there will be firm deaths and births, and the firm size distribution may also change.

The requirement that consumption grows at a constant rate in the BGP implies that $r(t) = r^*$ (from (14.14)). Moreover, in BGP, we must also have $z(\nu, t \mid q) = z(q)$ and $\hat{z}(\nu, t \mid q) = \hat{z}(q)$. These together imply that in BGP $\dot{V}(\nu, t \mid q) = 0$ and $V(\nu, t \mid q) = V(q)$. Finally, since profits and costs are both proportional to quality q, we can also see that $\hat{z}(q) = \hat{z}$ and $V(q) = vq$ (see the proof of Proposition 14.6, which in fact shows that , $\hat{z}(\nu, t \mid q) = \hat{z}(t)$ and $V(\nu, t \mid q) = v(t)\, q$ in any equilibrium, even outside the BGP). These results enable a straightforward characterization of the BGP and the dynamic equilibrium. While $\hat{z}(q) = \hat{z}$ for all q, it is not necessarily true that $z(q) = z$ for all q. In fact, as we will see the equilibrium only pins down the average R&D intensity of incumbents.

Let us first look for an "interior" BGP equilibrium (we will verify below that such an interior BGP exists and is the unique equilibrium). This implies that incumbents undertake research, thus

$$\phi\left(V(\nu, t \mid \lambda q) - V(\nu, t \mid q)\right) = q(\nu, t). \tag{14.45}$$

Given the linearity of V in q, this implies the following convenient equation for the value of a firm with a machine of quality q:

$$V(q) = \frac{q}{\phi(\lambda - 1)}. \tag{14.46}$$

Moreover, from the free-entry condition (again holding as equality since the BGP is interior),

$$V(q) = \frac{\beta L q}{r^* + \hat{z}\eta(\hat{z})}. \tag{14.47}$$

Moreover, from the free-entry condition (again holding as equality from the fact that the equilibrium is interior):

$$\eta(\hat{z})\, V(\kappa q) = q \text{ or } V(q) = \frac{q}{\kappa\eta(\hat{z})}.$$

Combining this expression with (14.45) and (14.46),

$$\frac{\phi(\lambda - 1)}{\kappa\eta(\hat{z})} = 1.$$

This implies that the BGP R&D level by entrants $\hat{z}^*$ is implicitly defined by

$$\hat{z}(q) = \hat{z}^* \equiv \eta^{-1}\left(\frac{\phi(\lambda - 1)}{\kappa}\right) \text{ for all } q > 0. \tag{14.48}$$

Next, combining this with (14.47), the BGP interest rate is

$$r^* = \phi(\lambda - 1)\beta L - \hat{z}^*\eta(\hat{z}^*). \tag{14.49}$$

From (14.14), the growth rate of consumption and output is therefore given by

$$g^* = \frac{1}{\theta}\left(\phi(\lambda - 1)\beta L - \hat{z}^*\eta(\hat{z}^*) - \rho\right). \tag{14.50}$$

Equation (14.50) already has some interesting implications. In particular, it determines the relationship between the rate of innovation by entrants $\hat{z}^*$ and the BGP growth rate g^*. In standard Schumpeterian models, this relationship is positive. In contrast, here:

PROPOSITION 14.5. *There is a negative relationship between $\hat{z}^*$ and g^*.*

Equation (14.50), together with (14.48), determines the BGP growth rate of the economy, but does not specify how much of productivity growth is driven by creative destruction (innovation by entrants) and how much of it by productivity improvements by incumbents. To determine this, is useful to repeat the analysis as in Section 14.1. Recall, at this point, that $z(\nu, t \mid q)$ is not a function of ν, but could still depend on q. Consequently, the law of motion of average quality, $Q(t)$, is

$$\begin{aligned} Q(t+\Delta t) &= (\lambda \phi z(t) \Delta t) Q(t) + (\kappa \hat{z}(t) \eta(\hat{z}(t)) \Delta t) Q(t) \\ &\quad + (1 - \phi z(t) \Delta t - \hat{z}(t) \eta(\hat{z}(t)) \Delta t) Q(t) + o(\Delta t), \end{aligned} \tag{14.51}$$

where

$$z(t) = \frac{\int_0^1 z(\nu, t \mid q) q(\nu, t) d\nu}{Q(t)} \tag{14.52}$$

is the average R&D effort of incumbents that time t. Now subtracting $Q(t)$ from both sides of (14.51), dividing by Δt and taking the limit as $\Delta t \to 0$,

$$\frac{\dot{Q}(t)}{Q(t)} = (\lambda - 1) \phi z(t) + (\kappa - 1) \hat{z}(t) \eta(\hat{z}(t)). \tag{14.53}$$

Therefore, an alternative expression for the growth rate of the economy, which decomposes growth into the component coming from incumbent firms (the first term) and that coming from new entrants (the second term) is given as

$$g^* = (\lambda - 1) \phi z^* + (\kappa - 1) \hat{z}^* \eta(\hat{z}^*), \tag{14.54}$$

where z^* is the average BGP R&D effort of incumbents. The fact that this average R&D effort is constant in BGP follows from the fact from (14.53) together with the fact that in BGP the growth rate of average quality is g^* and the R&D effort by entrants on each machine line is $\hat{z}^*$. While (14.50) pins down the BGP growth rate of output and consumption, eq. (14.54) determines how much of it is driven by innovation by incumbents and how much of it by innovation by entrants. Moreover, we can also verify that this economy does not have any transitional dynamics (see the proof of Proposition 14.6). Therefore, if an equilibrium with growth exists, it will involve growth at the rate g^*. To ensure that such an equilibrium exists, we need to verify that R&D is profitable both for entrants and incumbents. The condition that the BGP interest rate, r^*, given by (14.49), should be greater than the discount rate ρ is sufficient for there to be positive aggregate growth. In addition, this interest rate should not be so high that the transversality condition of the consumers is violated. Finally, we need to ensure that there is also innovation by incumbents. The following condition ensures all three of these requirements (see the proof of Proposition 14.6):

$$\phi(\lambda - 1) \beta L - (\theta(\kappa - 1) + 1) \hat{z}^* \eta(\hat{z}^*) > \rho > (1 - \theta)(\phi(\lambda - 1) \beta L - \hat{z}^* \eta(\hat{z}^*)), \tag{14.55}$$

with $\hat{z}^*$ given by (14.48).

In addition, our main interest is with how much of productivity growth and innovation are driven by incumbents and how much by new entrants. This can now be easily obtained from (14.50) and (14.54) as

$$(\lambda - 1)\,\phi z^* = g^* - (\kappa - 1)\,\hat{z}^*\eta\,(\hat{z}^*)\,, \tag{14.56}$$

where g^* is given by (14.50) and $\hat{z}^*$ by (14.48).

Another set of interesting implications of this model concern firm-size dynamics. The size of a firm can be measured by its sales, which is equal to

$$x\,(\nu, t \mid q) = qL \text{ for all } \nu \text{ and } t.$$

To determine the law of motion of firm sales and thus the firm-size dynamics, we need to know how incumbent R&D effort varies with q. To start with, let us suppose that in BGP this R&D effort is independent of q, so that $z\,(\nu, t \mid q) = z^*$ for all q. From the analysis so far, it is clear that such an equilibrium will exist (for a justification of the focus on this equilibrium, see Exercise 14.25; this issue will be discussed further below). Then, the quality of each incumbent firm will increase at the flow rate ϕz^*, with z^* given by (14.56). At the same time, each incumbent is also replaced at the flow rate $\hat{z}^*\eta\,(\hat{z}^*)$. Therefore, for Δt sufficiently small, the stochastic process for the size of a particular firm is given by

$$x\,(\nu, t + \Delta t \mid q) = \begin{cases} \lambda x\,(\nu, t \mid q) & \text{with probability } \phi z^* \Delta t + o\,(\Delta t) \\ 0 & \text{with probability } \hat{z}^*\eta\,(\hat{z}^*)\,\Delta t + o\,(\Delta t) \\ x\,(\nu, t \mid q) & \text{with probability } (1 - \phi z^* \Delta t - \hat{z}^*\eta\,(\hat{z}^*)\,\Delta t) + o\,(\Delta t)\,, \end{cases} \tag{14.57}$$

for all ν and t, where $\hat{z}^*$ is by (14.48) and z^* by (14.56). Therefore, firms have random growth, and surviving firms expand on average. However, firms also face a probability of replacement (extinction). In particular, let us denote the probability that a particular incumbent firm that started production in machine line ν at time s will be bankrupt by time $t \geq s$ by $P\,(t \mid s, \nu)$. Then, we have $\lim_{t\to\infty} P\,(t \mid s, \nu) = 1$, so that each firm will necessarily die eventually.

PROPOSITION 14.6. *Consider the above-described economy starting with an initial condition $Q\,(0) > 0$. Suppose that (14.32) and (14.55) are satisfied and focus on equilibrium in which all incumbents exert the same level of R&D effort. Then, there exists a unique equilibrium. In this equilibrium growth is always balanced, and technology, $Q\,(t)$, aggregate output, $Y\,(t)$, and aggregate consumption, $C\,(t)$, grow at the rate g^* as in (14.50) with $\hat{z}^*$ given by (14.48). Equilibrium growth is driven both by innovation by incumbents and by creative destruction by entrants. Any given firm expands on average as long as it survives, but is eventually replaced by a new entrant with probability one.*

PROOF. First, note that in an interior BGP where $\phi\,(V\,(\nu, t \mid \lambda q) - V\,(\nu, t \mid q)) = q$, V must be linear in q, thus $V\,(\nu, t \mid q) = vq$ as used in the text. Given this observation, the characterization of the BGP follows from the argument preceding the proposition. In particular, $\hat{z}^*$ is uniquely determined by (14.48) and (14.50) gives the unique BGP growth rate. To ensure that this is indeed an interior BGP, we need to check four conditions:

(1) *Positive $\hat{z}^*$:* this follows from (14.32), e.g., see observation [A] below.

(2) *Positive growth:* this requires $g^* = (\phi(\lambda-1)\beta L - \hat{z}^*\eta(\hat{z}^*) - \rho)/\theta > 0$. Since $(\theta(\kappa-1)+1) > 1$, the first inequality in (14.55), is sufficient for this.

(3) *Positive z^*:* combining (14.50) together with (14.54),

$$\begin{aligned} z^* &= \frac{g^* - (\kappa-1)\hat{z}^*\eta(\hat{z}^*)}{(\lambda-1)\phi} \\ &= \frac{(\phi(\lambda-1)\beta L - \hat{z}^*\eta(\hat{z}^*) - \rho)/\theta - (\kappa-1)\hat{z}^*\eta(\hat{z}^*)}{(\lambda-1)\phi} > 0, \end{aligned}$$

in view of the first inequality (14.55).

(4) *The transversality condition:* condition (14.15) should hold so that the maximization problem of the representative household is well-defined. The condition $r^* > g^*$ is necessary and sufficient to ensure (14.15). The second inequality in (14.55) ensures that this inequality holds.

Therefore, the BGP is interior and is uniquely defined.

I next prove that the BGP also gives the unique dynamic equilibrium path. Let us start with two observations.

[A] Because of the Inada conditions, (14.32), the free-entry condition (14.13) must hold as equality for all ν, t and q, so that

$$\eta(\hat{z}(\nu,t \mid q))\, V(\nu,t \mid \kappa q) = q \text{ for all } \nu,\ q \text{ and } t. \tag{14.58}$$

Since this equation holds for all t and the right-hand side is differentiable in q and t, so must be the left-hand side. Differentiating with respect to t,

$$\frac{\partial \hat{z}(\nu,t \mid q)/\partial t}{\hat{z}(\nu,t \mid q)} = \frac{1}{\varepsilon_\eta(\hat{z}(\nu,t \mid q))} \frac{\dot{V}(\nu,t \mid \kappa q)}{V(\nu,t \mid \kappa q)} \text{ for all } q \text{ and } t, \tag{14.59}$$

where

$$\varepsilon_\eta(\hat{z}) \equiv -\frac{\eta'(\hat{z})\,\hat{z}}{\eta(\hat{z})} > 0$$

is the elasticity of the η function.

[B] The value of a firm with a machine of quality q at time t can be written as

$$V(\nu,t \mid q) = \int_t^\infty \exp\left[-\int_t^s \left(r(s') + \hat{z}(\nu,s' \mid q)\eta\left(\hat{z}(\nu,s' \mid q)\right)\right) ds'\right] \pi(\nu,s \mid q)\, ds. \tag{14.60}$$

For any finite q, this value is finite since, from observation [A], $(r(t) + \hat{z}(\nu,t \mid q)\eta(\hat{z}(\nu,t \mid q))) > 0$ for all ν, q and t, and because $\pi(\nu,s \mid q) = \beta q L \in (0,\infty)$. In addition, in view of (14.32), $V(\nu,t \mid q)$ is uniformly bounded away from 0 for any path of interest rates that does not limit to infinity. Moreover, $r(t) \to \infty$ is also impossible, since, from (14.14), it would imply $\dot{C}(t)/C(t) \to \infty$, violating the transversality condition, (14.15).

Now consider the following two cases.

Case 1: Suppose that (14.45), that is, $\phi(V(\nu,t \mid \lambda q) - V(\nu,t \mid q)) = q$, holds for all ν, q and t. Then, as in BGP, $V(\nu,t \mid q)$ is linear in q and can be written as $V(\nu,t \mid q) = v(t)\,q$

for some function $v(t)$. This implies that (14.45) can be written as $\phi(\lambda-1)v(t)=1$ for all t. Differentiating this equation with respect to time, $\dot{v}(t)=0$ and $v(t)=v$ for all t. Moreover, from observation [A] above, the free-entry condition (14.43) must hold as equality, so $\eta\left(\hat{z}(\nu,t\mid\kappa^{-1}q)\right)v(t)=1$ for all t, which is only possible if $\hat{z}(\nu,t\mid q)=\hat{z}(t)$ for all q and t. Combining these conditions with (14.42) yields

$$r(t)v=\beta L-\hat{z}\eta(\hat{z})v$$

for all t, which implies that $r(t)$ must be constant as well. Therefore, all variables must immediately take their BGP values, $r(t)=r^*$ and $\hat{z}(t)=\hat{z}^*$ for all t. Moreover, we are focusing on equilibria where $z(\nu,t\mid q)$ is independent of q, and thus we must have $z(\nu,t\mid q)=z^*$ for all t. This establishes that the economy must be in the BGP equilibrium starting at $t=0$.

Case 2: Suppose that (14.45) does not hold for some $\nu\in\mathcal{N}\subset[0,1]$, q and t. Since either $\phi(V(\nu,t\mid\lambda q)-V(\nu,t\mid q))=q$ or $\phi(V(\nu,t\mid\lambda q)-V(\nu,t\mid q))<q$ and $z(\nu,t\mid q)=0$, and because from observation [A] above, $\eta\left(\hat{z}(\nu,t\mid\kappa^{-1}q)\right)V(\nu,t\mid q)=\kappa^{-1}q$, the value function (14.42) in this case can be written as

$$\frac{\dot{V}(\nu,t\mid q)}{V(\nu,t\mid q)}=r(t)+\hat{z}(\nu,t\mid q)\eta(\hat{z}(\nu,t\mid q))-\kappa\beta L\eta\left(\hat{z}(\nu,t\mid\kappa^{-1}q)\right).$$

Combining this with (14.59),

$$\frac{\partial\hat{z}(\nu,t\mid\kappa^{-1}q)/\partial t}{\hat{z}(\nu,t\mid\kappa^{-1}q)}=\frac{1}{\varepsilon_\eta(\hat{z}(\nu,t\mid\kappa^{-1}q))}\left[r(t)+\hat{z}(\nu,t\mid q)\eta(\hat{z}(\nu,t\mid q))-\kappa\beta L\eta\left(\hat{z}(\nu,t\mid\kappa^{-1}q)\right)\right].$$

Similarly,

$$\frac{\partial\hat{z}(\nu,t\mid q)/\partial t}{\hat{z}(\nu,t\mid q)}=\frac{1}{\varepsilon_\eta(\hat{z}(\nu,t\mid q))}\left[r(t)+\hat{z}(\nu,t\mid\kappa q)\eta(\hat{z}(\nu,t\mid\kappa q))-\beta L\eta(\hat{z}(\nu,t\mid q))\right],$$

and so on. It is then straightforward to see that the system of differential equations for $\hat{z}(\nu,t\mid q)$ for all q is unstable, in the sense that if $\partial\hat{z}(\nu,t\mid q)/\partial t>0$, then it will grow continuously and if $\partial\hat{z}(\nu,t\mid q)/\partial t<0$, then it will shrink continuously. This implies that we must have $\partial\hat{z}(\nu,t\mid q)/\partial t=0$ for all $\nu\in[0,1]$. Suppose, to obtain a contradiction, that there exists (a positive measure) set $\mathcal{N}\subset[0,1]$ such that for all $\nu\in\mathcal{N}$, $\partial\hat{z}(\nu,t\mid q)/\partial t>0$. Then, $\hat{z}(\nu,t\mid q)\to\infty$ for all $\nu\in\mathcal{N}$, and thus $\eta(\hat{z}(\nu,t\mid q))\to 0$. But to ensure the free-entry condition, (14.43), which from observation [A] always holds as equality, we need $\eta(\hat{z}(\nu,t\mid q))V(\nu,t\mid\kappa q)=q$ and thus $V(\nu,t\mid\kappa q)\to\infty$. But this is impossible in view of observation [B] above. Next, suppose that there exists $\mathcal{N}\subset[0,1]$ such that for all $\nu\in\mathcal{N}$, $\partial\hat{z}(\nu,t\mid q)/\partial t<0$. Then, $\hat{z}(\nu,t\mid q)\to 0$ and thus $\eta(\hat{z}(\nu,t\mid q))\to\infty$, which from the free-entry condition implies $V(\nu,t\mid\kappa q)\to 0$ and contradicts observation [B] that V is uniformly bounded away from zero. Therefore, $\partial\hat{z}(\nu,t\mid q)/\partial t=0$ for all ν, q and t, which implies

$$r(t)+\hat{z}(\nu,t\mid\kappa q)\eta(\hat{z}(\nu,t\mid\kappa q))-\beta L\eta(\hat{z}(\nu,t\mid q))=0 \text{ for all } \nu, q \text{ and } t.$$

This is only possible if $r(t)$ is constant and thus equal to r^* and $\hat{z}(\nu,t\mid q)=\hat{z}^*$, establishing the result that also in this case there are no transitional dynamics. Moreover, since

$\hat{z}(\nu, t \mid q) = \hat{z}^*$ for all ν, q and t, (14.60) implies that $V(\nu, t \mid q)$ is linear in q in this case as well.

Finally, the result that surviving firms expand on average and that all firms die with probability 1 follows from eq. (14.57). □

Proposition 14.6 focuses on equilibria in which all incumbents exert the same R&D effort. Exercise 14.26 shows that the same conclusions hold when we do not focus on such equilibria.

14.3.3. Some Numbers. I will now try to flesh out the implications of this model on the decomposition of productivity growth between incumbents and entrants. Although some of the parameters of the current model are difficult to pin down with our current knowledge of the technology of R&D, some simple back-of-the-envelope calculations are still informative. Let us choose the following standard numbers:

$$g^* = 0.02,\ \rho = 0.01,\ r^* = 0.05, \text{ and } \theta = 2,$$

where the last number, the intertemporal elasticity of substitution, is pinned down by the choice of the other three numbers. The first three numbers refer to annual rates (implicitly defining $\Delta t = 1$ as one year). There is much greater uncertainty concerning the remaining parameters. Let us normalize $\phi = L = 1$. For the rest, I will report a number of different possibilities. As a benchmark, As a benchmark, I take $\beta = 2/3$, which implies that two thirds of national income accrues to labor and one third to profits. The requirement in (14.36) then implies that $\kappa > 1.7$. I also take the benchmark value of $\kappa = 3$, so that entry by new firms is sufficiently "radical" as suggested by some of the qualitative accounts of the innovation process (e.g., Freeman, 1982, Scherer, 1984). Innovation by incumbents is taken to be correspondingly smaller $\lambda = 1.2$. This implies that productivity gains from a radical innovation is about ten times that of a standard "incremental" innovation by incumbents (i.e., $(\kappa - 1)/(\lambda - 1) = 10$). I will then show how the results change when the magnitudes of radical and incremental innovations are varied. For the function $\eta(z)$, I adopt the following frequently-used form:

$$\eta(z) = Bz^{-\alpha},$$

and choose the benchmark value of $\alpha = 0.5$. The remaining two parameters, ϕ and B will be chosen to ensure $g^* = 0.02$. I start with the benchmark value of $\phi = 0.4$, but this value will need to be modified in some of the variations in order to satisfy condition (14.55) above or to ensure $g^* = 0.02$. With these numbers, (14.48) implies

$$\hat{z}^* \eta(\hat{z}^*) = 0.0033,$$

and

$$(\kappa - 1)\hat{z}^* \eta(\hat{z}^*) = 0.0067.$$

The value for $\hat{z}^* \eta(\hat{z}^*)$ also implies that there is entry of a new firm (creative destruction) in each machine line on average once every 7.5 years (recall that $r^* = 0.05$ as the annual interest

rate so that $r^*/\hat{z}^*\eta(\hat{z}^*) \approx 7.46$). Next, using (14.54), the contribution of productivity growth by continuing firms is

$$\begin{aligned}(\lambda-1)\phi z^* &= g^* - (\kappa-1)\hat{z}^*\eta(\hat{z}^*) \\ &= 0.0133.\end{aligned}$$

Therefore, in this benchmark parameterization, over two thirds of productivity growth comes from innovation by incumbents. Moreover, $\phi z^* = 0.0667$, so that there are on average 1.2 incremental innovations per year by an incumbent in a particular machine line ($r^*/\phi z^* \approx 1.2$). Using alternative values of the parameters κ, λ, β and α leads to broadly similar conclusions, though depending on the exact parameterization the contribution of entrants to productivity growth can be larger or smaller.

14.3.4. The Effects of Policy on Growth. Let us now use this model to analyze the effects of policies on equilibrium productivity growth and its decomposition between incumbents and entrants. Since the model has a Schumpeterian structure (with quality improvements as the engine of growth and creative destruction playing a major role), it may be conjectured that entry barriers (or taxes on potential entrance) will have negative effects on economic growth as in the baseline model earlier in this chapter. To investigate whether this is the case, let us suppose that there is a tax τ_e on R&D expenditure by entrants and a tax τ_i on R&D expenditure by incumbents (naturally, these can be taken to be negative and interpreted as subsidies as well). Note also that the tax on entrants, τ_e, can be interpreted as a more strict patent policy than the one in the baseline model, where the entrant did not have to pay the incumbent for partially benefiting from its accumulated knowledge. Nevertheless, to keep the analysis brief, I only focus on the case in which tax revenues are collected by the government rather than rebated back to the incumbent as patent fees.

Repeating the analysis above yields the following equilibrium conditions:

$$\eta(\hat{z}^*)V(\kappa q) = (1+\tau_e)q \text{ or } V(q) = \frac{q(1+\tau_e)}{\kappa\eta(\hat{z}^*)}. \tag{14.61}$$

The equation that determines the optimal R&D decisions of incumbents, (14.45), is also modified because of the tax rate τ_i and becomes

$$\phi(V(\lambda q) - V(q)) = (1+\tau_i)q. \tag{14.62}$$

Now combining (14.61) with (14.62),

$$\phi\left(\frac{(\lambda-1)(1+\tau_e)}{\kappa\eta(\hat{z}^*)(1+\tau_i)}\right) = 1.$$

Consequently, the BGP R&D level by entrants $\hat{z}^*$, when their R&D is taxed at the rate τ_e, is given by

$$\hat{z}^* \equiv \eta^{-1}\left(\frac{\phi(\lambda-1)(1+\tau_e)}{\kappa(1+\tau_i)}\right). \tag{14.63}$$

Equation (14.47) still applies, so that the BGP interest rate is $r^* = (1+\tau_e)^{-1}\kappa\eta(\hat{z}^*)\beta L - \hat{z}^*\eta(\hat{z}^*)$, which, by substituting for (14.63), can be written as

$$r^* = \frac{\phi(\lambda-1)\beta L}{1+\tau_i} - \eta(\hat{z}^*)\hat{z}^*,$$

and the BGP growth rate is

$$g^* = \frac{1}{\theta}\left(\frac{\phi(\lambda-1)\beta L}{1+\tau_i} - \eta(\hat{z}^*)\hat{z}^* - \rho\right),$$

with $\hat{z}^*$ given by (14.63). The following is now immediate:

PROPOSITION 14.7. *The growth rate of the economy is (strictly) decreasing in the tax rate on incumbents, i.e., $dg^*/d\tau_i < 0$, and is (strictly) increasing in the tax rate on entrants, i.e., $dg^*/d\tau_e > 0$.*

The result in this proposition is rather surprising and extreme. As shown above, in Schumpeterian models, making entry more difficult, either with entry barriers or by taxing R&D by entrants, has negative effects on economic growth. Despite the Schumpeterian nature of the current model, here blocking entry increases equilibrium growth. Moreover, as Exercise 14.24 shows, in the decentralized equilibrium of this economy there tends to be too much entry, so a tax on entry also tends to improve welfare. The intuition for this result is related to the main departure of this model from the standard Schumpeterian models. In contrast to the baseline Schumpeterian models, the engine of growth is still quality improvements, but these are undertaken both by incumbents and entrants. Entry barriers, by protecting incumbents, increase their value and greater value by incumbents encourages more R&D investments and faster productivity growth. Taxing entrants makes incumbents more profitable and this encourages further innovation by the incumbents. Taxes on entrants or entry barriers also further increase the contribution of incumbents to productivity growth.

Nevertheless, the result in this proposition should be interpreted with caution. The model in this section is special in that the R&D technology of incumbents is a linear. This linearity is important for the results in Proposition 14.7. Exercise 14.26 shows that the equilibrium can be characterized even when $\phi(z)$ is a concave function of z , and in this case, the effect of taxes on entrants is ambiguous, because it encourages R&D by incumbents and discourages R&D by entrants. Therefore, Proposition 14.7 should be read as emphasizing a particular new channel in the stark as possible way, not as a realistic description of how innovation will respond to tax policies.

14.4. Step-by-Step Innovations*

In the baseline Schumpeterian model and also in the extended Schumpeterian model of the previous section, new entrants could undertake innovation on any machine and did not need to have developed any knowhow on a particular line of business. This led to a simple structure, in many ways parallel to the models of expanding varieties studied in the previous chapter. However, quality improvements in practice may have a major cumulative

aspect. For example, it may be that only firms that have already reached a certain level of knowledge in a particular product or machine line can engage in further innovations. This is in fact consistent with qualitative accounts of technological change and competition in specific industries. Abernathy (1980, p. 70), for instance, concludes his study of a number of diverse industries by stating that: "each of the major companies seems to have made more frequent contributions in a particular area," and argues that this is because previous innovations in a field facilitate future innovations. This aspect is entirely missing from the baseline model of Schumpeterian growth, where any firm can engage in research to develop the next higher-quality machine (and in addition Arrow's replacement effect implies that incumbents do not undertake R&D, though this aspect was relaxed and generalized in the previous section). A more realistic description of the research process may involve only a few firms engaging in continuous and cumulative innovation, and in competition in a particular product or machine line.

This section presents a model of cumulative innovation of this type. Following Aghion, Harris, Howitt and Vickers (2001), we will refer to this as a model of *step-by-step innovation*. Such models are not only useful in providing a different conceptualization of the process of Schumpeterian growth, but they also enable us to endogenize the equilibrium market structure and allow a richer analysis of the effects of competition and intellectual property rights policy. Together with the model of innovation by incumbents and entrants presented in the previous section, this model enables us to have a framework in which existing firm (continuing establishments) contribute to productivity growth and build on their own past innovations (consistent with the empirical evidence as discussed in Section 18.1 in Chapter 18). In fact, the model in this section has a number of distinctive features relative to those presented so far in this and the previous chapter. For instance, these models predict that weaker patent protection and greater competition should reduce economic growth. Nevertheless, existing empirical evidence suggests that typically industries that are more competitive experience faster growth (or at the very least, there is a non-monotonic relationship between competition and economic growth, see, for example, Blundell (1999), Nickell (1999) and Aghion, Bloom, Blundell, Griffith and Howitt (2005)). Schumpeterian models with an endogenous market structure show that the effects of competition and intellectual property rights on economic growth are more complex, and greater competition (and weaker intellectual property rights protection) sometimes increases the growth rate of the economy.

14.4.1. Preferences and Technology. Consider the following continuous-time economy with a unique final good. The economy is populated by a continuum of measure 1 of individuals, each with 1 unit of labor endowment, which they supply inelastically. To simplify the analysis, I assume that the instantaneous utility function takes a logarithmic form. Thus the representative household has preferences given by

$$\int_0^\infty \exp\left(-\rho t\right) \log C\left(t\right) dt, \tag{14.64}$$

where $\rho > 0$ is the discount rate and $C\left(t\right)$ is consumption at date t.

Let $Y(t)$ be the total production of the final good at time t. The economy is closed and the final good is used only for consumption (there is no investment or spending on machines), so that $C(t) = Y(t)$. The standard Euler equation from (14.64) then implies that

$$g(t) \equiv \frac{\dot{C}(t)}{C(t)} = \frac{\dot{Y}(t)}{Y(t)} = r(t) - \rho, \tag{14.65}$$

where this equation defines $g(t)$ as the growth rate of consumption and thus output, and $r(t)$ is the interest rate at date t.

The final good Y is produced using a continuum 1 of intermediate goods according to the Cobb-Douglas production function

$$\ln Y(t) = \int_0^1 \ln y(\nu, t)\, d\nu, \tag{14.66}$$

where $y(\nu, t)$ is the output of νth intermediate at time t. Throughout, I take the price of the final good (or the ideal price index of the intermediates) as the numeraire and denote the price of intermediate ν at time t by $p^y(\nu, t)$. I also assume that there is free entry into the final good sector. These assumptions, together with the Cobb-Douglas production function (14.66), imply that each final good producer will have the following demand for intermediates

$$y(\nu, t) = \frac{Y(t)}{p^y(\nu, t)}, \text{ for all } \nu \in [0, 1]. \tag{14.67}$$

Intermediate $\nu \in [0, 1]$ come in two different varieties, each produced by one of two infinitely-lived firms. These two varieties are perfect substitutes and compete a la Bertrand. No other firm is able to produce in this industry. Firm $i = 1$ or 2 in industry ν has the following technology

$$y(\nu, t) = q_i(\nu, t)\, l_i(\nu, t) \tag{14.68}$$

where $l_i(\nu, t)$ is the employment level of the firm and $q_i(\nu, t)$ is its level of technology at time t. The only difference between the two firms is their technology, which will be determined endogenously. As in the models studied so far, each consumer in the economy holds a balanced portfolio of the shares of all firms. Consequently, the objective function of each firm is to maximize expected profits.

The production function for intermediate goods, (14.68), implies that the marginal cost of producing intermediate ν for firm i at time t is

$$MC_i(\nu, t) = \frac{w(t)}{q_i(\nu, t)} \tag{14.69}$$

where $w(t)$ is the wage rate in the economy at time t.

Let us denote the *technological leader* in each industry by i and the *follower* by $-i$, so that:

$$q_i(\nu, t) \geq q_{-i}(\nu, t).$$

Bertrand competition between the two firms implies that all intermediates will be supplied by the leader at the limit price (see Exercise 14.28):

$$p_i^y(\nu,t) = \frac{w(t)}{q_{-i}(\nu,t)}. \tag{14.70}$$

Equation (14.67) then implies the following demand for intermediates:

$$y(\nu,t) = \frac{q_{-i}(\nu,t)}{w(t)} Y(t). \tag{14.71}$$

R&D by the leader or the follower stochastically leads to innovation. We assume that when the leader innovates, its technology improves by a factor $\lambda > 1$. The follower, on the other hand, can undertake R&D to catch up with the frontier technology. Let us assume that because this innovation is for the follower's variant of the product and results from its own R&D efforts, it does not constitute infringement of the patent of the leader, and the follower does not have to make any payments to the technological leader in the industry.

R&D investments by the leader and the follower may have different costs and success probabilities. Nevertheless, I simplify the analysis here by assuming that they have the same costs and the same probability of success. In particular, in all cases, each firm (in every industry) has access to the following R&D technology (innovation possibilities frontier):

$$z_i(\nu,t) = \Phi(h_i(\nu,t)), \tag{14.72}$$

where $z_i(\nu,t)$ is the flow rate of innovation at time t and $h_i(\nu,t)$ is the number of workers hired by firm i in industry ν to work in the R&D process at t. Let us assume that Φ is twice continuously differentiable and satisfies $\Phi'(\cdot) > 0$, $\Phi''(\cdot) < 0$, $\Phi'(0) < \infty$ and that there exists $\bar{h} \in (0,\infty)$ such that $\Phi'(h) = 0$ for all $h \geq \bar{h}$. The assumption that $\Phi'(0) < \infty$ implies that there is no Inada condition when $h_i(\nu,t) = 0$. The last assumption, on the other hand, ensures that there is an upper bound on the flow rate of innovation. Recalling that the wage rate for labor is $w(t)$, the cost for R&D is therefore $w(t) G(z_i(\nu,t))$ where

$$G(z_i(\nu,t)) \equiv \Phi^{-1}(z_i(\nu,t)), \tag{14.73}$$

and the assumptions on Φ immediately imply that G is twice differentiable and satisfies $G'(\cdot) > 0$, $G''(\cdot) > 0$, $G'(0) > 0$ and $\lim_{z\to\bar{z}} G'(z) = \infty$, where $\bar{z} \equiv \Phi(\bar{h})$ is the maximal flow rate of innovation (with $\bar{h}$ defined above).

Let us next turn to the evolution of technologies within each industry. Suppose that leader i in industry ν at time t has a technology level of

$$q_i(\nu,t) = \lambda^{n_i(\nu,t)}, \tag{14.74}$$

and that the follower $-i$'s technology at time t is

$$q_{-i}(\nu,t) = \lambda^{n_{-i}(\nu,t)}, \tag{14.75}$$

where, naturally, $n_i(\nu,t) \geq n_{-i}(\nu,t)$. Let us denote the technology gap in industry ν at time t by $n(\nu,t) \equiv n_i(\nu,t) - n_{-i}(\nu,t)$. If the leader undertakes an innovation within a time interval of Δt, then the technology gap rises to $n(\nu,t+\Delta t) = n(\nu,t)+1$ (the probability of two or more innovations within the interval Δt is again $o(\Delta t)$). If, on the other hand, the

follower undertakes an innovation during the interval Δt, then $n(\nu, t+\Delta t)=0$. In addition, let us assume that there is an intellectual property rights (IPR) policy of the following form: the patent held by the technological leader expires at the exponential rate $\kappa<\infty$, in which case, the follower can close the technology gap.

Given this specification, the law of motion of the technology gap in industry ν can be expressed as

(14.76)
$$n(\nu,t+\Delta t)=\begin{cases} n(\nu,t)+1 & \text{with probability} \quad z_i(\nu,t)\Delta t+o(\Delta t) \\ 0 & \text{with probability} \quad (z_{-i}(\nu,t)+\kappa)\Delta t+o(\Delta t) \\ n(\nu,t) & \text{with probability} \quad 1-(z_i(\nu,t)+z_{-i}(\nu,t)+\kappa)\Delta t-o(\Delta t)), \end{cases}$$

where $o(\Delta t)$ again represents second-order terms, in particular, the probabilities of more than one innovations within an interval of length Δt. The terms $z_i(\nu,t)$ and $z_{-i}(\nu,t)$ are the flow rates of innovation by the leader and the follower, while κ is the flow rate at which the follower is allowed to copy the technology of the leader. In the first line, when $n(\nu,t)=0$, so that the two firms are *neck-and-neck*, $z_i(\nu,t)$ should be taken as equal to $2z(\nu,t)$, since the two firms will undertake the same amount of research effort given by $z(\nu,t)$ and the technology gap will increase to 1 if one of them is successful, which has probability $2z(\nu,t)$.

We next write the instantaneous "operating" profits for the leader (the profits exclusive of R&D expenditures and license fees). Profits of leader i in industry ν at time t are

$$\begin{aligned} \Pi_i(\nu,t) &= [p_i^y(\nu,t)-MC_i(\nu,t)]\,y_i(\nu,t) \\ &= \left(\frac{w(t)}{q_{-i}(\nu,t)}-\frac{w(t)}{q_i(\nu,t)}\right)\frac{Y(t)}{p_i^y(\nu,t)} \\ &= \left(1-\lambda^{-n(\nu,t)}\right)Y(t) \end{aligned} \tag{14.77}$$

where recall that $n(\nu,t)$ is the technology gap in industry j at time t. The first line simply uses the definition of operating profits as price minus marginal cost times quantity sold. The second line uses the fact that the equilibrium limit price of firm i is $p_i^y(\nu,t)=w(t)/q_{-i}(\nu,t)$ as given by (14.70), and the final equality uses the definitions of $q_i(\nu,t)$ and $q_{-i}(\nu,t)$ from (14.74) and (14.75). The expression in (14.77) also implies that there will be zero profits in an industry that is neck-and-neck, that is, in industries with $n(j,t)=0$. Followers also make zero profits, since they have no sales. Since the profits $\Pi_i(\nu,t)$ in (14.77) only depend on $Y(t)$ and the gap between the leader and the follower, $n(\nu,t)$, I sometimes denote them by $\Pi_n(t)$.

The Cobb-Douglas aggregate production function in (14.66) is responsible for the simple form of the profits (14.77), since it implies that profits only depend on the technology gap of the industry and aggregate output. This will simplify the analysis below by making the technology gap in each industry the only industry-specific payoff-relevant state variable.

The objective function of each firm is to maximize the net present discounted value of net profits (operating profits minus R&D expenditures and plus or minus patent fees). In

doing this, each firm will take the processes of interest rates, $[r(t)]_{t=0}^{\infty}$, aggregate output levels, $[Y(t)]_{t=0}^{\infty}$, wages, $[w(t)]_{t=0}^{\infty}$, the R&D decisions of all other firms and policies as given. Note that as in the baseline model of Schumpeterian growth in Section 14.1, even though technology and output in each sector are stochastic, total output, $Y(t)$, given by (14.66) is nonstochastic.

14.4.2. Equilibrium. Let $\boldsymbol{\mu}(t) \equiv \{\mu_n(t)\}_{n=0}^{\infty}$ denote the distribution of industries over different technology gaps, with $\sum_{n=0}^{\infty} \mu_n(t) = 1$. For example, $\mu_0(t)$ denotes the fraction of industries in which the firms are neck-and-neck at time t. Throughout, I focus on Markov Perfect Equilibria (MPE), where strategies are only functions of the payoff-relevant state variables. MPE is a natural equilibrium concept in this context, since it does not allow for implicit collusive agreements between the follower and the leader (see Appendix Chapter C for references and a further discussion of MPE). While such collusive agreements may be likely when there are only two firms in the industry, in most industries there are many more firms and also many potential entrants, making collusion more difficult. The focus on MPE allows us to drop the dependence on industry ν, thus we refer to R&D decisions by z_n for the technological leader that is n steps ahead and by z_{-n} for a follower that is n steps behind. Let us denote the list of decisions by the leader and follower with technology gap n at time t by $\xi_n(t) \equiv \langle z_n(t), p_i^y(\nu,t), y_i(\nu,t)\rangle$ and $\xi_{-n}(t) \equiv z_{-n}(t)$. Throughout, $\boldsymbol{\xi}$ will indicate the whole sequence of decisions at every state, $\boldsymbol{\xi}(t) \equiv \{\xi_n(t)\}_{n=-\infty}^{\infty}$.[6]

An allocation in this economy is then given by time paths of decisions for a leader that is $n = 0, 1, ..., \infty$ steps ahead, $[\xi_n(t)]_{t=0}^{\infty}$, time paths of R&D decisions for a follower that is $n = 1, ..., \infty$ steps behind, $\left[\xi_{-n}(t)\right]_{t=0}^{\infty}$, time path of wages and interest rates $[w(t), r(t)]_{t=0}^{\infty}$, and time paths of industry distributions over technology gaps $[\boldsymbol{\mu}(t)]_{t=0}^{\infty}$.

A MPE is represented by time paths $[\boldsymbol{\xi}^*(t), w^*(t), r^*(t), Y^*(t)]_{t=0}^{\infty}$ such that (i) $\left[p_i^{y*}(\nu,t)\right]_{t=0}^{\infty}$ and $[y_i^*(\nu,t)]_{t=0}^{\infty}$ implied by $[\boldsymbol{\xi}^*(t)]_{t=0}^{\infty}$ satisfy (14.70) and (14.71); (ii) R&D policies $[\mathbf{z}^*(t)]_{t=0}^{\infty}$ are best responses to themselves, i.e., $[\mathbf{z}^*(t)]_{t=0}^{\infty}$ maximizes the expected profits of firms taking aggregate output $[Y^*(t)]_{t=0}^{\infty}$, factor prices $[w^*(t), r^*(t)]_{t=0}^{\infty}$, and the R&D policies of other firms $[\mathbf{z}^*(t)]_{t=0}^{\infty}$ as given; (iii) aggregate output $[Y^*(t)]_{t=0}^{\infty}$ is given by (14.66); and (iv) the labor and capital markets clear at all times given the factor prices $[w^*(t), r^*(t)]_{t=0}^{\infty}$.

Let us now characterize this equilibrium. Since only the technological leader produces, labor demand in industry ν with technology gap $n(\nu,t) = n$ can be expressed as

$$l_n(t) = \frac{\lambda^{-n} Y(t)}{w(t)} \text{ for } n \geq 0. \tag{14.78}$$

[6]There are two sources of abuse of notation here. First, pricing and output decisions, given by (14.70) and (14.71), depend on the aggregate level of output $Y(t)$ as well. However, profits, as given by (14.77), and other choices do not depend on $Y(t)$, and I suppress this dependence without any affect on the analysis. Second, the sequences $[p_i^{y*}(\nu,t)]_{t=0}^{\infty}$ and $[y_i^*(\nu,t)]_{t=0}^{\infty}$ are stochastic, while the rest of the objects specified above are not. Since the stochastic nature of these sequences has no effect on the analysis, we suppress this feature as well.

In addition, there is demand for labor coming from R&D of both followers and leaders in all industries. Using (14.72) and the definition of the G function, we can express industry demands for R&D labor as

$$h_n(t) = \begin{cases} G(z_n(t)) + G(z_{-n}(t)) & \text{if } n \geq 1 \\ 2G(z_0(t)) & \text{if } n = 0 \end{cases}, \tag{14.79}$$

where $z_{-n}(t)$ refers to the R&D effort of a follower that is n steps behind. Moreover, this expression takes into account that in an industry with neck-and-neck competition, i.e., with $n = 0$, there will be twice the demand for R&D coming from the two "symmetric" firms.

The labor market clearing condition can then be expressed as:

$$1 \geq \sum_{n=0}^{\infty} \mu_n(t) \left[\frac{1}{\omega(t) \lambda^n} + G(z_n(t)) + G(z_{-n}(t)) \right], \tag{14.80}$$

and $\omega(t) \geq 0$, with complementary slackness, where

$$\omega(t) \equiv \frac{w(t)}{Y(t)} \tag{14.81}$$

is the labor share at time t. The labor market clearing condition, (14.80), uses the fact that total supply is equal to 1, and the demand cannot exceed this amount. If demand falls short of 1, then the wage rate, $w(t)$, and thus the labor share, $\omega(t)$, have to be equal to zero (though this will never be the case in equilibrium). The right-hand side of (14.80) consists of the demand for production (the terms with ω in the denominator), the demand for R&D workers from the neck-and-neck industries ($2G(z_0(t))$ when $n = 0$) and the demand for R&D workers coming from leaders and followers in other industries ($G(z_n(t)) + G(z_{-n}(t))$ when $n > 0$).

The relevant index of aggregate quality in this economy is no longer the average, but reflects the Cobb-Douglas aggregator in the production function,

$$\ln Q(t) \equiv \int_0^1 \ln q(\nu, t) \, d\nu. \tag{14.82}$$

Given this, the equilibrium wage can be written as (see Exercise 14.29):

$$w(t) = Q(t) \lambda^{-\sum_{n=0}^{\infty} n\mu_n(t)}. \tag{14.83}$$

14.4.3. Steady-State Equilibrium. Let us now focus on steady-state (Markov Perfect) equilibria, where the distribution of industries $\boldsymbol{\mu}(t) \equiv \{\mu_n(t)\}_{n=0}^{\infty}$ is stationary, $\omega(t)$ defined in (14.81) and g^*, the growth rate of the economy, is constant over time. If the economy is in steady state at time $t = 0$, then by definition, $Y(t) = Y_0 e^{g^* t}$ and $w(t) = w_0 \exp(g^* t)$. The two equations also imply that $\omega(t) = \omega^*$ for all $t \geq 0$. Let us assume that the parameters are such that the steady-state growth rate g^* is positive, but not large enough to violate the transversality conditions. This implies that net present values of each firm at all points in time will be finite and enable us to write the maximization problem of a leader that is $n > 0$ steps ahead recursively.

Standard arguments imply that the value function for a firm that is n steps ahead (or $-n$ steps behind) is given by

(14.84)
$$\begin{aligned} r(t) V_n(t) - \dot{V}_n(t) &= \max_{z_n(t)} \{ [\Pi_n(t) - w^*(t) G(z_n(t))] \\ &\quad + z_n(t) [V_{n+1}(t) - V_n(t)] + \left(z^*_{-n}(t) + \kappa\right) [V_0(t) - V_n(t)] \}. \end{aligned}$$

In steady state, the net present value of a firm that is n steps ahead, $V_n(t)$, will also grow at a constant rate g^* for all $n \geq 0$. Let us then define normalized values as

$$v_n(t) \equiv \frac{V_n(t)}{Y(t)} \tag{14.85}$$

for all n which will be independent of time in steady state, i.e., $v_n(t) = v_n$.

Using (14.85) and the fact that from (14.65), $r(t) = g(t) + \rho$, the steady-state value function (14.84) can be written as:

$$\begin{aligned} \rho v_n &= \max_{z_n} \{ \left(1 - \lambda^{-n}\right) - \omega^* G(z_n) + z_n [v_{n+1} - v_n], \\ &\quad + \left[z^*_{-n} + \kappa\right] [v_0 - v_n] \} \text{ for all } n \geq 1, \end{aligned} \tag{14.86}$$

where z^*_{-n} is the equilibrium value of R&D by a follower that is n steps behind, and ω^* is the steady-state labor share (while z_n is now explicitly chosen to maximize v_n).

Similarly the value for neck-and-neck firms is

$$\rho v_0 = \max_{z_0} \{ -\omega^* G(z_0) + z_0 [v_1 - v_0] + z^*_0 [v_{-1} - v_0] \}, \tag{14.87}$$

while the values for followers are given by

$$\rho v_{-n} = \max_{z_{-n}} \{ -\omega^* G(z_{-n}) + [z_{-n} + \kappa] [v_0 - v_{-n}] \}.$$

It is clear that these value functions and profit-maximizing R&D decision for followers should not depend on how many steps behind the leader they are, since a single innovation is sufficient to catch-up with the leader. Therefore,

$$\rho v_{-1} = \max_{z_{-1}} \{ -\omega^* G(z_{-1}) + [z_{-1} + \kappa] [v_0 - v_{-1}] \}, \tag{14.88}$$

where v_{-1} represents the value of any follower (regardless how many steps behind it is). The maximization problems involved in the value functions are straightforward and immediately yield the following profit-maximizing R&D decisions

$$z^*_n = \max \left\{ G'^{-1} \left(\frac{[v_{n+1} - v_n]}{\omega^*} \right), 0 \right\} \tag{14.89}$$

$$z^*_{-n} = \max \left\{ G'^{-1} \left(\frac{[v_0 - v_{-n}]}{\omega^*} \right), 0 \right\} \tag{14.90}$$

$$z^*_0 = \max \left\{ G'^{-1} \left(\frac{[v_1 - v_0]}{\omega^*} \right), 0 \right\}, \tag{14.91}$$

where $G'^{-1}(\cdot)$ is the inverse of the derivative of the G function, and since G is twice differentiable and strictly concave, G'^{-1} is differentiable and strictly increasing. These equations

therefore imply that innovation rates, the z_n^*'s, are increasing in the incremental value of moving to the next step and decreasing in the cost of R&D, as measured by the normalized wage rate, ω^*. Note also that since $G'(0) > 0$, these R&D levels can be equal to zero, which is taken care of by the max operator.

The response of innovation rates, z_n^*, to the increments in values, $v_{n+1} - v_n$, is the key economic force in this model. For example, a policy that reduces the patent protection of leaders that are $n+1$ steps ahead (by increasing κ) will make being $n+1$ steps ahead less profitable, thus reduce $v_{n+1} - v_n$ and z_n^*. This corresponds to the standard *disincentive effect* of relaxing IPR protection. However, relaxing IPR protection may also create a beneficial *composition effect*; this is because, typically, $\{v_{n+1} - v_n\}_{n=0}^{\infty}$ is a decreasing sequence, which implies that z_{n-1}^* is higher than z_n^* for $n \geq 1$ (see Proposition 14.11 below). Weaker patent protection (in the form of shorter patent lengths) will shift more industries into the neck-and-neck state and potentially increase the equilibrium level of R&D in the economy.

Given the equilibrium R&D decisions, the steady-state distribution of industries across states $\boldsymbol{\mu}^*$ has to satisfy the following accounting identities:

$$\left(z_{n+1}^* + z_{-1}^* + \kappa\right) \mu_{n+1}^* = z_n^* \mu_n^* \text{ for } n \geq 1, \tag{14.92}$$

$$\left(z_1^* + z_{-1}^* + \kappa\right) \mu_1^* = 2 z_0^* \mu_0^*, \tag{14.93}$$

$$2 z_0^* \mu_0^* = z_{-1}^* + \kappa. \tag{14.94}$$

The first expression equates exit from state $n+1$ (which takes the form of the leader going one more step ahead or the follower catching-up the leader) to entry into this state (which takes the form of a leader from the state n making one more innovation). The second equation, (14.93), performs the same accounting for state 1, taking into account that entry into this state comes from innovation by either of the two firms that are competing neck-and-neck. Finally, eq. (14.94) equates exit from state 0 with entry into this state, which comes from innovation by a follower in any industry with $n \geq 1$.

The labor market clearing condition in steady state can then be written as

$$1 \geq \sum_{n=0}^{\infty} \mu_n^* \left[\frac{1}{\omega^* \lambda^n} + G\left(z_n^*\right) + G\left(z_{-n}^*\right) \right] \text{ and } \omega^* \geq 0, \tag{14.95}$$

with complementary slackness.

The next proposition characterizes the steady-state growth rate in this economy:

PROPOSITION 14.8. *The steady-state growth rate is given by*

$$g^* = \ln \lambda \left[2 \mu_0^* z_0^* + \sum_{n=1}^{\infty} \mu_n^* z_n^* \right]. \tag{14.96}$$

PROOF. Equations (14.81) and (14.83) imply

$$Y(t) = \frac{w(t)}{\omega(t)} = \frac{Q(t) \lambda^{-\sum_{n=0}^{\infty} n \mu_n^*(t)}}{\omega(t)}.$$

Since $\omega(t) = \omega^*$ and $\{\mu_n^*\}_{n=0}^{\infty}$ are constant in steady state, $Y(t)$ grows at the same rate as $Q(t)$. Therefore,

$$g^* = \lim_{\Delta t \to 0} \frac{\ln Q(t + \Delta t) - \ln Q(t)}{\Delta t}.$$

During an interval of length Δt, in a fraction μ_n^* of the industries with technology gap $n \geq 1$ the leaders innovate at a rate $z_n^* \Delta t + o(\Delta t)$ and in the fraction μ_0^* of the industries with technology gap of $n = 0$, both firms innovate, so that the total innovation rate is $2z_0^* \Delta t + o(\Delta t)$). Since each innovation increases productivity by a factor λ, we obtain the preceding equation. Combining these observations,

$$\ln Q(t + \Delta t) = \ln Q(t) + \ln \lambda \left[2\mu_0^* z_0^* \Delta t + \sum_{n=1}^{\infty} \mu_n^* z_n^* \Delta t + o(\Delta t) \right].$$

Subtracting $\ln Q(t)$, dividing by Δt and taking the limit as $\Delta t \to 0$ gives (14.96). □

This proposition clarifies that the steady-state growth comes from two sources:

(1) R&D decisions of leaders or of firms in neck-and-neck industries.
(2) The distribution of industries across different technology gaps, $\boldsymbol{\mu}^* \equiv \{\mu_n^*\}_{n=0}^{\infty}$.

The latter channel reflects the composition effect discussed above. This type of composition effect implies that the relationship between competition and growth (or intellectual property rights protection and growth) is more complex than in the models discussed so far, because such policies will change the equilibrium market structure (the composition of industries).

A steady-state equilibrium is given by $\langle \boldsymbol{\mu}^*, \mathbf{v}, \mathbf{z}^*, \omega^*, g^* \rangle$ such that the distribution of industries $\boldsymbol{\mu}^*$ satisfy (14.92), (14.93) and (14.94), the values $\mathbf{v} \equiv \{v_n\}_{n=-\infty}^{\infty}$ satisfy (14.86), (14.87) and (14.88), the R&D decisions $\mathbf{z}^*$ are given by, (14.89), (14.90) and (14.91), the steady-state labor share ω^* satisfies (14.95) and the steady-state growth rate g^* is given by (14.96). I next provide a characterization of the steady-state equilibrium. The first result is a technical one that is necessary for this characterization.

PROPOSITION 14.9. *In a steady state equilibrium, $v_{-1} \leq v_0$ and $\{v_n\}_{n=0}^{\infty}$ forms a bounded and strictly increasing sequence converging to some positive value v_∞.*

PROOF. Let $\{z_n\}_{n=-1}^{\infty}$ be the R&D decisions of a firm and $\{v_n\}_{n=-1}^{\infty}$ be the sequence of values, taking the decisions of other firms and the industry distributions, $\{z_n^*\}_{n=-1}^{\infty}$, $\{\mu_n^*\}_{n=-1}^{\infty}$, ω^* and g^*, as given. By choosing $z_n = 0$ for all $n \geq -1$, the firm guarantees $v_n \geq 0$ for all $n \geq -1$. Let π_n denote normalized flow profits when the technology gap is n. Then, since $\pi_n \leq 1$ for all $n \geq -1$, $v_n \leq 1/\rho$ for all $n \geq -1$, establishing that $\{v_n\}_{n=-1}^{\infty}$ is a bounded sequence, with $v_n \in [0, 1/\rho]$ for all $n \geq -1$.

Proof of $v_1 > v_0$: Suppose, first, $v_1 \leq v_0$, then (14.91) implies $z_0^* = 0$, and by the symmetry of the problem in equilibrium, (14.87) implies $v_0 = v_1 = 0$. As a result, from (14.90) we obtain $z_{-1}^* = 0$. Equation (14.86) then implies that when $z_{-1}^* = 0$, $v_1 \geq \left(1 - \lambda^{-1}\right) / (\rho + \kappa) > 0$, yielding a contradiction and proving that $v_1 > v_0$.

Proof of $v_{-1} \leq v_0$: Suppose, to obtain a contradiction, that $v_{-1} > v_0$. Then, (14.90) implies $z^*_{-1} = 0$, which leads to $v_{-1} = \kappa v_0/(\rho + \kappa)$, contradicting $v_{-1} > v_0$ since $\kappa/(\rho + \kappa) < 1$ (given that $\kappa < \infty$).

Proof of $v_n < v_{n+1}$: Suppose, to obtain a contradiction, that $v_n \geq v_{n+1}$. Now (14.89) implies $z^*_n = 0$, and (14.86) becomes

$$\rho v_n = \left(1 - \lambda^{-n}\right) + z^*_{-1}\left[v_0 - v_n\right] + \kappa\left[v_0 - v_n\right]. \tag{14.97}$$

Also from (14.86), the value for state $n+1$ satisfies

$$\rho v_{n+1} \geq \left(1 - \lambda^{-n-1}\right) + z^*_{-1}\left[v_0 - v_{n+1}\right] + \kappa\left[v_0 - v_{n+1}\right]. \tag{14.98}$$

Combining the two previous expressions,

$$\begin{aligned}&\left(1 - \lambda^{-n}\right) + z^*_{-1}\left[v_0 - v_n\right] + \kappa\left[v_0 - v_n\right] \geq \\ &1 - \lambda^{-n-1} + z^*_{-1}\left[v_0 - v_{n+1}\right] + \kappa\left[v_0 - v_{n+1}\right].\end{aligned}$$

Since $\lambda^{-n-1} < \lambda^{-n}$, this implies $v_n < v_{n+1}$, contradicting the hypothesis that $v_n \geq v_{n+1}$, and establishing the desired result, $v_n < v_{n+1}$.

Consequently, $\{v_n\}_{n=-1}^{\infty}$ is nondecreasing and $\{v_n\}_{n=0}^{\infty}$ is (strictly) increasing. Since a nondecreasing sequence in a compact set must converge, $\{v_n\}_{n=-1}^{\infty}$ converges to its limit point, v_∞, which must be strictly positive, since $\{v_n\}_{n=0}^{\infty}$ is strictly increasing and has a nonnegative initial value. This completes the proof. □

A potential difficulty in the analysis of the current model is that we have to determine R&D levels and values for an infinite number of firms, since the technology gap between the leader and the follower can, in principle, take any value. However, the next result shows that we can restrict attention to a finite sequence of values:

PROPOSITION 14.10. *There exists* $n^* \geq 1$ *such that* $z^*_n = 0$ *for all* $n \geq n^*$.

PROOF. See Exercise 14.30. □

The next proposition provides the most important economic insights of this model and shows that $\mathbf{z}^* \equiv \{z^*_n\}_{n=0}^{\infty}$ is a decreasing sequence, which implies that technological leaders that are further ahead undertake less R&D. Intuitively, the benefits of further R&D investments are decreasing in the technology gap, since greater values of the technology gap translate into smaller increases in the equilibrium markup (recall (14.77)). The fact that leaders that are sufficiently ahead of their competitors undertake little R&D is the main reason why composition effects play an important role in this model. For example, all else equal, closing the technology gap across a range of industries will increase R&D spending and equilibrium growth (though, as discussed in the previous section, this may not always increase welfare, especially if there is a strong business stealing effect).

PROPOSITION 14.11. *In any steady-state equilibrium,* $z^*_{n+1} \leq z^*_n$ *for all* $n \geq 1$ *and moreover,* $z^*_{n+1} < z^*_n$ *if* $z^*_n > 0$. *Furthermore,* $z^*_0 > z^*_1$ *and* $z^*_0 \geq z^*_{-1}$.

PROOF. From eq. (14.89),

$$\delta_{n+1} \equiv v_{n+1} - v_n < v_n - v_{n-1} \equiv \delta_n \tag{14.99}$$

is sufficient to establish that $z^*_{n+1} \leq z^*_n$.

Let us write:

$$\bar{\rho} v_n = \max_{z_n} \left\{ \left(1 - \lambda^{-n}\right) - \omega^* G\left(z_n\right) + z_n \left[v_{n+1} - v_n\right] + \left(z^*_{-1} + \kappa\right) v_0 \right\}, \tag{14.100}$$

where $\bar{\rho} \equiv \rho + z^*_{-1} + \kappa$. Since z^*_{n+1}, z^*_n and z^*_{n-1} are maximizers of the value functions v_{n+1}, v_n and v_{n-1}, (14.100) implies:

$$\begin{aligned}
\bar{\rho} v_{n+1} &= 1 - \lambda^{-n-1} - \omega^* G\left(z^*_{n+1}\right) + z^*_{n+1}\left[v_{n+2} - v_{n+1}\right] + \left(z^*_{-1} + \kappa\right) v_0, \\
\bar{\rho} v_n &\geq 1 - \lambda^{-n} - \omega^* G\left(z^*_{n+1}\right) + z^*_{n+1}\left[v_{n+1} - v_n\right] + \left(z^*_{-1} + \kappa\right) v_0, , \\
\bar{\rho} v_n &\geq 1 - \lambda^{-n} - \omega^* G\left(z^*_{n-1}\right) + z^*_{n-1}\left[v_{n+1} - v_n\right] + \left(z^*_{-1} + \kappa\right) v_0, , \\
\bar{\rho} v_{n-1} &= 1 - \lambda^{-n+1} - \omega^* G\left(z^*_{n-1}\right) + z^*_{n-1}\left[v_n - v_{n-1}\right] + \left(z^*_{-1} + \kappa\right) v_0.
\end{aligned} \tag{14.101}$$

Now taking differences with $\bar{\rho} v_n$ and using the definition of δ_n's,

$$\begin{aligned}
\bar{\rho}\delta_{n+1} &\leq \lambda^{-n}\left(1 - \lambda^{-1}\right) + z^*_{n+1}\left(\delta_{n+2} - \delta_{n+1}\right) \\
\bar{\rho}\delta_n &\geq \lambda^{-n+1}\left(1 - \lambda^{-1}\right) + z^*_{n-1}\left(\delta_{n+1} - \delta_n\right).
\end{aligned}$$

Therefore,

$$\left(\bar{\rho} + z^*_{n-1}\right)\left(\delta_{n+1} - \delta_n\right) \leq -k_n + z^*_{n+1}\left(\delta_{n+2} - \delta_{n+1}\right),$$

where $k_n \equiv (\lambda - 1)^2 \lambda^{-n-1} > 0$. Now to obtain a contradiction, suppose that $\delta_{n+1} - \delta_n \geq 0$. From the previous equation, this implies $\delta_{n+2} - \delta_{n+1} > 0$ since k_n is strictly positive. Repeating this argument successively, we have that if $\delta_{n'+1} - \delta_{n'} \geq 0$, then $\delta_{n+1} - \delta_n > 0$ for all $n \geq n'$. However, we know from Proposition 14.9 that $\{v_n\}_{n=0}^{\infty}$ is strictly increasing and converges to a constant v_∞. This implies that $\delta_n < 0$ for sufficiently large n (with $\delta_n \uparrow 0$), which contradicts the hypothesis that $\delta_{n+1} - \delta_n \geq 0$ for all $n \geq n' \geq 0$, and establishes that $z^*_{n+1} \leq z^*_n$. To see that the inequality is strict when $z^*_n > 0$, it suffices to note that (14.99) has already been established, i.e., $\delta_{n+1} - \delta_n < 0$, thus if eq. (14.89) has a positive solution, then we necessarily have $z^*_{n+1} < z^*_n$.

Proof of $z^*_0 \geq z^*_{-1}$: (14.87) can be written as

$$\rho v_0 = -\omega^* G\left(z^*_0\right) + z^*_0\left[v_{-1} + v_1 - 2v_0\right]. \tag{14.102}$$

We have $v_0 \geq 0$ from Proposition 14.9. Suppose $v_0 > 0$. Then, (14.102) implies $z^*_0 > 0$ and

$$\begin{aligned}
v_{-1} + v_1 - 2v_0 &> 0 \\
v_1 - v_0 &> v_0 - v_{-1}.
\end{aligned} \tag{14.103}$$

This inequality combined with (14.91) and (14.90) yields $z^*_0 > z^*_{-1}$. Suppose next that $v_0 = 0$. Inequality (14.103) now holds as a weak inequality and implies that $z^*_0 \geq z^*_{-1}$. Moreover, since $G(\cdot)$ is strictly convex and z^*_0 is given by (14.91), (14.102) then implies $z^*_0 = 0$ and thus $z^*_{-1} = 0$.

Proof of $z_0^* > z_1^*$: See Exercise 14.31. □

This proposition therefore shows that the highest amount of R&D is undertaken in neck-and-neck industries. This explains why composition effects can increase aggregate innovation. Exercise 14.32 shows how a relaxation of intellectual property rights protection can increase the growth rate in the economy.

So far, I have not provided a closed-form solution for the growth rate in this economy. It turns out that this is generally not possible, because of the endogenous market structure in these types of models. Nevertheless, it can be proved that a steady-state equilibrium exists in this economy, though the proof is somewhat more involved and does not generate additional insights for our purposes (see Acemoglu and Akcigit, 2006).

An important feature of this model is that equilibrium markups are endogenous and evolve over time as a function of competition between the firms producing in the same product line. More importantly, Proposition 14.11 implies that when a particular firm is sufficiently ahead of its rival, it undertakes less R&D. Therefore, this model, contrary to the baseline Schumpeterian model and also contrary to all expanding varieties models, implies that greater competition may lead to higher growth rates. Greater competition generated by closing the gap between the followers and leaders induces the leaders to undertake more R&D in order to escape the competition from the followers.

14.5. Taking Stock

This chapter introduced the basic Schumpeterian model of economic growth to emphasize the importance of competition among firms both in the innovation process and in the product market. Schumpeterian growth introduces a process of creative destruction, where new products or machines replace older models and new firms replace incumbent producers.

The baseline model features process innovations leading to quality improvements. The description of economic growth that emerges from this model is, in many ways, more realistic than the expanding variety models. In particular, technological progress does not always correspond to new products or machines complementing existing ones, but involves the creation of higher-quality producers replacing incumbents. Arrow's replacement effect, discussed in Chapter 12, implies that there is a strong incentive for new entrants to undertake research because the new, higher-quality products will replace the products of incumbents, leading to Schumpeterian creative destruction as the engine of economic growth. Even though the description of economic growth in this model is richer than the expanding varieties model, the mathematical structure turns out to be quite similar to the models in the previous chapters. In reduced form, the model again resembles an AK economy. The main difference is that now the growth rate of the economy, through the rate of replacement of old products, affects the value of innovation. Nevertheless, in the baseline version of the model, the effects of various policy interventions are the same as in the expanding product variety model of the previous chapter.

An important insight of Schumpeterian models is that growth comes with potential conflict of interest. The process of creative destruction destroys the monopoly rents of previous incumbents. This raises the possibility that distortionary policies may arise endogenously as a way of protecting the rents of politically powerful incumbents. Models of creative destruction therefore raise the political economy issues that are central for understanding the fundamental causes of economic growth and provide us insights about both the endogenous nature of technology and about the potential resistance to technological change.

Schumpeterian models also enable us to make greater contact with the industrial organization of innovation. In particular, the process of creative destruction implies that the market structure may be evolving endogenously over time. Nevertheless, the baseline Schumpeterian models have a number of shortcomings, and addressing these is an interesting and important area for future research. An important discrepancy between the baseline models and the data is that, while the models predict all productivity growth to be driven by the process of creative destruction and entry, in the data much of productivity growth is driven by innovation by continuing firms and plants. Section 14.3 provided a first look at how the baseline models can be extended to account for these patterns and to provide a richer framework for the analysis of the industrial organization of innovation. A second important shortcoming of the baseline models is that they predict that markups are constant and there is always a single firm supplying the entire market. These implications can also be relaxed by considering a richer framework, for example, by allowing cumulative or step-by-step innovation and competition between multiple firms that engage in innovation. Section 14.4 showed how the baseline model can be extended to incorporate step-by-step (cumulative) innovations and how such a model leads to endogenously evolving monopoly markups and market structure. Perhaps more interestingly, in models that incorporate different aspects of the industrial organization of innovation, the effects of competition and patent protection on economic growth are potentially quite different from both the baseline model of Schumpeterian growth and from models of expanding varieties. This suggests that Schumpeterian models might provide a useful framework for the analysis of a range of industrial policies, including anti-trust policies, licensing and intellectual property rights policies.

14.6. References and Literature

The baseline model of Schumpeterian growth presented in Section 14.1 is based on the work by Aghion and Howitt (1992). Similar models have also been developed by Segerstrom, Anant and Dinopoulos (1990), Grossman and Helpman (1991a,b). Aghion and Howitt (1998) provide an excellent survey of many Schumpeterian models of economic growth and numerous extensions. The specific modeling assumptions made in the presentation here draw on Acemoglu (1998), which also uses the aggregate production function with proportional quality improvements and costs of production and R&D increasing proportionally with quality. The original Aghion and Howitt (1992) approach is very similar to that used in Section 14.2.

Aghion and Howitt (1992) also discuss uneven growth and potential growth cycles, which were presented in Section 14.2. Uneven growth and cycles are also possible in expanding product or input variety models as shown by Matsuyama (1999, 2001). I only discussed the possibility of such cycles in the context of Schumpeterian growth, since the forces leading to such cycles are more pronounced in these models.

The effect of creative destruction on unemployment is first studied in Aghion and Howitt (1994). The implications of creative destruction for firm-specific investments are discussed in Francois and Roberts (2001) and in Martimort and Verdier (2003).

The model in Section 14.3 draws on Acemoglu (2008b) and is a first attempt to introducing productivity growth driven both by incumbents and entrants (see also Barro and Sala-i-Martin, 2004, for a model in which incumbents undertake R&D because they have cost advantage). A related paper is Klette and Kortum (2004). Klette and Kortum construct a richer model of firm and aggregate innovation dynamics based on expanding product varieties. Their key assumption is that firms with more products have an advantage in discovering more new products. With this assumption, their model generates the same patterns of firm growth as the simple model and Section 14.3, and also matches certain additional facts about propensity to patent and the differential survival probabilities of firms by size. One disadvantage of this approach is that the firm size distribution is not driven by the dynamics of continuing plans (in fact, if new products are interpreted as new plants, the Klette-Kortum model predicts that all productivity growth will be driven by entry of new plants, though this may be an extreme interpretation, since some new products may be produced in existing plants). Lentz and Mortensen (2006) extend Klette and Kortum's model by introducing additional sources of heterogeneity and estimate this extended model on Danish data. Klepper (1996) documents various facts about the firm size, entry and exit decisions and innovation, and provides a simple descriptive model that can account for these facts. None of these papers consider a Schumpeterian growth featuring innovation both by incumbents and entrants that can be easily mapped to decomposing the contribution of new and continuing plants (firms) to productivity growth.

Step-by-step or cumulative innovations have been analyzed in Aghion, Harris and Vickers (1999) and Aghion, Harris, Howitt and Vickers (2001). The model presented here is a simplified version of Acemoglu and Akcigit (2006), which includes a detailed analysis of the implications of intellectual property rights policy and licensing in this class of models. The proof of existence of a steady-state equilibrium under a somewhat more general environment is provided in that paper. The notion of Markov Perfect Equilibrium used in Section 14.4 is a standard equilibrium concept in dynamic games and is a refinement of subgame perfect equilibrium, that restricts strategies to depend only on payoff-relevant state variables. This concept is further discussed in the Appendix Chapter C and in Fudenberg and Tirole (1994).

Blundell (1999), Nickell (1999) and Aghion, Bloom, Blundell, Griffith and Howitt (2005) provide evidence that greater competition may encourage economic growth and technological progress. The latter paper shows that industries where the technology gap between firms

is smaller are typically more innovative. Aghion, Harris, Howitt and Vickers (2001) and Aghion, Bloom, Blundell, Griffith and Howitt (2005) show that in step-by-step models of innovation greater competition may increase growth. Aghion, Dewatripont and Ray (2000) provide another reason why competition may encourage growth. In their model competitive pressure improves managerial incentives and efficiency.

14.7. Exercises

EXERCISE 14.1. Prove that in the baseline model of Schumpeterian growth in Section 14.1, BGP R&D towards different types of machines must be equal to some z^*. [Hint: use (14.13) together with (14.12)].

EXERCISE 14.2.
(1) Prove that in the baseline model of Schumpeterian growth in Section 14.1, all R&D will be undertaken by entrants, and there will never be R&D by incumbents. [Hint: rewrite (14.12) by allowing for a choice of R&D investments].
(2) Now suppose that the flow rate of success of R&D is $\phi\eta/q$ for an incumbent as opposed to η/q for an entrant. Show that for any value of $\phi \leq 1/(\lambda - 1)$, the incumbent will still choose zero R&D. Explain this result.

EXERCISE 14.3. The baseline endogenous technological change models, including the model of Schumpeterian growth in this chapter, assume that new products are protected by perpetual patents. This exercise asks you to show that this is not strictly necessary in the logic of these models. Suppose that there is no patent protection for any innovation, but copying an innovator requires a fixed cost $\varepsilon > 0$. Any firm, after paying this cost, has access to the same technology as the innovator. Prove that in this environment there will be no copying and all the results of the model with fully-enforced perpetual patents apply.

EXERCISE 14.4. Complete the proof of Proposition 14.1. In particular, verify that the equilibrium growth rate is unique, strictly positive and such that the transversality condition (14.15) is satisfied.

EXERCISE 14.5. Prove Proposition 14.2.

EXERCISE 14.6. Modify the baseline model of Section 14.1 so that the aggregate production function for the final good is

$$Y(t) = \frac{1}{1-\beta}\left[\int_0^1 (q(\nu,t)x(\nu,t \mid q))^{1-\beta}\, d\nu\right] L^{\beta}.$$

All the other features of the model remain unchanged.

(1) Show that with this production function, a BGP does not exist. Explain why this is.
(2) What would you change in the model to ensure the existence of a BGP.

EXERCISE 14.7. * In the baseline Schumpeterian growth model, instead of (14.3), suppose that the production function of the final good sector is given by

$$Y(t) = \frac{1}{1-\beta}\left[\int_0^1 q(\nu,t)^{\zeta_1} x(\nu,t \mid q)^{1-\beta} d\nu\right] L^{\beta}.$$

Suppose also that producing one unit of an intermediate could of quality q costs ψq^{ζ_2} and that one unit of final good devoted to research on the machine line with quality q generates a flow rate of innovation equal to η/q^{ζ_3}. Characterize the equilibrium of this economy and determine what combinations of the parameters ζ_1, ζ_2 and ζ_3 will ensure balanced growth.

Exercise 14.8. (1) Verify that Theorem 7.14 from Chapter 7 can be applied to the social planner's problem in the baseline Schumpeterian model of Section 14.1.

(2) Derive eq. (14.24).

Exercise 14.9. Show that condition (14.5) is sufficient to ensure that a firm that innovates will set the unconstrained monopoly price. [Hint: first suppose that the innovator sets the monopoly price $\psi q/(1-\beta)$ for a product of quality q. Then, consider the firm with the next highest quality, $\lambda^{-1}q$. Suppose that this firm sells at marginal cost, $\psi\lambda^{-1}q$. Then, find the value of λ such that final good producers are indifferent between buying a machine of quality q at the price $\psi q/(1-\beta)$ versus a machine of quality $\lambda^{-1}q$ at the price $\psi\lambda^{-1}q$.]

Exercise 14.10. Analyze the baseline model of Schumpeterian growth in Section 14.1 assuming that (14.5) is not satisfied.

(1) Show that monopolists will set a limit price.
(2) Characterize the BGP equilibrium growth rate.
(3) Characterize the Pareto optimal allocation and compare it to the equilibrium allocation. How does the comparison differ from the case in which innovations were drastic?
(4) Now consider a hypothetical economy in which the previous highest-quality producer disappears so that the monopolist can charge a markup of $1/(1-\beta)$ instead of the limit price. Show that the BGP growth rate in this hypothetical economy is strictly greater than the growth rate characterized in 2 above. Explain this result.

Exercise 14.11. Suppose that there is constant exponential population growth at the rate n. Modify the baseline model of Section 14.1 so that there is no scale effect and the economy grows at the constant rate (with positive growth of income per capita). [Hint: suppose that one unit of final good spent on R&D for improving the machine of quality q leads to flow rate of innovation equal to η/q^{ϕ}, where $\phi > 1$].

Exercise 14.12. * Consider a version of the model of Schumpeterian growth in which the x's do not depreciate fully after use (similar to Exercise 13.22 in the previous chapter). Preferences and the rest of the production structure are the same as in the baseline model in Section 14.1.

(1) Define the equilibrium in BGP allocations.
(2) Formulate the maximization problem of a monopolist with the highest quality machine.
(3) Show that, contrary to Exercise 13.22, the results are different than those in Section 14.1. Explain why depreciation of machines was not important in the expanding varieties model, but it is important in Schumpeterian models.

Exercise 14.13. Consider a version of the model of Schumpeterian growth in which innovations reduce costs instead of increasing quality. In particular, suppose that the aggregate production function is given by

$$Y(t) = \frac{1}{1-\beta}\left[\int_0^1 x(\nu,t)^{1-\beta}d\nu\right]L^{\beta},$$

and the marginal cost of producing machine variety ν at time t is given by $MC(\nu,t)$. Every innovation reduces this cost by a factor λ.

(1) Define the equilibrium and BGP allocations.
(2) Specify a form of the innovation possibilities frontier that is consistent with balanced growth.
(3) Derive the BGP growth rate of the economy and show that there are no transitional dynamics.
(4) Compare the BGP growth rate to the Pareto optimal growth rate of the economy. Can there be excessive innovations?
(5) What are the similarities and differences between this model and the baseline model presented in Section 14.1.

Exercise 14.14. Consider the model in Section 14.2, with R&D performed by workers. Suppose instead that the aggregate production function for the final good is given by

$$Y(t) = \frac{1}{1-\beta}\left[\int_0^1 q(\nu,t)x(\nu,t \mid q)^{1-\beta}d\nu\right]L_E(t)^{\beta},$$

where $L_E(t)$ denotes the number of workers employed in final good production at time t.

(1) Show that in this case, there will only be R&D for the machine with the highest $q(\nu,t)$.
(2) How would you modify the model so that the equilibrium has balanced R&D across sectors?

Exercise 14.15. Consider the model of Schumpeterian growth in Section 14.1, with one difference: conditional on success an innovation generates a random improvement of λ over the previous technology, where the distribution function of λ is $H(\lambda)$ and has support $\left[(1-\beta)^{-(1-\beta)/\beta}, \bar{\lambda}\right]$.

(1) Define the equilibrium and BGP allocations.
(2) Characterize the BGP and specify restrictions on parameters so that the transversality condition is satisfied.
(3) Why did we assume that the lower support of λ is $(1-\beta)^{-1}$? How would the analysis change if this were relaxed?
(4) Show that there are no transitional dynamics in this economy.
(5) Compare the BGP growth rate to the Pareto optimal growth rate of the economy. Can there be excessive innovations?

Exercise 14.16. In the model of Section 14.2 show that the economy experiences no growth of output for intervals of average length $1/\eta L_R^*$.

Exercise 14.17. (1) Prove Proposition 14.4, in particular verifying that the allocation described there is unique, that the average growth rate is given by $g^* = \ln \lambda \eta L_R^*$ and that condition (14.27) is necessary and sufficient for the existence of the equilibrium described in the proposition.

(2) Explain why the growth rate features $\ln \lambda$ rather than $\lambda - 1$ as in the model of Section 14.1.

Exercise 14.18. Consider the one-sector Schumpeterian model in discrete time. Suppose as in the model in Section 14.2 that consumers are risk neutral, there is no population growth, and the final good sector has the production function given by (14.25). There is a linear production technology for intermediate goods whereby any type of intermediate good (that has been invented) can be produced at the marginal cost of ψ units of the final good. Assume also that the R&D technology is such that $L_R > 0$ workers employed in research at time t will necessarily lead to an innovation for time $t + 1$, and the number of workers used in research simply determines the quality of the innovation via the function $\Lambda(L_R)$, that is, if date t quality is q, the new, date $t+1$ machine will have quality $q' = \Lambda(L_R) q$. Assume that there will be innovation only if $L_R > 0$ and $\Lambda(\cdot)$ is strictly increasing, differentiable, strictly concave and satisfies the appropriate Inada conditions.

(1) Define equilibrium and BGP allocations.

(2) Characterize the BGP and specify restrictions on parameters so that the transversality condition is satisfied. [Hint: to simplify the algebra, you may wish to assume that once the new machine is invented, the old one cannot be used any longer, so that there is no limit pricing].

(3) Compare the BGP growth rate to the Pareto optimal growth rate of the economy. Show that the size of innovations is always too small relative to the size of innovations in the Pareto optimal allocation.

Exercise 14.19. * Consider the one-sector Schumpeterian model in discrete time analyzed in the previous exercise, except that now $\Lambda(\cdot)$ denotes the probability of innovation and each innovation improves the quality of a machine q to λq, where $\lambda > 1$. Suppose that when a new innovation arrives a fraction φ of workers employed in the final good production will not be able to adapt to this new technology and will need to remain unemployed for one time period to "retool".

(1) Define the equilibrium and BGP allocations. [Hint: also specify the number of unemployed workers in equilibrium].

(2) Defined the appropriate generalization of the BGP for this economy and determine the number of unemployed workers in this equilibrium.

(3) Show that the economy will experience bursts of unemployment, followed by periods of full employment.

(4) Show that a decline in ρ will increase the average growth rate and the average unemployment rate in the economy.

Exercise 14.20. * Derive eq.'s (14.28)-(14.29).

Exercise 14.21. * Consider the model discussed in subsection 14.2.2.

(1) Choose a functional form for $\eta(\cdot)$ such that eq.'s (14.29) have solutions L_R^1 and $L_R^2 \neq L_R^1$. Explain why, when such solutions exist, there is a perfect foresight equilibrium with two-period endogenous cycles.
(2) Show that even when solutions exist, there also exists a steady-state equilibrium with constant research.
(3) Show that when such solutions do not exist, there exists an equilibrium which exhibits oscillatory transitional dynamics converging to the steady state characterized in 2 above.

Exercise 14.22. * Show that the qualitative results of the model in subsection 14.2.2 generalize when there is a single firm undertaking research, thus internalizing the effect of L_R on $\eta(L_R)$.

Exercise 14.23. Suppose that in the model of Section 14.3 incumbents also have access to the radical innovation technology used by entrants. Show that there cannot existing equilibrium in which incumbents undertake positive R&D with this technology. [Hint: use the free-entry condition for entrants together with to condition that makes such investments profitable for incumbents and derive a contradiction].

Exercise 14.24. Set up the social planner's problem (of maximizing the utility of the representative household) in Section 14.3.

(1) Show that this maximization problem corresponds to a concave current-value Hamiltonian and derive the unique solution to this problem. Show that this solution involves the consumption of the representative household growing at a constant rate at all points.
(2) Show that the social planner will tend to increase growth because she avoids the monopoly markup over machines.
(3) Show that the social planner will tend to choose lower entry because of the negative externality in the research process.
(4) Give numerical examples in which the growth rate in the Pareto optimal allocation is greater than or less than the decentralized growth rate.

Exercise 14.25. Consider the model of Section 14.3 and suppose that the R&D technology of the incumbents for innovation is such that if an incumbent with a machine of quality q spends an amount zq for incremental innovations, then the flow rate of innovation is $\phi(z)$ (and this innovation again increases the quality of the incumbent's machine to λq). Assume that $\phi(z)$ is strictly increasing, strictly concave, differentiable, and satisfies $\lim_{z\to 0}\phi'(z) = \infty$ and $\lim_{z\to\infty}\phi'(z) = 0$.

(1) Focus on steady-state equilibria and conjecture that $V(q) = vq$. Using this conjecture, show that incumbents will choose R&D intensity z^* such that $(\lambda - 1)v = \phi'(z^*)$. Combining this equation with the free-entry condition for entrants and the

equation for growth rate given by (14.54), show that there exists a unique BGP equilibrium (under the conjecture that $V(q)$ is linear).

(2) Is it possible for an equilibrium to involve different levels of z for incumbents with different quality machines?

(3) In light of your answer to 2, what happens if we consider the "limiting case" of this model where $\phi(z)$= constant?

(4) Show that this equilibrium involves positive R&D both by incumbents and entrants.

(5) Now introduce taxes on R&D by incumbents and entrants at the rates τ_i and τ_e. Show that, in contrast to the results in Proposition 14.7, the effects of both taxes on growth are ambiguous. What happens if $\eta(z)$ =constant?

EXERCISE 14.26. * Prove Proposition 14.6 for the case in which $z(\nu, t \mid q)$ could differ across incumbents with different levels of q. Show that the same BGP as in Proposition 14.6 applies and is essentially unique, in the sense that average incumbent R&D effort is always equal to z^*.

EXERCISE 14.27. Consider the model of Section 14.3, but modify the production function to $Y(t) = \left[\int_0^1 q(\nu,t)x(\nu,t \mid q)^{1-\beta}d\nu\right] L^{\beta}/(1-\beta)$ and assume that production of an input of quality q requires ψq units of the final good as in the baseline model of Section 14.1. Show that the equilibrium growth rate and the decomposition of productivity growth between incumbents and entrants are identical in this case to the results in Section 14.3, but there are no firm size dynamics. Explain why this is. Are there dynamics of profits? How does the distribution of profits across firm evolve over time?

EXERCISE 14.28. * Derive eq. (14.70).

EXERCISE 14.29. * Derive eq. (14.83). [Hint: write $\ln Y(t) = \int_0^1 \ln q(\nu,t)\, l(\nu,t)\, d\nu = \int_0^1 \left[\ln q_i(\nu,t) + \ln \frac{Y(t)}{w(t)}\lambda^{-n(\nu)}\right] d\nu$ and rearrange this equation]

EXERCISE 14.30. * Prove Proposition 14.10.

EXERCISE 14.31. * Complete the proof of Proposition 14.11, in particular, prove that $z_0^* > z_1^*$ [Hint: use similar arguments to the first part of the proof.]

EXERCISE 14.32. * Consider a steady-state equilibrium in the model of Section 14.4. Suppose that $\kappa = 0$ and

$$G'(0) < \frac{1-\lambda}{\rho}$$

Let

$$z^* \equiv G'^{-1}\left(\frac{1-\lambda}{\rho}\right)$$

and suppose also that

$$G'(0) < \frac{z^*(1-\lambda)/\rho + G(z^*)}{\rho + z^*}.$$

(1) Show that in this case the steady state equilibrium has zero growth.

(2) Show that $\kappa > 0$ will lead to a positive growth rate. Interpret this result and contrast it to the negative effects of relaxing the protection of intellectual property rights in the baseline model of Schumpeterian growth.

Exercise 14.33. * Modify the model presented in Section 14.4 such that followers can now use the innovation of the technological leader and immediately leapfrog the leader, but in this case they have to pay a license fee of ζ to the leader.

(1) Characterize the BGP in this case
(2) Write the value functions.
(3) Explain why licensing can increase the growth rate of the economy in this case, and contrast this result with the one in Exercise 12.9, where licensing was never used in equilibrium. What is the source of the difference between the two sets of results?

Exercise 14.34. (1) What is the effect of competition on the rate of growth of the economy in a standard product variety model of endogenous growth? What about the quality-ladder model? Explain the intuition.

(2) Now consider the following one-period model. There are two Bertrand duopolists, producing a homogeneous good. At the beginning of each period, duopolist 1's marginal cost of production is determined as a draw from the uniform distribution $[0, \bar{c}_1]$ and the marginal cost of the second duopolist is determined as an independent draw from $[0, \bar{c}_2]$. Both cost realizations are observed and then prices are set. Demand is given by $Q = A - P$, with $A > 2 \max\{\bar{c}_1, \bar{c}_2\}$.
   (a) Characterize the equilibrium pricing strategies and calculate expected ex ante profits of the two duopolists.
   (b) Now imagine that both duopolists start with a cost distribution $[0, \bar{c}]$, and can undertake R&D at cost μ. If they do, with probability η, their cost distribution shifts to $[0, \bar{c} - \alpha]$ where $\alpha < \bar{c}$. Find the conditions under which one of the duopolists will invest in R&D and the conditions under which both will.
   (c) What happens when $\bar{c}$ declines? Interpreting the decline in $\bar{c}$ as increased competition, discuss the effect of increased competition on innovation incentives. Why is the answer different from that implied by the baseline endogenous technological change models of expanding varieties or Schumpeterian growth?

CHAPTER 15

Directed Technological Change

The previous two chapters introduced the basic models of endogenous technological change. These models provide us with a tractable framework for the analysis of aggregate technological change, but focus on a single type of technological change. Even when there are multiple types of machines, these all play the same role in increasing aggregate productivity. Consequently, technological change in these models is always "neutral" (that is, Hicks-neutral as defined in Chapter 2). There are two important respects in which these models are incomplete. First, technological change in practice is often not neutral: it benefits some factors of production and some agents in the economy more than others. Only in special cases, such as in economies with Cobb-Douglas aggregate production functions, these types of biases can be ignored. The study of why technological change is sometimes biased towards certain factors or sectors is both important for understanding the nature of endogenous technology and also because it clarifies the distributional effects of technological change, which determine which groups will embrace new technologies and which will oppose them. Second, limiting the analysis to only one type of technological change potentially obscures the different competing effects that determine the nature of technological change.

The purpose of this chapter is to extend the models of the last two chapters to consider *directed technological change*, which endogenizes the direction and bias of new technologies that are developed and adopted. Models of directed technological change not only generate new insights about the nature of endogenous technological progress, but also enable us to ask and answer new questions about recent and historical technological developments.

I start with a brief discussion of a range of economic problems in which considering the endogenous bias of technology is important and also present some of the general economic insights that will be important in models of directed technological change. The main results are presented in 15.3. The rest of the chapter generalizes the results and presents a few of their applications. Section 15.6 uses these models to return to the question raised in Chapter 2 concerning why technological change might take a purely labor-augmenting (Harrod-neutral) form. Section 15.8 presents an alternative approach to this question suggested by Jones (2005).

15.1. Importance of Biased Technological Change

To see the potential importance of the biased technological change, let us first review a number of examples.

1. Perhaps the most important example of biased technological change is the so-called *skill-biased technological change*, which plays an important role in the analysis of recent labor market developments and changes in the wage structure. Figure 15.1 plots a measure of the relative supply of skills (defined as the number of college equivalent workers divided by noncollege equivalents) and a measure of the return to skills, the college premium. It shows that over the past 60 years, the US relative supply of skills has increased rapidly, but there has been no tendency for the returns to college to fall in the face of this large increase in supply—on the contrary, there has been an increase in the college premium over this time period. The standard explanation for this pattern is that new technologies over the post-war period have been *skill-biased*. In fact, at some level this has to be so; if skilled and unskilled workers are imperfect substitutes, an increase in the relative supply of skills, without technological change, will necessarily reduce the skill premium.

The figure also shows that beginning in the late 1960s, the relative supply of skills increased much more rapidly than before, and the skill premium increased very rapidly beginning precisely in the late 1970s. The standard explanation for this increase is an acceleration in the skill bias of technological change that happens to be coincidental with the significant changes in the relative supply of skills.

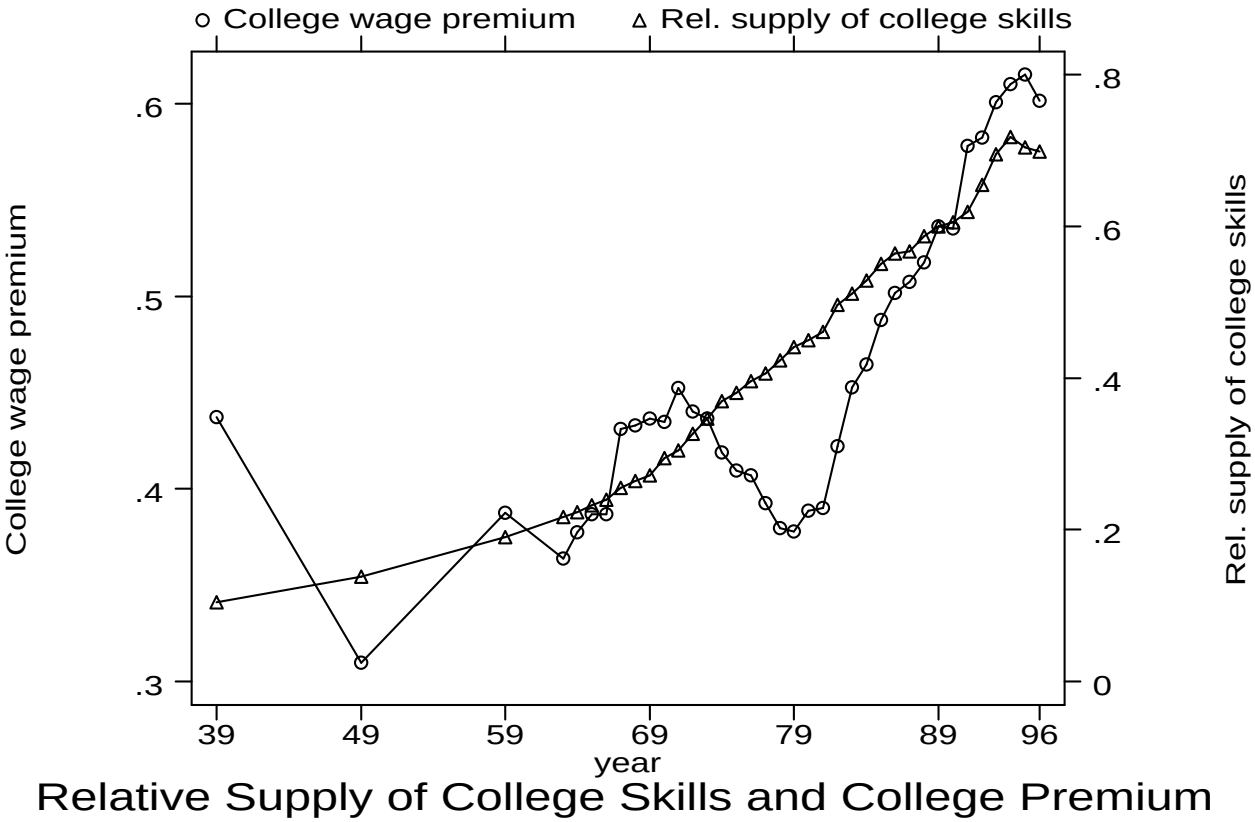

Figure 15.1. Relative supply of college graduates and the college premium in the U.S. labor market.

An obvious question is why technological changes have been skill-biased over the past 60 years or even 100 years? Relatedly, why does it appear that skill-biased technological change accelerated starting in the 1970s, precisely when the supply of skills increased rapidly? While some economists are happy to treat the bias of technological change as exogenous, this is not entirely satisfactory. We have seen that understanding the endogenous nature of technology is important for our study of cross-country income differences and the process of modern economic growth. It is unlikely that, while the amount of aggregate technological change is endogenous, the bias of technological change is entirely exogenous. It is therefore important to study the determinants of endogenous bias of technological change and ask why technological change has become more skill-biased in recent decades.

2. This conclusion is strengthened when we look at the historical process of technological change. In contrast to the developments during the recent decades, technological changes during the late 18th and early 19th centuries appear to have been *unskill-biased*. The artisan shop was replaced by the factory and later by interchangeable parts and the assembly line. Products previously manufactured by skilled artisans started to be produced in factories by workers with relatively few skills, and many previously complex tasks were simplified, reducing the demand for skilled workers. Mokyr (1990, p. 137) summarizes this process as follows:

> "First in firearms, then in clocks, pumps, locks, mechanical reapers, typewriters, sewing machines, and eventually in engines and bicycles, interchangeable parts technology proved superior and replaced the skilled artisans working with chisel and file."

Even though the types of skills valued in the labor market during the 19th century were different from those supplied by college graduates in today's labor markets, the juxtaposition of technological change biased towards college graduates in the recent past and biased against the most skilled workers of the time in the 19th century is both puzzling and intriguing. It raises the question: why was technological change, that has been generally skill-biased over the 20th century, biased towards unskilled workers in the 19th century?

3. As another example, consider the potential effect of labor market conditions on technological change. Beginning in the late 1960s and the early 1970s, both unemployment and the share of labor in national income increased rapidly in a number of continental European countries. During the 1980s, unemployment continued to increase, but the labor share started a steep decline, and in many countries it fell below its initial level. Blanchard (1997) interprets the first phase as the response of these economies to a wage-push by workers, and the second phase as a possible consequence of *capital-biased* technological changes. Is there a connection between capital-biased technological changes in European economies and the wage push preceding it?

4. The other obvious example of the importance of direct technological change is the common restriction to Harrod-neutral (purely labor-augmenting) technological progress in

growth models. Recall from Chapters 2 and 8 that if technological change is not labor-augmenting, equilibrium growth will not be balanced. But a range of evidence suggests that modern economic growth has been relatively balanced. Is there any reason to expect technological change to be endogenously labor-augmenting?

5. Finally, the past several decades have experienced a large increase in the volume of international trade and a rapid process of globalization. Do we expect globalization to affect the types of technologies that are being developed and used?

This chapter shows that a framework of directed technological change can provide potential answers to these questions. The main insight is to think of profit incentives as affecting not only the amount but also the direction of technological change. Before presenting detailed models, let us review the basic arguments, which are quite intuitive.

Imagine an economy which has two different factors of production, say L and H (corresponding, for example, to unskilled and skilled workers), and two different types of technologies that can complement (augment) either one or the other factor. We would expect that whenever the profitability of H-augmenting technologies is greater than the L-augmenting technologies, more of the former type will be developed by profit-maximizing (research) firms. But then, what determines the relative profitability of developing different technologies? The answer to this question summarizes most of the economics in the models of directed technological change. Two potentially counteracting effects shape the relative profitabilities of different types of technologies:

(1) *The price effect*: there will be stronger incentives to develop technologies when the goods produced by these technologies command higher prices.
(2) *The market size effect*: it is more profitable to develop technologies that have a larger market (for example, for the reasons discussed in Chapter 12).

An important result of the analysis in this chapter will be that this market size effect is powerful enough to outweigh the price effect. In fact, under fairly general conditions the following two results will hold:

- **Weak Equilibrium (Relative) Bias**: an increase in the relative supply of a factor always induces technological change that is biased in favor of this factor.
- **Strong Equilibrium (Relative) Bias**: if the elasticity of substitution between factors is sufficiently large, an increase in the relative supply of a factor induces sufficiently strong technological change biased towards itself that the endogenous-technology relative demand curve of the economy becomes *upward-sloping.*

To explain these concepts in a little more detail, suppose that the (inverse) relative demand curve takes the form $w_H/w_L = D(H/L, A)$, where w_H/w_L is the relative price of the H factor relative to the L factor, H/L is the relative supply of the H factor and A is a technology term. Technology A is *H-biased* if D is increasing in A, so that a higher A increases the relative demand for the H factor. The standard microeconomic theory implies

that D is *always* decreasing in H/L. Equilibrium bias concerns the behavior of A as H/L changes, so let us write this as $A(H/L)$. As a normalization, suppose that A is indeed H-biased. *Weak equilibrium bias* simply implies that $A(H/L)$ is increasing (nondecreasing) in H/L. *Strong equilibrium bias*, on the other hand, implies that $A(H/L)$ is sufficiently responsive to an increase in H/L that the total effect of the change in relative supply H/L is to increase w_H/w_L. In other words, let the endogenous-technology relative demand curve be $w_H/w_L = D(H/L, A(H/L)) \equiv \tilde{D}(H/L)$. Then, strong equilibrium bias corresponds to this endogenous-technology relative demand curve, $\tilde{D}$, being increasing.

At first, both the weak and the strong equilibrium bias results appear surprising. However, they become quite intuitive once the logic of directed technological change is understood. Moreover, they have a range of important implications. In particular, subsection 15.3.3 below shows how the weak and the strong relative bias results provide us with potential answers to the questions posed at the beginning of this section.

15.2. Basics and Definitions

Before studying directed technological change, it is useful to clarify the difference between factor-augmenting and factor-biased technological changes, which are sometimes confused in the literature. Suppose that the production side of the economy can be represented by an aggregate production function,

$$Y(t) = F(L(t), H(t), A(t)),$$

where $L(t)$ is labor, and $H(t)$ denotes another factor of production, which could be skilled labor, capital, land or some intermediate goods, and $A(t)$ represents technology. Without loss of generality imagine that $\partial F/\partial A > 0$, so a greater level of A corresponds to "better technology". Recall that technological change is *L-augmenting* if

$$\frac{\partial F(L,H,A)}{\partial A} \equiv \frac{L}{A}\frac{\partial F(L,H,A)}{\partial L}.$$

This is clearly equivalent to the production function taking the more special form, $F(AL, H)$. In the case where L corresponds to labor and H to capital, this is also equivalent to Harrod-neutral technological change. Conversely, H-augmenting technological change is defined similarly, and corresponds to the production function taking the special form $F(L, AH)$.

Though often equated with factor-augmenting changes, the concept of factor-biased technological change is *very* different. We say that technological change change is *L-biased*, if it increases the *relative* marginal product of factor L compared to factor H. Mathematically, this corresponds to

$$\frac{\partial \frac{\partial F(L,H,A)/\partial L}{\partial F(L,H,A)/\partial H}}{\partial A} \geq 0.$$

Put differently, biased technological change shifts out the relative demand curve for a factor, so that its relative marginal product (relative price) increases at given factor proportions (given relative quantity of factors). Conversely, technological change is H-biased if this inequality holds in reverse. Figure 15.2 plots the effect of an H-biased (skill-biased)

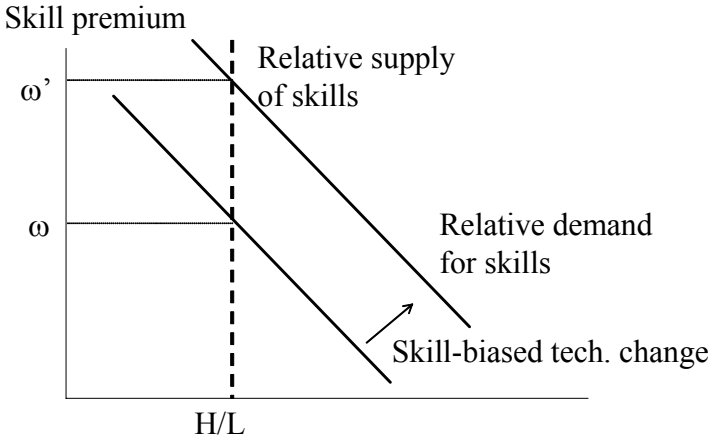

FIGURE 15.2. The effect of H-biased technological change on relative demand and relative factor prices.

technological change on the relative demand for factor H and on its relative price, the skill premium.

These concepts can be further clarified using the CES production function (recall Example 2.3 in Chapter 2), which can be written as

$$Y(t)=\left[\gamma_L\left(A_L(t)L(t)\right)^{\frac{\sigma-1}{\sigma}}+\gamma_H\left(A_H(t)H(t)\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}},$$

where $A_L(t)$ and $A_H(t)$ are two separate technology terms, the γ_is determine the importance of the two factors in the production function, and $\gamma_L+\gamma_H=1$. Finally, $\sigma\in(0,\infty)$ is the elasticity of substitution between the two factors. Recalled that when $\sigma=\infty$, the two factors are perfect substitutes, and the production function is linear; when $\sigma=1$, the production function is Cobb-Douglas; and when $\sigma=0$, there is no substitution between the two factors, and the production function is Leontieff. Let us refer to $\sigma>1$ as "gross substitutes," and and to $\sigma<1$ as "gross complements". While there are multiple definitions of complementarity in the microeconomics literature, this terminology is useful to distinguish the two cases in which $\sigma<1$ and $\sigma>1$, which will have very different implications in the current context.

Clearly, by construction, $A_L(t)$ is L-augmenting, while $A_H(t)$ is H-augmenting. Interestingly, whether technological change that is L-augmenting (or H-augmenting) is L-biased or H-biased depends on the elasticity of substitution, σ. Let us first calculate the relative marginal product of the two factors (see Exercise 15.1):

$$\frac{MP_H}{MP_L}=\gamma\left(\frac{A_H(t)}{A_L(t)}\right)^{\frac{\sigma-1}{\sigma}}\left(\frac{H(t)}{L(t)}\right)^{-\frac{1}{\sigma}}, \tag{15.1}$$

where $\gamma\equiv\gamma_H/\gamma_L$. The relative marginal product of H is decreasing in its relative abundance, $H(t)/L(t)$. This is simply the consequence of the usual *substitution effect*, leading to a

negative relationship between relative supplies and relative marginal products (or prices) and thus to a downward-sloping relative demand curve (see Figure 15.3). The effect of $A_H(t)$ on this relative marginal product depends on σ, however. If $\sigma > 1$, an increase in $A_H(t)$ (relative to $A_L(t)$) increases the relative marginal product of H. In contrast, when $\sigma < 1$, an increase in $A_H(t)$ reduces the relative marginal product of H. Therefore, when the two factors are gross substitutes, H-augmenting technological change is also H-biased. In contrast, when the two factors are gross complements, the relationship is *reversed*, and H-augmenting technological change is now L-biased. Naturally, when $\sigma = 1$, the production function is Cobb-Douglas, and neither a change in $A_H(t)$ nor in $A_L(t)$ is biased towards any of the factors. Note also for future reference that by virtue of the fact that σ is the elasticity of substitution between the two factors,

$$\sigma = -\left(\frac{d\log\left(MP_H/MP_L\right)}{d\log\left(H/L\right)}\right)^{-1}$$

The intuition for why, when $\sigma < 1$, H-augmenting technological change is L-biased is simple but important: with gross complementarity ($\sigma < 1$), an increase in the productivity of H increases the demand for labor, L, by more than the demand for H. As a result, the marginal product of labor increases by more than the marginal product of H. This can be seen most clearly in the extreme case where $\sigma \to 0$, so that the two factors become Leontieff. In this case, starting from a situation in which $\gamma_L A_L(t) L(t) = \gamma_H A_H(t) H(t)$, a small increase in $A_H(t)$ will create an "excess of the services" of the H factor (and thus "excess demand" for L), and the price of factor H will fall to 0.

I have so far defined the meaning of H-biased and L-biased technological changes. It is also useful to define two concepts that will play a major role in the remainder of this chapter. There is *weak equilibrium bias* of technology if an increase in the relative supply of H, H/L, induces technological change biased towards H. Mathematically, this is equivalent to:

$$\frac{\partial MP_H/MP_L}{\partial A_H/A_L}\frac{dA_H/A_L}{dH/L} \geq 0.$$

From (15.1), it is clear that this condition will hold if

$$\frac{d\left(A_H(t)/A_L(t)\right)^{\frac{\sigma-1}{\sigma}}}{dH/L} \geq 0,$$

so that in response to the change in relative supplies, $A_H(t)/A_L(t)$ changes in a direction that is biased towards the factor that has become more abundant.

On the other hand, there is *strong equilibrium bias* if an increase in H/L induces a sufficiently large change in the bias of technology so that, following the change in factor supplies, the marginal product of H *relative* to that of L increases. Mathematically, this is equivalent to

$$\frac{dMP_H/MP_L}{dH/L} > 0,$$

where I now use a strict inequality to distinguish strong equilibrium bias from the case in which relative marginal products are independent of relative supplies (e.g., because factors

are perfect substitutes). These equations make it clear that the major difference between weak and strong equilibrium bias is whether the relative marginal product of the two factors are evaluated at the initial relative supplies (in the case of weak bias) or at the new relative supplies (in the case of strong bias). Consequently, strong equilibrium bias is a much more demanding concept than weak equilibrium bias.

15.3. Baseline Model of Directed Technological Change

In this section, I present the baseline model of directed technological change, which uses the expanding varieties model of endogenous technological change and the lab-equipment specification of the innovation possibilities frontier. The former choice is motivated by the fact that the expanding varieties model is somewhat simpler to work with than the model of Schumpeterian growth introduced in the previous chapter. The lab-equipment specification, on the other hand, highlights that none of the results here depend on technological externalities. Section 15.4 will consider a model of directed technological with change knowledge spillovers. Exercise 15.19 shows that all of the results presented here generalize to a model of Schumpeterian growth, thus the assumption of expanding varieties is only adopted for convenience.

The baseline economy has a constant supply of two factors, L and H, and admits a representative household with the standard CRRA preferences given by

$$\int_0^{\infty} \exp\left(-\rho t\right) \frac{C\left(t\right)^{1-\theta}-1}{1-\theta} dt, \tag{15.2}$$

where, as usual, $\rho > 0$. The supply side is represented by an aggregate production function combining the outputs of two intermediate sectors with a constant elasticity of substitution:

$$Y\left(t\right) = \left[\gamma_L Y_L\left(t\right)^{\frac{\varepsilon-1}{\varepsilon}} + \gamma_H Y_H\left(t\right)^{\frac{\varepsilon-1}{\varepsilon}}\right]^{\frac{\varepsilon}{\varepsilon-1}}, \tag{15.3}$$

where $Y_L\left(t\right)$ and $Y_H\left(t\right)$ denote the outputs of two intermediate goods. As the indices indicate, the first is L-intensive, while the second is H-intensive. The parameter $\varepsilon \in [0, \infty)$ is the elasticity of substitution between these two intermediate goods. The assumption that (15.3) features a constant elasticity of substitution simplifies the analysis but is not crucial for the results. How relaxing this assumption affects the results is discussed at the end of this chapter.

The resource constraint of the economy at time t takes the form

$$C\left(t\right) + X\left(t\right) + Z\left(t\right) \leq Y\left(t\right), \tag{15.4}$$

where, as before, $X\left(t\right)$ denotes total spending on machines and $Z\left(t\right)$ is aggregate R&D spending.

The two intermediate goods are produced competitively with the following production functions:

$$Y_L\left(t\right) = \frac{1}{1-\beta}\left(\int_0^{N_L(t)} x_L\left(\nu, t\right)^{1-\beta} d\nu\right) L^{\beta} \tag{15.5}$$

and

$$Y_H(t) = \frac{1}{1-\beta}\left(\int_0^{N_H(t)} x_H(\nu,t)^{1-\beta}\, d\nu\right) H^{\beta}, \tag{15.6}$$

where $x_L(\nu,t)$ and $x_H(\nu,t)$ denote the quantities of the different types of machines (used in the production of one or the other intermediate good) and $\beta \in (0,1)$.[1] These machines are again assumed to depreciate after use. The parallel between these production functions and the aggregate production function of the economy in the baseline expanding product varieties model of Chapter 13 is obvious. There are two important differences, however. First, these are now production functions for intermediate goods rather than the final good. Second, the two production functions (15.5) and (15.6) use different types of machines. The range of machines complementing labor, L, is $[0, N_L(t)]$, while the range of machines complementing factor H is $[0, N_H(t)]$.

Again as in Chapter 13, all machines in both sectors are supplied by monopolists that have a fully-enforced perpetual patent on the machines. by denote the prices charged by these monopolists at time t by $p_L^x(\nu,t)$ for $\nu \in [0, N_L(t)]$ and $p_H^x(\nu,t)$ for $\nu \in [0, N_H(t)]$. Once invented, each machine can be produced at the fixed marginal cost ψ in terms of the final good, which we again normalize to $\psi \equiv 1-\beta$. This implies that total resources devoted to machine production at time t are

$$X(t) = (1-\beta)\left(\int_0^{N_L(t)} x_L(\nu,t)\, d\nu + \int_0^{N_H(t)} x_H(\nu,t)\, d\nu\right).$$

The *innovation possibilities frontier*, which determines how new machine varieties are created, is assumed to take a form similar to the lab-equipment specification in Chapter 13:

$$\dot{N}_L(t) = \eta_L Z_L(t) \text{ and } \dot{N}_H(t) = \eta_H Z_H(t), \tag{15.7}$$

where $Z_L(t)$ is R&D expenditure *directed* at discovering new labor-augmenting machines at time t, while $Z_H(t)$ is R&D expenditure directed at discovering H-augmenting machines. Total R&D spending is the sum of these two, that is,

$$Z(t) = Z_L(t) + Z_H(t).$$

The value of a monopolist that discovers one of these machines is again given by the standard formula for the present discounted value of profits:

$$V_f(\nu,t) = \int_t^{\infty} \exp\left[-\int_t^s r\left(s'\right) ds'\right] \pi_f(\nu,s) ds, \tag{15.8}$$

where $\pi_f(\nu,t) \equiv p_f^x(\nu,t) x_f(\nu,t) - \psi x_f(\nu,t)$ again denotes instantaneous profits for $f = L$ or H, and $r(t)$ is the market interest rate at time t. Once again, it is sometimes more convenient to work with the Hamilton-Jacobi-Bellman version of this value function, which takes the form:

$$r(t) V_f(\nu,t) - \dot{V}_f(\nu,t) = \pi_f(\nu,t). \tag{15.9}$$

[1]Note that the range of machines used in the two sectors are different (there are two disjoint sets of machines); I use the index ν to denote either set of machines for notational simplicity.

Throughout, the price of the final good is normalized to 1 at every instant, which is equivalent to setting the ideal price index of the two intermediates equal to one, that is,

$$\left[\gamma_L^{\varepsilon} p_L(t)^{1-\varepsilon} + \gamma_H^{\varepsilon} p_H(t)^{1-\varepsilon}\right]^{\frac{1}{1-\varepsilon}} = 1 \text{ for all } t, \tag{15.10}$$

where $p_L(t)$ is the price index of Y_L at time t and $p_H(t)$ is the price of Y_H. Intertemporal prices are therefore given by the interest rates, $[r(t)]_{t=0}^{\infty}$. Finally, factor prices for L and H are denoted by $w_L(t)$ and $w_H(t)$.

15.3.1. Characterization of Equilibrium. An allocation in this economy is defined by the following objects: time paths of consumption levels, aggregate spending on machines, and aggregate R&D expenditure $[C(t), X(t), Z(t)]_{t=0}^{\infty}$, time paths of available machine types, $[N_L(t), N_H(t)]_{t=0}^{\infty}$, time paths of prices and quantities of each machine and the net present discounted value of profits from that machine, $[p_L^x(\nu,t), x_L(\nu,t), V_L(\nu,t)]^{\infty}_{\substack{t=0,\\ \nu\in[0,N_L(t)]}}$ and $[p_H^x(\nu,t), x_H(\nu,t), V_H(\nu,t)]^{\infty}_{\substack{t=0,\\ \nu\in[0,N_H(t)]}}$, and time paths of factor prices, $[r(t), w_L(t), w_H(t)]_{t=0}^{\infty}$.

An equilibrium is an allocation in which all existing research firms choose $\left[p_f^x(\nu,t), x_f(\nu,t)\right]^{\infty}_{\substack{t=0,\\ \nu\in[0,N_f(t)]}}$ for $f = L, H$ to maximize profits, the evolution of $[N_L(t), N_H(t)]_{t=0}^{\infty}$ is determined by free entry, the time paths of factor prices, $[r(t), w_L(t), w_H(t)]_{t=0}^{\infty}$, are consistent with market clearing, and the time paths of $[C(t), X(t), Z(t)]_{t=0}^{\infty}$ are consistent with consumer optimization.

To characterize the (unique) equilibrium, let us first consider the maximization problem of producers in the two sectors. Since machines depreciate fully after use, these maximization problems are static and can be written as

$$\max_{L,[x_L(\nu,t)]_{\nu\in[0,N_L(t)]}} p_L(t) Y_L(t) - w_L(t) L - \int_0^{N_L(t)} p_L^x(\nu,t)\, x_L(\nu,t)\, d\nu, \tag{15.11}$$

and

$$\max_{H,[x_H(\nu,t)]_{\nu\in[0,N_H(t)]}} p_H(t) Y_H(t) - w_H(t) H - \int_0^{N_H(t)} p_H^x(\nu,t)\, x_H(\nu,t)\, d\nu. \tag{15.12}$$

The main difference from the maximization problem facing final good producers in Chapter 13 is the presence of prices $p_L(t)$ and $p_H(t)$, which reflect the fact that these sectors produce intermediate goods, whereas factor and machine prices are expressed in terms of the numeraire, the final good.

These two maximization problems immediately imply the following demand for machines in the two sectors:

$$x_L(\nu,t) = \left[\frac{p_L(t)}{p_L^x(\nu,t)}\right]^{1/\beta} L \quad \text{for all } \nu \in [0, N_L(t)] \text{ and all } t, \tag{15.13}$$

and

$$x_H(\nu,t) = \left[\frac{p_H(t)}{p_H^x(\nu,t)}\right]^{1/\beta} H \quad \text{for all } \nu \in [0, N_H(t)] \text{ and all } t. \tag{15.14}$$

Similar to the demands for machines in Chapter 13, these are isoelastic, so the maximization of the net present discounted value of profits implies that each monopolist should set a constant markup over marginal cost and thus a price of

$$p_L^x(\nu,t) = p_H^x(\nu,t) = 1 \text{ for all } \nu \text{ and } t.$$

Substituting these prices into (15.13) and (15.14),

$$x_L(\nu,t) = p_L(t)^{1/\beta} L \quad \text{for all } \nu \text{ and all } t,$$

and

$$x_H(\nu,t) = p_H(t)^{1/\beta} H \quad \text{for all } \nu \text{ and all } t.$$

Since these quantities do not depend on the identity of the machine, only on the sector that is being served, profits are also independent of the machine type. In particular,

$$\pi_L(t) = \beta p_L(t)^{1/\beta} L \text{ and } \pi_H(t) = \beta p_H(t)^{1/\beta} H. \tag{15.15}$$

This implies that the net present discounted values of monopolists only depend on which sector they are supplying and can be denoted by $V_L(t)$ and $V_H(t)$.

Next, combining these with (15.5) and (15.6) yields the *derived* production functions for the two intermediate goods:

$$Y_L(t) = \frac{1}{1-\beta} p_L(t)^{\frac{1-\beta}{\beta}} N_L(t) L \tag{15.16}$$

and

$$Y_H(t) = \frac{1}{1-\beta} p_H(t)^{\frac{1-\beta}{\beta}} N_H(t) H. \tag{15.17}$$

These derived production functions are similar to (13.12) in Chapter 13, except for the presence of the price terms.

Finally, the prices of the two intermediate goods are derived from the marginal product conditions of the final good technology, (15.3), which imply

$$\begin{aligned} p(t) &\equiv \frac{p_H(t)}{p_L(t)} = \gamma \left(\frac{Y_H(t)}{Y_L(t)} \right)^{-\frac{1}{\varepsilon}} \\ &= \gamma \left(p(t)^{\frac{1-\beta}{\beta}} \frac{N_H(t) H}{N_L(t) L} \right)^{-\frac{1}{\varepsilon}} \\ &= \gamma^{\frac{\varepsilon\beta}{\sigma}} \left(\frac{N_H(t) H}{N_L(t) L} \right)^{-\frac{\beta}{\sigma}}, \end{aligned} \tag{15.18}$$

where again $\gamma \equiv \gamma_H/\gamma_L$ and

$$\begin{aligned} \sigma &\equiv \varepsilon - (\varepsilon - 1)(1-\beta) \\ &= 1 + (\varepsilon - 1)\beta. \end{aligned}$$

is the (derived) elasticity of substitution between the two factors. The first line of this expression simply defines $p(t)$ as the relative price between the two intermediate goods and uses the fact that the ratio of the marginal productivities of the two intermediate goods must

be equal to this relative price. The second line substitutes from (15.16) and (15.17) above. Using the latter equation, the relative factor prices in this economy are obtained as

$$\begin{aligned}\omega(t) &\equiv \frac{w_H(t)}{w_L(t)} \\ &= p(t)^{1/\beta}\frac{N_H(t)}{N_L(t)} \\ &= \gamma^{\frac{\varepsilon}{\sigma}}\left(\frac{N_H(t)}{N_L(t)}\right)^{\frac{\sigma-1}{\sigma}}\left(\frac{H}{L}\right)^{-\frac{1}{\sigma}}. \qquad (15.19)\end{aligned}$$

The first line defines $\omega(t)$ as the relative wage of factor H compared to factor L. The second line uses the definition of marginal product combined with (15.16) and (15.17), and the third line, (15.19), uses (15.18). I refer to σ as the (derived) elasticity of substitution between the two factors, since it is exactly equal to

$$\sigma = -\left(\frac{d\log\omega(t)}{d\log(H/L)}\right)^{-1}.$$

To complete the description of equilibrium in the technology side, we need to impose the following free-entry conditions:

$$\eta_L V_L(t) \leq 1 \text{ and } \eta_L V_L(t) = 1 \text{ if } Z_L(t) > 0. \qquad (15.20)$$

and

$$\eta_H V_H(t) \leq 1 \text{ and } \eta_H V_H(t) = 1 \text{ if } Z_H(t) > 0. \qquad (15.21)$$

Finally, the consumer side is characterized by the same necessary conditions as usual:

$$\frac{\dot{C}(t)}{C(t)} = \frac{1}{\theta}(r(t) - \rho), \qquad (15.22)$$

and

$$\lim_{t\to\infty}\left[\exp\left(-\int_0^t r(s)\,ds\right)(N_L(t)V_L(t) + N_H(t)V_H(t))\right] = 0, \qquad (15.23)$$

which uses the fact that $N_L(t)V_L(t) + N_H(t)V_H(t)$ is the total value of corporate assets in this economy.

Let us next define the BGP equilibrium as an equilibrium path where consumption grows at the constant rate, g^*, and the relative price $p(t)$ is constant. From (15.10) this implies that $p_L(t)$ and $p_H(t)$ are also constant.

Let V_L and V_H be the BGP net present discounted values of new innovations in the two sectors. Then, (15.9) implies that

$$V_L = \frac{\beta p_L^{1/\beta} L}{r^*} \text{ and } V_H = \frac{\beta p_H^{1/\beta} H}{r^*}, \qquad (15.24)$$

where r^* is the BGP interest rate, while p_L and p_H are the BGP prices of the two intermediate goods. The comparison of these two values is of crucial importance. As discussed intuitively

above, the greater is V_H relative to V_L, the greater are the incentives to develop H-augmenting machines, N_H, rather than N_L. Taking the ratio of these two expressions,

$$\frac{V_H}{V_L} = \left(\frac{p_H}{p_L}\right)^{\frac{1}{\beta}} \frac{H}{L}.$$

This expression highlights the two effects on the direction of technological change discussed in Section 15.1.

(1) The price effect manifests itself because V_H/V_L is increasing in p_H/p_L. The greater is this relative price, the greater are the incentives to invent technologies complementing the H factor. Since goods produced by relatively scarce factors will be relatively more expensive, the price effect tends to favor technologies complementing scarce factors.

(2) The market size effect is a consequence of the fact that V_H/V_L is increasing in H/L. The market for a technology is the workers (or the other factors) that will be using and working with this technology. Consequently, an increase in the supply of a factor translates into a greater market for the technology complementing that factor. The market size effect encourages innovation for the more abundant factor.

The above discussion is incomplete, however, since prices are endogenous. Combining (15.24) together with (15.18), relative prices can be eliminated and the relative profitability of the technologies becomes

$$\frac{V_H}{V_L} = \gamma^{\frac{\varepsilon}{\sigma}} \left(\frac{N_H}{N_L}\right)^{-\frac{1}{\sigma}} \left(\frac{H}{L}\right)^{\frac{\sigma-1}{\sigma}}. \tag{15.25}$$

Note for future reference that an increase in the relative factor supply, H/L, will increase V_H/V_L as long as $\sigma > 1$ and it will reduce it if $\sigma < 1$. This shows that the elasticity of substitution between the factors, σ, regulates whether the price effect dominates the market size effect. Since σ is not a primitive, but a derived parameter, we would like to know when it is greater than 1. It is straightforward to check that

$$\sigma \gtreqless 1 \iff \varepsilon \gtreqless 1.$$

So the two factors will be gross substitutes when the two intermediate goods are gross substitutes in the production of the final good.

Next, using the two free-entry conditions (15.20) and (15.21), and assuming that both of them hold as equalities, we obtain the following BGP *technology market clearing* condition:

$$\eta_L V_L = \eta_H V_H. \tag{15.26}$$

Combining this with (15.25), we obtain the following BGP ratio of relative technologies

$$\left(\frac{N_H}{N_L}\right)^* = \eta^\sigma \gamma^\varepsilon \left(\frac{H}{L}\right)^{\sigma-1}, \tag{15.27}$$

where $\eta \equiv \eta_H/\eta_L$ and the *'s denote that this expression refers to the BGP value. The notable feature here is that relative productivities are determined by the innovation possibilities frontier and the relative supply of the two factors. In this sense, this model totally endogenizes

technology. Equation (15.27) contains most of the economics of directed technology. However, before discussing this, it is useful to characterize the BGP growth rate of the economy. This is done in the next proposition:

PROPOSITION 15.1. *Consider the directed technological change model described above. Suppose that*

$$\beta\left[\gamma_H^{\varepsilon}\left(\eta_H H\right)^{\sigma-1}+\gamma_L^{\varepsilon}\left(\eta_L L\right)^{\sigma-1}\right]^{\frac{1}{\sigma-1}} > \rho \text{ and} \tag{15.28}$$

$$(1-\theta)\,\beta\left[\gamma_H^{\varepsilon}\left(\eta_H H\right)^{\sigma-1}+\gamma_L^{\varepsilon}\left(\eta_L L\right)^{\sigma-1}\right]^{\frac{1}{\sigma-1}} < \rho.$$

Then, there exists a unique BGP equilibrium in which the relative technologies are given by (15.27), and consumption and output grow at the rate

$$g^*=\frac{1}{\theta}\left(\beta\left[\gamma_H^{\varepsilon}\left(\eta_H H\right)^{\sigma-1}+\gamma_L^{\varepsilon}\left(\eta_L L\right)^{\sigma-1}\right]^{\frac{1}{\sigma-1}}-\rho\right). \tag{15.29}$$

PROOF. The derivation of (15.29) is provided by the argument preceding the proposition. Exercise 15.2 asks you to check that condition (15.28) ensures that free-entry conditions (15.20) and (15.21) must hold, to verify that this is the unique relative equilibrium technology, to calculate the BGP equilibrium growth rate and to verify that the transversality condition is satisfied. □

It can also be verified that there are simple transitional dynamics in this economy whereby starting with technology levels $N_H(0)$ and $N_L(0)$, there always exists a unique equilibrium path and it involves the economy monotonically converging to the BGP equilibrium of Proposition 15.1. This is stated in the next proposition:

PROPOSITION 15.2. *Consider the directed technological change model described above. Starting with any $N_H(0)>0$ and $N_L(0)>0$, there exists a unique equilibrium path. If $N_H(0)/N_L(0)<(N_H/N_L)^*$ as given by (15.27), then $Z_H(t)>0$ and $Z_L(t)=0$ until $N_H(t)/N_L(t)=(N_H/N_L)^*$. If $N_H(0)/N_L(0)<(N_H/N_L)^*$, then $Z_H(t)=0$ and $Z_L(t)>0$ until $N_H(t)/N_L(t)=(N_H/N_L)^*$.*

PROOF. See Exercise 15.3. □

More interesting than the aggregate growth rate and the transitional dynamics behavior of the economy are the results concerning the direction of technological change and its effects on relative factor prices. These are studied in the next subsection.

15.3.2. Directed Technological Change and Factor Prices. Let us start by studying (15.27). This equation implies that, in BGP, there is a positive relationship between the relative supply of the H factor, H/L, and the relative factor-augmenting technologies, N_H^*/N_L^* only when $\sigma>1$. In contrast, if the derived elasticity of substitution, σ, is less than 1, the relationship is reversed. This might suggest that, depending on the elasticity of substitution between factors (or between the intermediate goods), changes in factor supplies may induce technological changes that are biased in favor or against the factor that is becoming more

abundant. However, this conclusion is not correct. Recall from Section 15.2 that N_H^*/N_L^* refers to the ratio of factor-augmenting technologies, or to the ratio of *physical* productivities. What matters for the bias of technology is *the value of marginal product* of factors, which is affected by changes in relative prices. We have already seen that the relationship between factor-augmenting technologies and factor-biased technologies is reversed precisely when σ is less than 1. Thus, when $\sigma > 1$, an increase in N_H^*/N_L^* is relatively biased towards H, while when $\sigma < 1$, it is a decrease in N_H^*/N_L^* that is relatively biased towards H.

This immediately establishes the following *weak equilibrium bias result*:

PROPOSITION 15.3. *Consider the directed technological change model described above. There is always weak equilibrium (relative) bias in the sense that an increase in H/L always induces relatively H-biased technological change.*

Recall that weak bias was defined in Section 15.2 with a weak inequality, so that the proposition is also correct when $\sigma = 1$, even though in this case it can be verified easily from (15.27) that N_H^*/N_L^* does not depend on H/L.

Proposition 15.3 is the basis of the discussion about induced biased technological change in Section 15.1, and already gives us a range of insights about how changes in the relative supplies of skilled workers may be at the root of the skill-biased technological change. These implications are further discussed in the next subsection.

The results of this proposition reflect the strength of the market size effect discussed above. Recall that the price effect creates a force favoring factors that become relatively scarce. In contrast, the market size effect, which is related to the non-rivalry of ideas discussed in Chapter 12, suggests that technologies should change in a way that favors factors that are becoming relatively abundant. Proposition 15.3 shows that the market size effect always dominates the price effect.

Proposition 15.3 is only informative about the direction of the induced technological change, but does not specify whether this induced effect will be strong enough to make the endogenous-technology relative demand curve for factors upward-sloping. Recall that in basic producer theory, all demand curves, and thus relative demand curves, are downward-sloping as well. However, as hinted in Section 15.1, directed technological change can lead to the seemingly paradoxical result that relative demand curves can be upward-sloping once the endogeneity of technology is taken into account. To obtain this result, let us substitute for $(N_H/N_L)^*$ from (15.27) into the expression for the relative wage given technologies, (15.19), and obtain the following BGP relative factor price ratio (see Exercise 15.4):

$$\omega^* \equiv \left(\frac{w_H}{w_L}\right)^* = \eta^{\sigma-1}\gamma^{\varepsilon}\left(\frac{H}{L}\right)^{\sigma-2}. \tag{15.30}$$

Inspection of this equation immediately establishes conditions for *strong equilibrium (relative) bias.*

PROPOSITION 15.4. *Consider the directed technological change model described above. Then, if $\sigma > 2$, there is strong equilibrium (relative) bias in the sense that an increase in*

H/L raises the relative marginal product and the relative wage of the factor H compared to factor L.

Figure 15.3 illustrates the results of Propositions 15.3 and 15.4, referring to H as skilled labor and L as unskilled labor as in the first application discussed in Section 15.1.

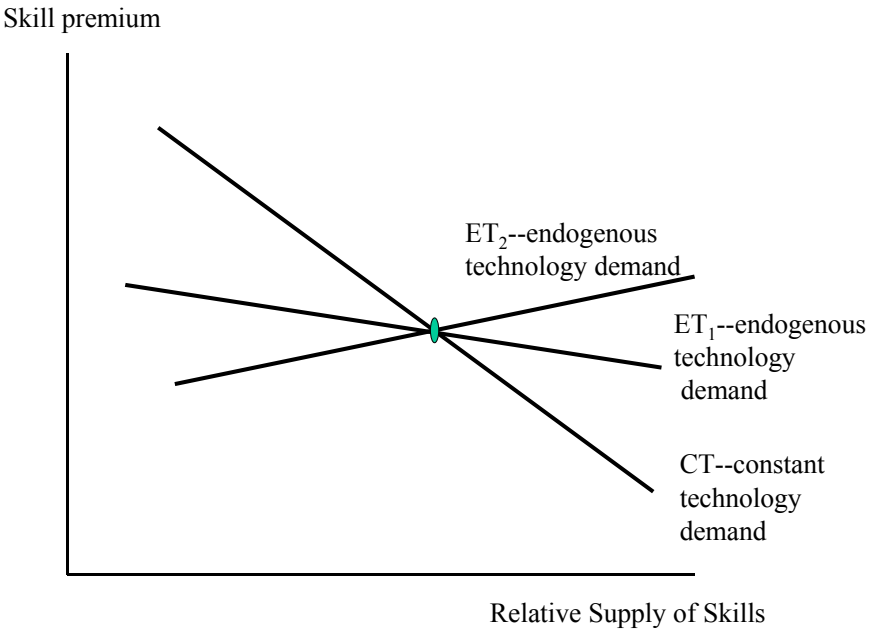

FIGURE 15.3. The relationship between the relative supply of skills and the skill premium in the model of directed technical change.

The curve marked with CT corresponds to the constant-technology relative demand from eq. (15.19). It is always downward-sloping because it holds the relative technologies, N_H/N_L, constant, and thus only features the usual substitution effect. The fact that this curve is downward-sloping follows from basic producer theory. The curve marked as ET_1 applies when technology is endogenous, but the condition in Proposition 15.4, that $\sigma > 2$, is not satisfied. Proposition 15.3 states that even in this case an increase in H/L will induce skill-biased (H-biased) technological change. This implies that when H/L is higher than its initial level, the induced-technology effect will shift the constant-technology demand curve CT to the right (i.e., as technology changes, it will be another CT curve, above the original one, that will apply). When it is below, the same effect will shift CT to the left. Consequently, the locus of points that the endogenous-technology demand, ET_1, is shallower than CT. (This can be also verified comparing the relative demand functions with constant and endogenous technology, (15.19) and (15.30), and noting that $\sigma - 2$ is never less than $-1/\sigma$). There is an obvious analogy between this result and Samuelson's *LeChatelier principle*, which states that long-run demand curves, which apply when all factors can adjust, must be more elastic than the short-run demand curves which hold some factors constant. We can think of the

endogenous-technology demand curve as adjusting the "factors of production" corresponding to technology. However, the analogy is imperfect because the effects here are caused by general equilibrium changes, while the LeChatelier principle and the basic producer theory focus on partial equilibrium effects. In fact, in basic producer theory, with or without the LeChatelier effects, all demand curves *must be downward-sloping*, whereas here ET_2, which applies when the conditions of Proposition 15.4 hold, is upward-sloping; higher levels of relative supply of skills correspond to higher skill premia.

A complementary intuition for this result can be obtained by going back to the importance of non-rivalry of ideas as discussed in Chapter 12. Here, as in the basic endogenous technology models of the last two chapters, the non-rivalry of ideas leads to an aggregate production function that exhibits increasing returns to scale (in all factors including technologies). It is the increasing returns to scale in the production possibilities set of the economy that leads to potentially upward-sloping relative demand curves. Put differently, the market size effect, which results from the non-rivalry of ideas and is at the root of aggregate increasing returns, can create sufficiently strong induced technological change to increase the relative marginal product and the relative price of the factor that has become more abundant.

15.3.3. Implications. The results of Propositions 15.3 and 15.4 are not only of theoretical interest, but also shed light on a range of important empirical patterns. As already discussed above, one of the most interesting applications is to changes in the skill premium. For this application, suppose that H stands for college-educated workers. In the US labor market, the skill premium has shown no tendency to decline despite a very large increase in the supply of college educated workers. On the contrary, following a brief period of decline during the 1970s in the face of the very large increase in the supply of college-educated workers, the skill (college) premium has increased very sharply throughout the 1980s and the 1990s, to reach a level not experienced in the postwar era. Figure 15.1 above showed these general patterns.

In the labor economics and parts of the macroeconomics literature, the most popular explanation for these patterns is skill-biased technological change. For example, computers or new IT technologies are argued to favor skilled workers relative to unskilled workers. But why should the economy adopt and develop more skill-biased technologies throughout the past 30 years, or more generally throughout the entire 20th century? This question becomes more relevant once we remember that during the 19th century many of the technologies that were fueling economic growth, such as the factory system and the major spinning and weaving innovations, were unskill-biased rather than skill-biased.

Thus, in summary, the following stylized facts are relevant:

- Secular skill-biased technological change increasing the demand for skills throughout the 20th century.
- Possible acceleration in skill-biased technological change over the past 25 years.

- A range of important technologies biased against skill workers during the 19th century.

Propositions 15.3 and 15.4 provide us with a framework for thinking about these issues. Proposition 15.3 implies the increase in the relative supply of skills will always induce skill-biased technological change. In addition, Proposition 15.4 implies that, when $\sigma > 2$, the long-run (endogenous-technology) relationship between the relative supply of skills and the skill premium is positive. More specifically:

1. According to Propositions 15.3 and 15.4, the increase in the number of skilled workers that has taken place throughout the 20th century should cause steady skill-biased technological change. Therefore, models of directed technological change offer a natural explanation for the secular skill-biased technological developments of the past century.

2. The more rapid increase in the number of skilled workers over the past 25 years, shown in Figure 15.1, should also induce an acceleration in skill-biased technological change. How this class of models might account for the dynamics of factor prices in the face of endogenously changing technologies is discussed below.

3. Can the framework also explain the prevalence of skill-replacing/labor-biased technological change in the late 18th and 19th centuries? While we know less about both changes in relative supplies and technological developments during these historical periods, available evidence suggests that there were large increases in the number of unskilled workers available to be employed in the factories during this time periods. Bairoch (1988, p. 245), for example, describes this rapid expansion of the supply of unskilled labor as follows:

> " ... between 1740 and 1840 the population of England ... went up from 6 million to 15.7 million. ... while the agricultural labor force represented 60-70% of the total work force in 1740, by 1840 it represented only 22%."

Habakkuk's well-known account of 19th-century technological development (pp. 136-137) also emphasizes the increase in the supply of unskilled labor in English cities, and attributes it to five sources. First, "technical changes in agriculture increased the supply of labor available to industry" (p. 137). Second, "population was increasing very rapidly" (p. 136). Third, labor reserves of rural industry came to the cities. Fourth, "there was a large influx of labor from Ireland" (p. 137). Finally, enclosures released substantial labor from agriculture.[2] According to our model of directed technological change, this increase in the relative supply of unskilled labor should have encouraged unskill-biased technological change, and this is consistent with the patterns discussed above.

In addition to accounting for the recent skill-biased technological developments and for the historical technologies that appear to have been biased towards unskilled workers, this

[2]Although there is general agreement that the supply of unskilled labor in British cities increased between the 18th and 19th centuries, there is disagreement among economic historians about the importance of various factors. For example, Habakkuk's emphasis on the importance of enclosures is challenged by Allen (1992, 2004). The exact contribution of the enclosures movement to the increase in the supply of unskilled labor is secondary for the argument here.

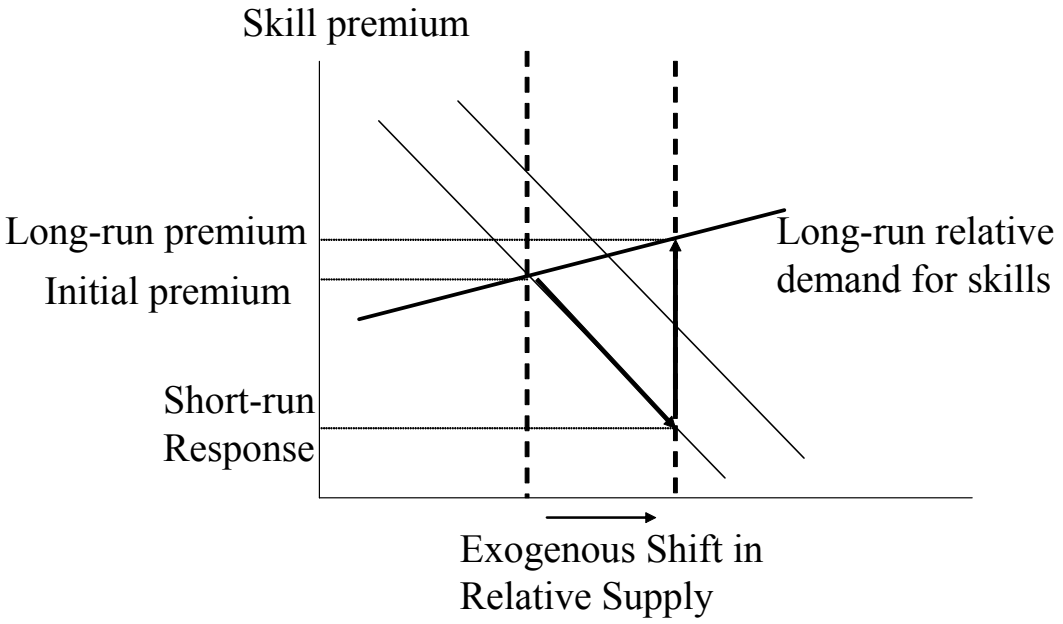

FIGURE 15.4. Dynamics of the skill premium in response to an exogenous increase in the relative supply of skills, with an upward-sloping endogenous-technology relative demand curve.

framework also gives a potential interpretation for the dynamics of the college premium during the 1970s and 1980s. It is reasonable to presume that the equilibrium skill bias of technologies, N_H/N_L, is a sluggish variable determined by the slow buildup and development of new technologies (as the analysis of transitional dynamics in Proposition 15.2 shows). In this case, a rapid increase in the supply of skills would first reduce the skill premium as the economy would be moving along a constant technology (constant N_H/N_L) curve as shown in Figure 15.4. After a while technology would start adjusting, and the economy would move back to the upward sloping relative demand curve, with a relatively sharp increase in the college premium. This approach can therefore explain both the decline in the college premium during the 1970s and its subsequent large surge, and relates both of these phenomena to the large increase in the supply of skilled workers.

If, on the other hand, $\sigma < 2$, then the long-run relative demand curve will be downward sloping, though again it will be shallower than the short-run relative demand curve. Following the increase in the relative supply of skills there will again be an initial decline in the college premium, and as technology starts adjusting the skill premium will increase. But it will end up below its initial level. To explain the larger increase in the college premium in the 1980s, in this case we would need some exogenous skill-biased technological change. Figure 15.5 draws this case.

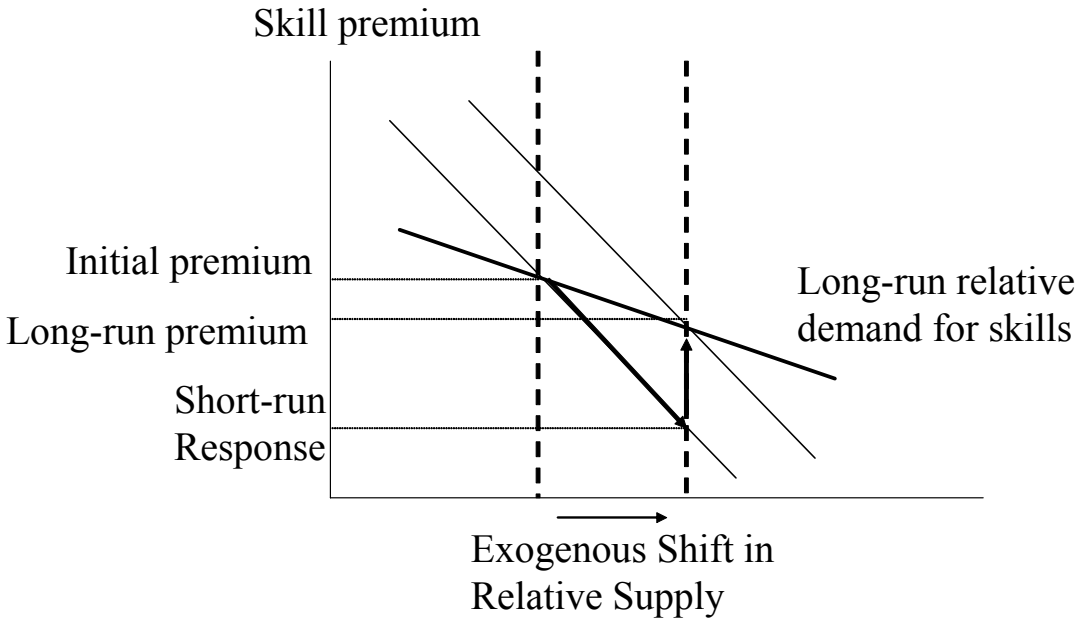

FIGURE 15.5. Dynamics of the skill premium in response to an increase in the relative supply of skills, with a downward-sloping endogenous-technology relative demand curve.

Consequently, a model of directed technological change can shed light both on the secular skill bias of technology and on the relatively short-run changes in technology-induced factor prices. Other implications of these results will be discussed below. However, before doing this, a couple of further issues are worth noting. First, Proposition 15.4 shows that upward-sloping relative demand curves arise only when $\sigma > 2$. In the context of substitution between skilled and unskilled workers, an elasticity of substitution much higher than 2 is unlikely. Most estimates put the elasticity of substitution between 1.4 and 2. One would like to understand whether $\sigma > 2$ is a feature of the specific model discussed here and how different assumptions about the technology of production or the innovation possibilities frontier affect this result. This issue will be discussed in Section 15.4. Second, we would like to understand the relationship between the market size effect and the scale effects, in particular, whether the results on induced technological change are an artifact of the scale effect (which many economists do not view as an attractive feature of endogenous technological change models). Section 15.5 shows that this is not the case and exactly the same results apply when scale effects are removed. Third, we would like to apply these ideas to investigate whether there are reasons for technological change to be endogenously labor-augmenting in the neoclassical growth model. This will be investigated in Section 15.6. Finally, it is also useful to contrast equilibrium allocation to the Pareto optimal allocation. We will start with this latter comparison in the next subsection.

15.3.4. Pareto Optimal Allocations. The analysis of Pareto optimal allocation is very similar to the analysis of optimal growth in Chapter 13. For this reason, we will present only a sketch of the argument. As in that analysis, it is straightforward to see that the social planner would not charge a markup on machines, thus

$$x_L^S(\nu,t) = (1-\beta)^{-1/\beta} p_L(t)^{1/\beta} L \text{ and } x_H^S(\nu,t) = (1-\beta)^{-1/\beta} p_H(t)^{1/\beta} H.$$

Combining these with the production function and some algebra establish that net output, which can be used for consumption or research, is equal to (see Exercise 15.6):

$$Y^S(t) = (1-\beta)^{-1/\beta} \beta \left[\gamma_L^{\varepsilon} \left(N_L^S(t) L \right)^{\frac{\sigma-1}{\sigma}} + \gamma_H^{\varepsilon} \left(N_H^S(t) H \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}. \tag{15.31}$$

In view of this, the current-value Hamiltonian for the social planner can be written as

$$\mathcal{H}\left(N_L^S, N_H^S, Z_L^S, Z_H^S, C^S, \mu_L, \mu_H\right) = \frac{C^S(t)^{1-\theta}-1}{1-\theta} + \mu_L(t)\eta_L Z_L^S(t) + \mu_H(t)\eta_H Z_H^S(t),$$

subject to

$$C^S(t) = (1-\beta)^{-1/\beta} \left[\gamma_L^{\varepsilon} \left(N_L^S(t) L \right)^{\frac{\sigma-1}{\sigma}} + \gamma_H^{\varepsilon} \left(N_H^S(t) H \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} - Z_L^S(t) - Z_H^S(t).$$

The necessary conditions for this problem give the following characterization of the Pareto optimal allocation in this economy.

PROPOSITION 15.5. *The stationary solution of the Pareto optimal allocation involves relative technologies given by (15.27) as in the decentralized equilibrium. The stationary growth rate is higher than the equilibrium growth rate and is given by*

$$g^S = \frac{1}{\theta}\left((1-\beta)^{-1/\beta} \beta \left[(1-\gamma)^{\varepsilon} (\eta_H H)^{\sigma-1} + \gamma^{\varepsilon} (\eta_L L)^{\sigma-1} \right]^{\frac{1}{\sigma-1}} - \rho \right) > g^*,$$

where g^* is the BGP growth rate given in (15.29).

PROOF. See Exercise 15.7. □

15.4. Directed Technological Change with Knowledge Spillovers

I now consider the directed technological change model of the previous section with knowledge spillovers. This exercise has three purposes. First, it will show how the main results on the direction of technological change can be generalized to a model using the other common specification of the innovation possibilities frontier. Second, this analysis will show that the strong bias result in Proposition 15.4 can hold under somewhat weaker conditions. Third, this formulation will be essential for the study of labor-augmenting technological change in Section 15.6.

The lab-equipment specification of the innovation possibilities frontier is special in one respect: it does not feature for *state dependence*. State dependence refers to the phenomenon in which the path of past innovations affects the relative costs of different types of innovations. The lab-equipment specification implied that R&D spending always leads to the same increase in the number of L-augmenting and H-augmenting machines. I now introduce a specification

with knowledge spillovers, which allows for state dependence. Recall that, as discussed in Section 13.2 in Chapter 13, when there are scarce factors used for R&D, then growth cannot be sustained by continuously increasing the amount of these factors allocated to R&D. Therefore, in order to achieve sustained growth, these factors need to become more and more productive over time, because of spillovers from past research. Here for simplicity, let us assume that R&D is carried out by scientists and that there is a constant supply of scientists equal to S (Exercise 15.18 shows that the results are identical when workers can be employed in the R&D sector). With only one sector, the analysis in Section 13.2 in Chapter 13 indicates that sustained endogenous growth requires $\dot{N}/N$ to be proportional to S. With two sectors, instead, there is a variety of specifications with different degrees of state dependence, because productivity in each sector can depend on the state of knowledge in both sectors. A flexible formulation is the following:

(15.32)
$$\dot{N}_L(t) = \eta_L N_L(t)^{(1+\delta)/2} N_H(t)^{(1-\delta)/2} S_L(t) \text{ and } \dot{N}_H(t) = \eta_H N_L(t)^{(1-\delta)/2} N_H(t)^{(1+\delta)/2} S_H(t),$$

where $\delta \leq 1$, and $S_L(t)$ is the number of scientists working to produce L-augmenting machines, while $S_H(t)$ denotes the number of scientists working on H-augmenting machines. Clearly, market clearing for scientists requires that

$$S_L(t) + S_H(t) \leq S. \tag{15.33}$$

In this specification, δ measures the degree of state-dependence: when $\delta = 0$, there is no state-dependence—$\left(\partial \dot{N}_H/\partial S_H\right) / \left(\partial \dot{N}_L/\partial S_L\right) = \eta_H/\eta_L$ regardless of the levels of N_L and N_H—because both N_L and N_H create spillovers for current research in both sectors. In this case, the results are identical to those in the previous section. In contrast, when $\delta = 1$, there is an extreme amount of state-dependence. In this case, $\left(\partial \dot{N}_H/\partial S_H\right) / \left(\partial \dot{N}_L/\partial S_L\right) = \eta_H N_H/\eta_L N_L$, so an increase in the stock of L-augmenting machines today makes future labor-complementary innovations cheaper, but has no effect on the cost of H-augmenting innovations. This discussion clarifies the role of the parameter δ and the meaning of state dependence. In some sense, state dependence adds another layer of "increasing returns," this time not for the entire economy, but for specific technology lines. In particular, a significant amount of state dependence implies that when N_H is high relative to N_L, it becomes more profitable to undertake more N_H-type innovations.

With this formulation of the innovation possibilities frontier, the free-entry conditions become (see Exercise 15.8):

$$\begin{aligned} \eta_L N_L(t)^{(1+\delta)/2} N_H(t)^{(1-\delta)/2} V_L(t) &\leq w_S(t) \\ \text{and } \eta_L N_L(t)^{(1+\delta)/2} N_H(t)^{(1-\delta)/2} V_L(t) &= w_S(t) \text{ if } S_L(t) > 0. \end{aligned} \tag{15.34}$$

and

$$\begin{aligned} \eta_H N_L(t)^{(1-\delta)/2} N_H(t)^{(1+\delta)/2} V_H(t) &\leq w_S(t) \\ \text{and } \eta_H N_L(t)^{(1-\delta)/2} N_H(t)^{(1+\delta)/2} V_H(t) &= w_S(t) \text{ if } S_H(t) > 0, \end{aligned} \tag{15.35}$$

where $w_S(t)$ denotes the wage of a scientist at time t. When both of these free-entry conditions hold, BGP technology market clearing implies

$$\eta_L N_L(t)^{\delta} \pi_L = \eta_H N_H(t)^{\delta} \pi_H, \tag{15.36}$$

where δ captures the importance of state-dependence in the technology market clearing condition, and profits are not conditioned on time, since they refer to the BGP values, which are constant as in the previous section (recall (15.15)). When $\delta = 0$, this condition is identical to (15.26) in the previous section. Therefore, as claimed above, all of the results concerning the direction of technological change would be identical to those from the lab-equipment specification.

This is no longer true when $\delta > 0$. To characterize the results in this case, let us combine condition (15.36) with eq.'s (15.15) and (15.18). This yields the equilibrium relative technology as (see Exercise 15.9):

$$\left(\frac{N_H}{N_L}\right)^* = \eta^{\frac{\sigma}{1-\delta\sigma}} \gamma^{\frac{\varepsilon}{1-\delta\sigma}} \left(\frac{H}{L}\right)^{\frac{\sigma-1}{1-\delta\sigma}}, \tag{15.37}$$

where recall that $\gamma \equiv \gamma_H/\gamma_L$ and $\eta \equiv \eta_H/\eta_L$. This expression shows that the relationship between the relative factor supplies and relative physical productivities now depends on δ. This is intuitive: as long as $\delta > 0$, an increase in N_H reduces the relative costs of H-augmenting innovations, so for technology market equilibrium to be restored, π_L needs to fall relative to π_H. Substituting (15.37) into the expression for relative factor prices for given technologies, which is still given by (15.19), yields the following long-run (endogenous-technology) relationship between relative factor prices and relative factor supplies:

$$\omega^* \equiv \left(\frac{w_H}{w_L}\right)^* = \eta^{\frac{\sigma-1}{1-\delta\sigma}} \gamma^{\frac{(1-\delta)\varepsilon}{1-\delta\sigma}} \left(\frac{H}{L}\right)^{\frac{\sigma-2+\delta}{1-\delta\sigma}}. \tag{15.38}$$

It can be verified that, as claimed above, when $\delta = 0$ and thus there is no state-dependence in R&D, both of the previous expressions are identical to their counterparts in the previous section, (15.27) and (15.30).

The growth rate of this economy is determined by the number of scientists. In BGP, both sectors grow at the same rate, so we need $\dot{N}_L(t)/N_L(t) = \dot{N}_H(t)/N_H(t)$, or

$$\eta_H N_H(t)^{\delta-1} S_H(t) = \eta_L N_L(t)^{\delta-1} S_L(t).$$

Combining this equation with (15.33) and (15.37), we obtain the following BGP condition for the allocation of researchers between the two different types of technologies,

$$\eta^{\frac{1-\sigma}{1-\delta\sigma}} \left(\frac{1-\gamma}{\gamma}\right)^{-\frac{\varepsilon(1-\delta)}{1-\delta\sigma}} \left(\frac{H}{L}\right)^{-\frac{(\sigma-1)(1-\delta)}{1-\delta\sigma}} = \frac{S_L^*}{S - S_L^*}, \tag{15.39}$$

and the BGP growth rate (15.40) below. Notice that given H/L, the BGP researcher allocations, S_L^* and S_H^*, are uniquely determined. We summarize these results with the following proposition:

PROPOSITION 15.6. *Consider the directed technological change model with knowledge spillovers and state dependence in the innovation possibilities frontier. Suppose that*

$$(1-\theta)\frac{\eta_L\eta_H\left(N_H/N_L\right)^{(\delta-1)/2}}{\eta_H\left(N_H/N_L\right)^{(\delta-1)}+\eta_L}S<\rho,$$

where N_H/N_L is given by (15.37). Then, there exists a unique BGP equilibrium in which the relative technologies are given by (15.37), and consumption and output grow at the rate

$$g^*=\frac{\eta_L\eta_H\left(N_H/N_L\right)^{(\delta-1)/2}}{\eta_H\left(N_H/N_L\right)^{(\delta-1)}+\eta_L}S. \tag{15.40}$$

PROOF. See Exercise 15.10. □

In contrast to the model with the lab-equipments technology, transitional dynamics do not always take the economy to the BGP equilibrium, however. This is because of the additional increasing returns to scale mentioned above. With a high degree of state dependence, when $N_H(0)$ is very high relative to $N_L(0)$, it may no longer be profitable for firms to undertake further R&D directed at labor-augmenting (L-augmenting) technologies. Whether this is so or not depends on a comparison of the degree of state dependence, δ, and the elasticity of substitution, σ. The latter matters because it regulates how prices change as there is an abundance of one type of technology relative to another, and thus determines the strength of the price effect on the direction of technological change. The next proposition analyzes the transitional dynamics in this case.

PROPOSITION 15.7. *Consider the directed technological change model with knowledge spillovers and state dependence in the innovation possibilities frontier. Suppose that*

$$\sigma<1/\delta.$$

Then, starting with any $N_H(0)>0$ and $N_L(0)>0$, there exists a unique equilibrium path. If $N_H(0)/N_L(0)<\left(N_H/N_L\right)^*$ as given by (15.37), then $Z_H(t)>0$ and $Z_L(t)=0$ until $N_H(t)/N_L(t)=\left(N_H/N_L\right)^*$. $N_H(0)/N_L(0)<\left(N_H/N_L\right)^*$, then $Z_H(t)=0$ and $Z_L(t)>0$ until $N_H(t)/N_L(t)=\left(N_H/N_L\right)^*$.

If

$$\sigma>1/\delta,$$

then starting with $N_H(0)/N_L(0)>\left(N_H/N_L\right)^*$, the economy tends to $N_H(t)/N_L(t)\to\infty$ as $t\to\infty$, and starting with $N_H(0)/N_L(0)<\left(N_H/N_L\right)^*$, it tends to $N_H(t)/N_L(t)\to 0$ as $t\to\infty$.

PROOF. See Exercise 15.12. □

Of greater interest for the focus here are the results on the direction of technological change. The first result on weak equilibrium bias immediately generalizes from the previous section:

Proposition 15.8. *Consider the directed technological change model with knowledge spillovers and state dependence in the innovation possibilities frontier. Then, there is always weak equilibrium (relative) bias in the sense that an increase in H/L always induces relatively H-biased technological change.*

Proof. See Exercise 15.13. □

While the results regarding weak bias have not changed, inspection of (15.38) shows that it is now easier to obtain *strong equilibrium (relative) bias*.

Proposition 15.9. *Consider the directed technological change model with knowledge spillovers and state dependence in the innovation possibilities frontier. Then, if*

$$\sigma > 2 - \delta,$$

there is strong equilibrium (relative) bias in the sense that an increase in H/L raises the relative marginal product and the relative wage of the H factor compared to the L factor.

Intuitively, the additional increasing returns to scale coming from state dependence makes strong bias easier to obtain, because the induced technology effect is stronger. When a particular factor, say H, becomes more abundant, this encourages an increase in N_H relative to N_L (in the case where $\sigma > 1$). However, from state dependence, this makes further increases in N_H more profitable, thus has a larger effect on N_H/N_L. Since with $\sigma > 1$ greater values of N_H/N_L tend to increase the relative price of factor H compared to L, this tends to make the strong bias result easier to obtain.

Returning to the discussion of the implications of the strong bias results for the behavior of the skill premium in the US market, Proposition 15.9 implies that values of the elasticity of substitution between skilled and unskilled labor significantly less than 2 may be sufficient to generate strong equilibrium bias. How much lower than 2 the elasticity of substitution can be depends on the parameter δ. Unfortunately, this parameter is not easy to measure in practice, even though existing evidence suggests that there is some amount of state dependence in the R&D technologies. For example, this is confirmed by the empirical finding that most patents developed in a particular industry build upon and cite previous patents in the same industry.

15.5. Directed Technological Change without Scale Effects

This section shows that the market size effect and its implications for the direction of technological change are independent of whether or not there are scale effects. The market size effect here refers to the relative market sizes of the users of two different types of technologies, not necessarily to the scale of the entire economy, whereas the scale effect concerns the impact of the size of the population on the equilibrium growth rate. The results in this section show that it is possible to entirely separate the market size effect responsible for the weak and strong endogenous bias results from the scale effect.

Consider the knowledge-based R&D model of the previous section, but only with limited spillovers from past research. In particular, suppose that eq. (15.32) is modified to

$$\dot{N}_L = \eta_L N_L^{\lambda} S_L \text{ and } \dot{N}_H = \eta_H N_H^{\lambda} S_H, \tag{15.41}$$

where $\lambda \in (0,1]$. In the case where $\lambda = 1$, the knowledge-based R&D formulation of the previous section applies, but now with no state dependence. When $\lambda < 1$, the extent of spillovers from past research are limited, and this economy will not have steady growth in the absence of population growth.

Let us also modify the baseline environment by assuming that totalis population, including the population of scientists, grows at the exponential rate n. With a similar arguments to that in Section 13.3 in Chapter 13, it can be verified that aggregate output in this economy will grow at the rate (see Exercise 15.15):

$$g^* = \frac{n}{1-\lambda}. \tag{15.42}$$

Consequently, even with limited knowledge spillovers, there will be income per capita growth at the rate $n/(1-\lambda)$. As usual, this is because of the amplification to the externalities provided by population growth. It can also be verified that when $\lambda = 1$, there is no balanced growth and output would reach infinity in finite time (see Exercise 15.16).

The important point for the focus here concerns the market size effect on the direction of technological change. To investigate this issue, note that the technology market clearing condition implied by (15.41) is (see Exercise 15.17):

$$\eta_L N_L^{\lambda} \pi_L = \eta_H N_H^{\lambda} \pi_H, \tag{15.43}$$

which is parallel to (15.36). Exactly the same analysis as above implies that equilibrium relative technology can be derived as

$$\left(\frac{N_H}{N_L}\right)^* = \eta^{\frac{\sigma}{1-\lambda\sigma}} \gamma^{\frac{\varepsilon}{1-\lambda\sigma}} \left(\frac{H}{L}\right)^{\frac{\sigma-1}{1-\lambda\sigma}}. \tag{15.44}$$

Now combining this with (15.19)—which still determines the relative factor prices given technology—we obtain

$$\omega^* \equiv \left(\frac{w_H}{w_L}\right)^* = \eta^{\frac{\sigma-1}{1-\lambda\sigma}} \gamma^{\frac{(1-\lambda)\varepsilon}{1-\lambda\sigma}} \left(\frac{H}{L}\right)^{\frac{\sigma-2+\lambda}{1-\lambda\sigma}}. \tag{15.45}$$

This equation shows that even without scale effects, the same results as before continued to hold. Summarizing this:

PROPOSITION 15.10. *Consider the directed technological change model with no scale effects described above. Then, there is always weak equilibrium (relative) bias, meaning that an increase in H/L always induces relatively H-biased technological change.*

PROOF. See Exercise 15.8. □

Proposition 15.11. *Consider the directed technological change model with no scale effects described above. If*

$$\sigma > 2 - \lambda,$$

then there is strong equilibrium (relative) bias in the sense that an increase in H/L raises the relative marginal product and the relative wage of the H factor compared to the L factor.

15.6. Endogenous Labor-Augmenting Technological Change

One of the advantages of the models of directed technological change is that they allow us to investigate why technological change might be purely labor-augmenting as required for balanced growth. This section shows that models of directed technological change create a natural reason for technology to be more labor-augmenting than capital-augmenting. However, under most circumstances, the resulting equilibrium is not purely labor-augmenting and as a result, a BGP fails to exist. However, in one important special case, the model delivers long-run purely labor-augmenting technological changes exactly as in the neoclassical growth model, thus providing a rationale for one of the strong assumptions of the standard growth models.

In thinking about labor-augmenting technological change, it is useful to consider a two-factor model with H corresponding to capital, that is, $H(t) = K(t)$, in the aggregate production function (15.3). Correspondingly, let us use N_L and N_K to denote the varieties of machines in the two sectors. Let us also simplify the discussion by assuming that there is no depreciation of capital. Note also that in this case the price of the second factor, $K(t)$, is the same as the interest rate, $r(t)$, since investing in the capital stock of the economy is a way of transferring consumption from one instant to another.

Let us first note that in the context of capital-labor substitution, the empirical evidence suggests that an elasticity of substitution of $\sigma < 1$ is much more plausible (whereas in the case of substitution between skilled and unskilled labor, the evidence suggested that $\sigma > 1$). An elasticity less than 1 is not only consistent with the available empirical evidence, but it is also economically plausible. For example, with the CES production function an elasticity of substitution between capital and labor greater than 1 would imply that production is possible without labor or without capital, which appears counterintuitive. It would also contradict are standard assumption on aggregate production functions, Assumption 1.

Now, recall that when $\sigma < 1$, factor-augmenting and factor-biased technologies are reversed. Therefore, *labor-augmenting technological change corresponds to capital-biased technological change.* Then, the question becomes: under what circumstances would the economy generate relatively capital-biased technological change? And also, when will the equilibrium technology be sufficiently capital biased that it corresponds to Harrod-neutral technological change? What distinguishes capital from labor is the fact that it accumulates. In other words, most growth models feature some type of capital-deepening, with $K(t)/L$ increasing as the economy grows. This implies that in contrast to the analysis so far, which looked

at the effect of one-time changes in relative supplies, the focus now must be on the implications of continuous changes in the relative supply of capital on technological change. In light of this observation, the answer to the first question above is straightforward: capital deepening, combined with Proposition 15.3, implies that technological change should be more labor-augmenting than capital-augmenting.

The next proposition summarizes the main idea of the previous paragraph. For simplicity, this proposition treats the increase in $K(t)/L$ as a sequence of one-time increases (full equilibrium dynamics are investigated in the next two propositions).

Proposition 15.12. *In the baseline model of directed technological change with $H(t) = K(t)$ as capital, if $K(t)/L$ is increasing over time and $\sigma < 1$, then $N_L(t)/N_K(t)$ will also increase over time, i.e., technological change will be "labor-augmenting".*

Proof. Equation (15.27) or eq. (15.37) together with $\sigma < 1$ implies that an increase in $K(t)/L$ will raise $N_L(t)/N_K(t)$. □

This result already gives us important economic insights. The reasoning of directed technological change indicates that there are natural reasons for technology to be more labor-augmenting than capital-augmenting. While this is encouraging, the next proposition shows that the results are not easy to reconcile with the fact that technological change should be *purely* labor-augmenting (Harrod neutral). To state this result in the simplest possible way and to facilitate the analysis in the rest of this section, let us simplify the analysis and suppose that capital accumulates at an exogenous rate, that is,

$$\frac{\dot{K}(t)}{K(t)} = s_K > 0. \tag{15.46}$$

Then, the next proposition shows a negative result on the possibility of purely labor-augmenting technological change.

Proposition 15.13. *Consider the baseline model of directed technological change with the knowledge spillovers specification and state dependence. Suppose that $\delta < 1$ and capital accumulates according to (15.46). Then, there exists no BGP.*

Proof. See Exercise 15.22. □

Intuitively, even though technological change is more labor-augmenting than capital-augmenting, there is still capital-augmenting technological change in equilibrium. This, combined with capital accumulation, is inconsistent with balanced growth. In fact, even a more negative result can be proved (see again Exercise 15.22): in any asymptotic equilibrium, the interest rate cannot be constant, thus consumption and output growth cannot be constant.

In contrast to these negative results, there is a special case that justifies the basic structure of the neoclassical growth model. This takes place when there is extreme state dependence, that is, $\delta = 1$. This case is, in many ways, quite natural and posits that spillovers are limited to the same class of technologies, so that $\dot{N}_L(t)/N_L(t) = \eta_L S_L(t)$ and $\dot{N}_H(t)/N_H(t) =$

$\eta_H S_H(t)$. In this case, it can be verified that technology market equilibrium implies the following relationship in BGP (see Exercise 15.23):

$$\frac{r(t)K(t)}{w_L(t)L} = \eta^{-1}. \tag{15.47}$$

Thus, directed technological change implies that in the long run, the share of capital is constant in national income. Long-run constant factor shares (combined with capital deepening) means that asymptotically all technological change must be purely-labor-augmenting. More specifically, recall from (15.19) that

$$\frac{r(t)}{w_L(t)} = \gamma^{\frac{\varepsilon}{\sigma}} \left(\frac{N_K(t)}{N_L(t)}\right)^{\frac{\sigma-1}{\sigma}} \left(\frac{K(t)}{L}\right)^{-\frac{1}{\sigma}},$$

where $\gamma \equiv \gamma_K/\gamma_L$ and γ_K replaces γ_H in the production function (15.3). Consequently,

$$\frac{r(t)K(t)}{w_L(t)L(t)} = \gamma^{\frac{\varepsilon}{\sigma}} \left(\frac{N_K(t)}{N_L(t)}\right)^{\frac{\sigma-1}{\sigma}} \left(\frac{K(t)}{L}\right)^{\frac{\sigma-1}{\sigma}}.$$

In this case, (15.47) combined with (15.46) implies that

$$\frac{\dot{N}_L(t)}{N_L(t)} - \frac{\dot{N}_K(t)}{N_K(t)} = s_K. \tag{15.48}$$

Moreover, it can be verified that the equilibrium interest rate is given by (see Exercise 15.24):

$$r(t) = \beta\gamma_K N_K(t) \left[\gamma_L \left(\frac{N_L(t)L}{N_K(t)K(t)}\right)^{\frac{\sigma-1}{\sigma}} + \gamma_L)\right]^{\frac{1}{\sigma-1}}. \tag{15.49}$$

Let us now define a *constant growth path* (CGP) as one in which consumption grows at a constant rate (though, differently from a BGP, the two sectors and the two technology stocks do not need grow at the same rate). From (15.22), this is only possible if $r(t)$ is constant and equal to some r^*. Equation (15.48) then implies that $(N_L(t)L)/(N_K(t)K(t))$ is constant, thus $N_K(t)$ must also be constant. Therefore, eq. (15.48) implies that for the economy to ultimately converge to a CGP, long-run technological change must be purely labor-augmenting. This is summarized in the following proposition:

PROPOSITION 15.14. *Consider the baseline model of directed technological change with the two factors corresponding to labor and capital. Suppose that the innovation possibilities frontier is given by the knowledge spillovers specification with extreme state dependence, that is, $\delta = 1$ and that capital accumulates according to (15.46). Then, there exists a CGP allocation in which there is only labor-augmenting technological change, the interest rate is constant and consumption and output grow at constant rates. Moreover, there cannot be any other CGP allocations.*

PROOF. Part of the proof is provided by the argument preceding the proposition. Exercise 15.25 asks you to complete the proof and show that no other constant CGP allocation can exist. □

Notice that Proposition 15.14 does not imply that all technological change must be Harrod neutral (purely labor-augmenting). Along the transition path, there can be (and in fact there will be) capital-augmenting technological change. However, in the long run (that is, asymptotically or as $t \to \infty$), all technological change will be labor-augmenting.

It can also be verified that the CGP allocation with purely labor-augmenting technological change is globally stable if $\sigma < 1$ (see Exercise 15.26). This is reasonable, especially in view of the results in Proposition 15.7, which indicated that the stability of equilibrium dynamics in the model with the knowledge spillovers requires $\sigma < 1/\delta$. Since here there is extreme state dependence ($\delta = 1$) stability requires $\sigma < 1$. Intuitively, if capital and labor were gross substitutes ($\sigma > 1$), the equilibrium would involve rapid accumulation of capital and capital-augmenting technological change, leading to an asymptotically increasing growth rate of consumption. However, when capital and labor are gross complements ($\sigma < 1$), capital accumulation would increase the price of labor more than proportionately and the profits from labor-augmenting technologies increase more than the profits from capital-augmenting ones. This encourages further labor-augmenting technological change. These strong price effects are responsible for the stability of the CGP allocation in Proposition 15.14. Intuitively, an elasticity of substitution between capital and labor that is less than 1 induces the economy to strive towards a balanced allocation of effective capital and labor units (where "effective" here refers to capital and labor units augmented with their complementary technologies). Since capital accumulates at a constant rate, a balanced allocation implies that the productivity of labor should increase faster, in particular, the economy should converge to an equilibrium path with purely labor-augmenting technological progress.

The results in Proposition 15.14 are important, since they provide a justification for the assumption in the Solow and neoclassical growth models that long-run technological change is purely labor-augmenting. Naturally, whether or not this is the case in practice is an empirical matter and is an interesting topic of future empirical research.

15.7. Generalizations and Other Applications

The results presented so far rely on a range of specific assumptions that are inherent in endogenous technological change models (e.g., Dixit-Stiglitz preferences and linear structure to ensure sustained growth). One may naturally wonder whether the results on weak and strong equilibrium bias generalize to situations in which these assumptions are relaxed. The answer is broadly yes. In Acemoglu (2007), I show that, as long as only factor-augmenting techological changes are possible, the main results presented here also apply in an environment in which production and cost functions take more general forms. In particular, in this general environment, there is always weak (relative) equilibrium bias in response to increases in relative supplies and there will be strong equilibrium bias when the elasticity of substitution is sufficiently high. However, once we allow for a richer menu of technological changes, these results do not necessarily hold. Nevertheless, the essence of the results appears to be much more general. In Acemoglu (2007), I define the complementary notions of weak and strong

absolute equilibrium bias, which refer to whether the equilibrium price of a factor change as the supply of that factor changes (rather than the price of a factor relative to the price of another factor, which is what I have focused on in this chapter). Under very weak regularity assumptions, there is always weak absolute equilibrium bias, in the sense that an increase in the supply of a factor always induces technological change biased in favor of that factor. Moreover, even though standard producer theory implies that an increase in the supply of a factor should reduce price, under plausible assumptions the induced technology effect can be strong enough that the price of the factor that has become more abundant can increase. In this case, there is strong absolute equilibrium bias and the (general equilibrium) demand curves for factors are upward sloping. Since these results require additional notation and somewhat different arguments, I will not present them here.

It is also useful to briefly discuss a number of other important applications of the models of directed technological change. To save space, these are not discussed in detail in the text and are left as exercises for the reader. In particular, Exercise 15.20 shows how this model can be used to shed light on the famous Habakkuk hypothesis in economic history, which relates the rapid technological progress in 19th-century United States to relative labor scarcity. Despite the importance of this hypothesis in economic history, there have been no compelling models of this process. This exercise shows why neoclassical models may have difficulties in explaining these patterns and how a model of directed technological change can account for this phenomenon as long as the elasticity of substitution is less than 1.

Exercise 15.21 shows the effects of international trade on the direction of technological change. It highlights that international trade will often affect the direction in which new technologies are developed, and this often works through the price effect emphasized above.

Exercise 15.27 returns to the discussion of the technological change and unemployment experiences of continental European countries we started with. It shows how a "wage push shock" can first increase equilibrium unemployment, and then induce endogenous capital-biased technological change, which reduces the demand for employment, further increasing unemployment.

Finally, Exercise 15.28 shows how the relative supply of factors can be endogenized and studies the two-way causality between relative supplies and relative technology can be studied.

15.8. An Alternative Approach to Labor-Augmenting Technological Change *

The models presented so far in this chapter are all based on the basic directed technological change framework developed in Acemoglu (1998, 2002). Section 15.6 showed how this approach can be used to provide conditions under which technological change will be endogenously labor-augmenting (recall that this type of technological progress is necessary for balanced growth). An alternative approach to this problem is suggested in the recent paper by Jones (2005). I now briefly discuss this alternative approach.

The models developed so far treat the different types of technologies (e.g., N_L and N_H in the previous sections) as state variables. Thus short-run production functions correspond to

the production possibilities sets for given state variables, while long-run production functions are derived from the production possibilities sets as the state variables themselves also adjust. Jones proposes a different approach, building on a classic paper by Houthakker (1955). Houthakker suggested that the aggregate production function should be derived as the "upper envelope" of different techniques (or "activities"). Each technique or activity corresponds to a particular way of combining capital and labor (thus to a Leontief production function of these two factors of production). However, when a producer has access to multiple ways of combining capital and labor, the resulting envelope will be different than Leontief. In a remarkable result, Houthakker showed that if the distribution of techniques is given by the Pareto distribution (which we already encountered in the previous chapter), this upper envelope will correspond to a Cobb-Douglas production function. Houthakker thus suggested a justification for Cobb-Douglas production functions based on "activity analysis".

Jones builds on and extends these insights. He argues that the long-run production function should be viewed as the upper envelope of different ideas generated over time. At a given point in time, the set of ideas that the society has access to is fixed and these ideas determine the short-run production function of the economy. In the long run, however, the society generates more ideas (either exogenously or via R&D) and the long-run production function is obtained as the upper envelope of this expanding set of ideas. Using a combination of Pareto distribution and Leontief production possibilities for a given idea, Jones shows that there will be a major difference between short-run and long-run production functions. In particular, as in Houthakker's analysis, the long-run production function will take a Cobb-Douglas form and will imply a constant share of capital in national income. However, this is not necessarily the case for short-run production functions. Then, with an argument similar to that in Section 15.6, the economy will adjust from the short-run to long-run production functions by undergoing a form of labor-augmenting technological change.

I now provide a brief sketch of Jones's model, focusing on the main economic insights. As pointed out above, the key building block of Jones's model are "ideas". An idea is a technique for combining capital and labor to produce output. At any given point in time, the economy will have access to a set of ideas. Let us denote the set of possible ideas by $\mathcal{I} \subset \mathbb{R}$ and the set of ideas available at time t by $\mathcal{I}(t) \subset \mathcal{I}$. Each idea $i \in \mathcal{I}$ is represented by a vector (a_i, b_i). The essence of the model is to construct the production possibilities set (frontier) of the economy from the set of available ideas. To do this, we first need to specify how a given idea is used for production. Let us suppose that there is a single final good, Y, which can be produced using any idea $i \in \mathcal{I}$ with a Leontief production function given by

$$Y(t) = \min\{b_i K(t), a_i L(t)\}, \tag{15.50}$$

where $K(t)$ and $L(t)$ are the amounts of capital and labor in the economy. In general, the economy may use multiple ideas, thus $K(t)$ and $L(t)$ should be indexed by i to denote the amount of capital and labor allocated to idea i. However, I will follow Jones and assume that at any point in time the economy will use a single idea. This is a restrictive assumption, but

it simplifies the model and its exposition significantly (see Exercise 15.30). The production function (15.50) makes it clear that a_i corresponds to labor-augmenting productivity of idea i and b_i is its capital-augmenting productivity.

Recall from Section 15.6 that the standard model delivers purely labor-augmenting technological change only under the (special) assumption of extreme state dependence or $\delta = 1$ (cfr. Proposition 15.14). In this model, we also have to make a special assumption, which, following Houthakker, is that ideas are independently and identically drawn from a Pareto distribution. In particular, suppose that each component of the idea, a_i and b_i, are drawn (independently) from two separate Pareto distributions. A random variable y has a Pareto distribution if its distribution function is given by $G(y) = 1 - \Gamma y^{-\alpha}$ for some parameter α. The assumption that a_i and b_i are independently drawn from Pareto distributions then implies that

$$\begin{aligned} \Pr[a_i \leq a] &= 1 - \left(\frac{a}{\gamma_a}\right)^{-\alpha}, \text{ for } a \geq \gamma_a > 0 \\ \Pr[b_i \leq b] &= 1 - \left(\frac{b}{\gamma_b}\right)^{-\beta}, \text{ for } b \geq \gamma_b > 0 \end{aligned}$$

where γ_a, γ_b, α and β are strictly positive constants, and $\alpha + \beta > 1$ (see Exercise 15.32 for the importance of this last inequality).

This Pareto assumption will play a crucial role in the result of this model. It is therefore appropriate to understand what the special features of the Pareto distribution are and why this distribution plays an important role in many different areas of economics. The Pareto distribution has two related special features. One is that its tails are relatively thick (and for this reason, the variance of a variable that is distributed Pareto is infinite, see Exercise 15.31). The second special feature is that if y has a Pareto distribution, then its expected value conditional on being greater than some y', $\mathbb{E}[y \mid y \geq y']$, is proportional to y'. Thus loosely speaking, the Pareto distribution has a certain degree of "proportionality" built into it. The expectation of something better in the future is proportional to what has been achieved today. This makes it quite convenient in the modeling of growth-related processes.

Now, given this structure, let us define the function

$$\mathbb{G}(b,a) \equiv \Pr[a_i \geq a \text{ and } b_i \geq b\,] = \left(\frac{b}{\gamma_b}\right)^{-\beta}\left(\frac{a}{\gamma_a}\right)^{-\alpha} \tag{15.51}$$

as the joint probability $a_i \geq a$ and $b_i \geq b$. Denote the level of aggregate output that can be produced using technique i with capital K and labor L by $\tilde{Y}_i(K,L)$. Before we know the realizations of a_i and b_i for idea i, this level of output is a random variable. Since the production function is Leontief, the distribution of $\tilde{Y}_i$ can be represented by the distribution

function of this variable,

$$
\begin{aligned}
H(y) \equiv \Pr\left[\tilde{Y}_i \leq y\right] &= 1 - \Pr\left[a_i L \geq y \text{ and } b_i K \geq y\right] \\
&= 1 - \mathbb{G}\left(\frac{y}{K}, \frac{y}{L}\right) \\
&= 1 - \gamma K^{\beta} L^{\alpha} y^{-(\alpha+\beta)},
\end{aligned} \tag{15.52}
$$

where the second line follows from the definition of the function $\mathbb{G}$ and the third line from (15.51) with $\gamma \equiv \gamma_a^{\alpha}\gamma_b^{\beta}$. This implies that the distribution of $\tilde{Y}_i$ is also Pareto (provided that $y \geq \min\{\gamma_b K, \gamma_a L\}$).

Let us next turn to the "global" production function, which describes the maximum amount of output that can be produced using any of the available techniques. In other words, let $\tilde{Y}(K,L) = \max_{i\in\mathcal{I}(t)} \tilde{Y}_i(K,L)$. Let $N(t)$ denote the total number of production techniques (ideas) that are available in the set $\mathcal{I}(t)$ (at time t). This equivalently implies that by time t there have been N distinct ideas that have been discovered. Since, by assumption, these N ideas are drawn independently, the "global" production function can be alternatively written as

$$
\tilde{Y}(t; N(t)) = F(K(t), L(t); N(t)) \equiv \max_{i=1,\ldots,N(t)} \min\left\{b_i K(t), a_i L(t)\right\}, \tag{15.53}
$$

where $\tilde{Y}(t; N(t))$ is also a random variable. Since the realization of the N ideas are random, output at time t, $\tilde{Y}(t; N(t))$, conditional on capital $K(t)$ and $L(t)$, is a random variable and we are interested in determining its distribution. Here the fact that the N draws of ideas are independent simplifies the analysis. In particular, the probability that the realization of $\tilde{Y}(t; N(t))$ is less than y is equal to the probability that each of the N ideas will produce less than y. Therefore,

$$
\begin{aligned}
\Pr\left[\tilde{Y}(t; N(t)) \leq y\right] &= H(y)^{N(t)}. \\
&= \left(1 - \gamma K(t)^{\beta} L(t)^{\alpha} y^{-(\alpha+\beta)}\right)^{N(t)}.
\end{aligned} \tag{15.54}
$$

Equation (15.54) makes it clear that as the number of ideas N gets large, the probability that Y is less than any level of y will go to zero. This is simply a restatement of the fact that output will grow without bound, which here follows from the fact that the Pareto distribution has unbounded support. Therefore, we cannot simply determine the distribution of output as $N(t) \to \infty$. Instead, we have to look at aggregate output normalized by an appropriate variable, such as its "expected value" (and apply a type of reasoning similar to the Central Limit Theorem). Given the Pareto distribution, the normalizing factor turns out to be $n(t) \equiv \left(\gamma N(t) K(t)^{\beta} L(t)^{\alpha}\right)^{\frac{1}{\alpha+\beta}}$, so that

$$
\begin{aligned}
\Pr\left[\tilde{Y}(t; N(t)) \leq \left(\gamma N(t) K(t)^{\beta} L(t)^{\alpha}\right)^{\frac{1}{\alpha+\beta}} y\right] &= \left(1 - \gamma K(t)^{\beta} L(t)^{\alpha} \left(n(t)\, y\right)^{-(\alpha+\beta)}\right)^{N(t)} \\
&= \left(1 - \frac{y^{-(\alpha+\beta)}}{N(t)}\right)^{N(t)},
\end{aligned} \tag{15.55}
$$

where the second line, (15.55), makes it clear that $n(t) \equiv \left(\gamma N(t) K(t)^{\beta} L(t)^{\alpha}\right)^{\frac{1}{\alpha+\beta}}$ was indeed the correct normalizing factor. Now recalling that $\lim_{N\to\infty}(1-x/N)^N = \exp(-x)$,

$$\lim_{N(t)\to\infty} \Pr\left[\tilde{Y}(t; N(t)) \leq \left(\gamma N(t) K(t)^{\beta} L(t)^{\alpha}\right)^{\frac{1}{\alpha+\beta}} y\right] = \exp(-y^{-(\alpha+\beta)}) \tag{15.56}$$

for $y > 0$. Equation (15.56) gives the famous *Fréchet distribution*, which is one of the three limiting distributions for extreme values.[3] More specifically, this implies that

$$\frac{\tilde{Y}(t; N(t))}{(\gamma N(t) K(t)^{\beta} L(t)^{\alpha})^{1/\alpha+\beta}} \sim \text{Fréchet}(\alpha+\beta),$$

so that the global production function, appropriately normalized, converges asymptotically to a Fréchet distribution. This means that as $N(t)$ becomes large (which will happen naturally as $t \to \infty$ and more ideas are discovered), the long-run or the "global" production function behaves approximately like

$$\tilde{Y}(t; N(t)) \approx \varepsilon(t) \left(\gamma N(t) K(t)^{\beta} L(t)^{\alpha}\right)^{\frac{1}{\alpha+\beta}} \tag{15.57}$$

where $\varepsilon(t)$ is a random variable drawn from a Fréchet distribution. The intuition for this result is similar to Houthakker's result that aggregation over different units producing with techniques drawn independently from a Pareto distribution leads to a Cobb-Douglas production function The implications are quite different, however. In particular, since the long-run production function behaves approximately as Cobb-Douglas, it implies that factor shares must be constant in the long run. However, the short-run production function (for a finite number of ideas) is not Cobb-Douglas. Therefore, as $N(t)$ increases, the production function evolves endogenously towards the Cobb-Douglas limit with constant factor shares, and as in the analysis in Section 15.6, this means that technological change must ultimately become purely labor-augmenting. Therefore, Jones's model shows that ideas related to Houthakker's derivation of a static production function also imply that the short-run production function will evolve endogenously on average with labor-augmenting technological change dominating the limiting behavior and making sure that the economy, in the long run, acts as if it has a Cobb-Douglas production function.

Although this is an interesting idea and the Pareto distribution appears in many important contexts and has various desirable properties, it is not clear whether it provides a compelling reason for technological change to be labor-augmenting in the long run. Labor-augmenting technological change should be an equilibrium outcome (resulting from the research and innovation incentives of firms and workers). The directed technological change models emphasized how these incentives play out under various scenarios. In the current model, the Cobb-Douglas production function arises purely from aggregation. There is no equilibrium interactions, price or market size effects. Related to this, the unit of analysis is

[3] In particular, a fascinating result in statistics shows that regardless of the distribution $F(y)$, if we take N independent draws from F and look at the probability distribution of the highest draw, then as $N \to \infty$, this distribution converges to one of the following three distributions: the Weibull, the Gumbel or the Frechet. In fact, the Frechet distribution is the most common one. See, for example, Billingsley (1995, Section 14).

unclear. The same argument can be applied to a single firm, to an industry, or to a region. Thus if we are happy with this argument for the economy as a whole, we may also wish to apply it to firms, industries, and regions, concluding that the long-run production function of every unit of production or every firm, industry and region should be Cobb-Douglas. However, existing evidence indicates that there are considerable differences in the production functions across industries and they can not be well-approximated by Cobb-Douglas production functions (see the overview of the evidence on industry and aggregate production functions in Acemoglu 2003a). This suggests that it would be interesting to combine the aggregation of different activities or "ideas" as in Houthakker and Jones's papers with some type of equilibrium structure, which might delineate at what level the aggregation should take place and why it may apply to (some) economies, but not necessarily to single firms or industries.

15.9. Taking Stock

This chapter introduced the basic models of directed technological change. These approaches differ from the endogenous technological change models of the previous two chapters because they not only determine the rate of aggregate technological change, but also endogenize the direction and bias of technological change. The bias of technological change is potentially important for the distributional consequences of the introduction of new technologies (who will be the losers and who will be the winners from technological progress?). Therefore, the bias of technological change will play an important role in the study of political economy of growth.

Equally important, models of directed technological change enable us to investigate a range of new questions. These include the sources of skill-biased technological change over the past 100 years, the causes of acceleration in skill-biased technological change during more recent decades, the causes of unskilled-biased technological developments during the 19th century, the impact of international trade on the direction of technological change, the relationship between labor market institutions and the types of technologies that are developed and adopted, and last but not least, an investigation of why technological change in neoclassical-type models may be largely labor-augmenting.

A relatively simple class of directed technological change models can shed light on all of these questions. These models are quite tractable and allow closed-form solutions for equilibrium relative technologies and long-run growth rates. Their implications for the empirical questions mentioned above stem from two important, and perhaps at first surprising, results, which we can refer to as *weak equilibrium bias* and *strong equilibrium bias* results. The first states that under fairly weak assumptions an increase in the relative supply of a factor *always* induces endogenous changes in technology that are relatively biased towards that factor. Consequently, any increase in the ratio of skilled to unskilled workers or in the capital-labor ratio will have major implications for the relative productivity of these factors. The more surprising result is the strong equilibrium bias one, which states that contrary to

basic producer theory, (relative) demand curves can slope up. In particular, if the elasticity of substitution between factors is sufficiently high, a greater relative supply of a factor causes sufficiently strong induced technological change to make the resulting relative price of the more abundant factor increase. In other words, the long run (endogenous-technology) relative demand curve becomes upward-sloping. The possibility that relative demand curves may be upward-sloping not only has a range of important empirical implications, but also illustrates the strength of endogenous technological change models, since such a result is not possible in the basic producer theory with exogenous technology.

The chapter has concluded with a number of applications of these ideas to a range of empirically important areas. Models of directed technological change are very much in their infancy and there are many theoretical dimensions in which further developments are possible. Perhaps more importantly, there are also numerous applications of these ideas.

Finally, this chapter has also been an important step in the investigation of the causes of cross-country income differences and sources of modern economic growth. Its main lesson is in clarifying the determinants of the nature of technological progress. Technology should not be thought of as a blackbox, but the outcome of decisions by firms, individuals and other agents in the economy. This implies that profit incentives will play a major role in both the aggregate rate of technological progress and also in the biases of the technologies that are being developed and adopted. Models of directed technological change illustrate this reasoning in a sharp way and show a range of its implications.

15.10. References and Literature

Models of directed technological change were developed in Acemoglu (1998, 2002a, 2003a,b, 2007a), Kiley (1999), and Acemoglu and Zilibotti (2001). These papers use the term *directed technical change*, but here I used the related term *directed technological change*, to emphasize continuity with the models of endogenous technological change studied in the previous chapters. The framework presented here builds on Acemoglu (2002a). A somewhat more general framework, with much less functional form restrictions, is presented in Acemoglu (2007a).

Other papers modeling the direction of technological change include Caselli and Coleman (2004), Xu (2001), Gancia (2003), Thoenig and Verdier (2003), Ragot (2003), Duranton (2004), Benabou (2005), and Jones (2005).

Models of directed technological change are closely related to the earlier literature on induced innovation. The induced innovation literature was started indirectly by Hicks, who in *The Theory of Wages* (1932), argued

> "A change in the relative prices of the factors of production is itself a spur to invention, and to invention of a particular kind—directed to economizing the use of a factor which has become relatively expensive." (pp. 124-5).

Hicks' reasoning, that technological change would attempt to economize on the more expensive factor, was criticized by Salter (1960) who pointed out that there was no particular

reason for saving on the more expensive factor—firms would welcome all cost reductions. Moreover, the concept of "more expensive factor" did not make much sense, since all factors were supposed to be paid their marginal product. An important paper by Kennedy (1964) introduced the concept of "innovation possibilities frontier" and argued that it is the form of this frontier—rather than the shape of a given neoclassical production function—that determines the factor distribution of income. Kennedy, furthermore, argued that induced innovations would push the economy to an equilibrium with a constant relative factor share (see also Samuelson, 1965, and Drandakis and Phelps, 1965). Around the same time, Habakkuk (1962) published his important treatise, *American and British Technology in the Nineteenth Century: the Search for Labor Saving Inventions*, where he argued that labor scarcity and the search for labor saving inventions were central determinants of technological progress. The flavor of Habakkuk's argument was one of induced innovations: labor scarcity increased wages, which in turn encouraged labor-saving technological change. Nevertheless, neither Habakkuk nor the induced innovation provided any micro-founded model of technological change or technology adoption. For example, in Kennedy's specification the production function at the firm level exhibited increasing returns to scale because, in addition to factor quantities, firms could choose "technology quantities," but this increasing returns to scale was not taken into account in the analysis. Similar problems are present in the other earlier works as well. It was also not clear who undertook the R&D activities and how they were financed and priced. These shortcomings reduced the interest in this literature, and there was little research for almost 30 years, with the exception of some empirical work, such as that by Hayami and Ruttan (1970) on technological change in American and Japanese agriculture.

The analysis in Acemoglu (1998) and the subsequent work in this area, instead, starts from the explicit micro-foundations of the endogenous technological change models discussed in the previous two chapters. The presence of monopolistic competition avoids the problems that the induced innovations literature had with increasing returns to scale.

Acemoglu (2002a and 2003b) show that the specific way in which endogenous technological change is modeled does not affect the major results on the direction of technological change. This is also illustrated in Exercises 15.19 and 15.29. In addition, even though the focus here has been on technological progress, in Acemoglu (2007a) I showed that all of the results generalize to models of technology adoption as well. In Acemoglu (2007a) I introduced the alternative concept of *weak absolute bias* and *strong absolute bias*, which look at the marginal product of a factor rather than relative marginal product, and proved that there are considerably more general theorems on weak and strong absolute bias. In the text here, I sometimes referred to weak relative bias and strong relative bias to distinguish the results here from the absolute bias results. The results in Acemoglu (2007a) also show that the constant elasticity of substitution aggregator of the output of the two sectors is unnecessary for the results. Nevertheless, I have kept the constant elasticity of substitution structure to simplify the exposition.

Changes in the US wage inequality over the past 60 years are surveyed in Katz and Autor (1999), Autor, Katz and Krueger (1998) and Acemoglu (2002b). The latter paper also discusses how models on directed technological change can provide a good explanation for changes in wage inequality over the past 100 years and also for changes in the direction of technological change in the US and UK economies over the past 200 years. There are many studies estimating the elasticity of substitution between skilled and unskilled workers. The estimates are typically between 1.4 and 2. See, for example, Katz and Murphy (1992), Krusell, Ohanian, Rios-Rull and Violante (1999), and Angrist (1995). A number of these estimates are summarized and discussed in Hamermesh (1993) and Acemoglu (2002b).

Evidence that 19th century technologies were generally labor-complementary (unskilled-biased) is provided by, among others, James and Skinner (1985) and Mokyr (1990), while Goldin and Katz (1998) argued the same for a range of important early 20th century technologies.

Blanchard (1997) discusses the persistence of European unemployment and argues that the phase during the 1990s can only be understood by changes in technology reducing demand for high-cost labor. This is the basis of Exercise 15.27 below. Caballero and Hammour (1999) provide an alternative and complementary explanation to that suggested here.

Acemoglu and Zilibotti (2001) discuss implications of directed technological change for cross-country income differences. We have not dwelled on this topic here, since this will be discussed in greater detail in Chapter 18.4 in the context of appropriate technologies, which is further discussed in Acemoglu (2003b).

Acemoglu (2003b) suggested that increased international trade can cause endogenous skill-biased technological change. Exercise 15.21 is based on this idea. Variants of this story have been developed by Xu (2001), Gancia (2003), Thoenig and Verdier (2003).

The model of long-run purely labor-augmenting technological change presented in Section 15.6 was first proposed in Acemoglu (2003a), and the model presented here is a simplified version of the model in that paper. Similar ideas were discussed informally in Kennedy (1964) and a recent paper by Funk (2002) provides microfoundations for Kennedy's ideas. Section 15.8 builds on Jones (2005), which in turn uses ideas that first appeared in Houthakker (1955). A related paper to Jones (2005) is Lagos (2001), who also uses an approach similar to Houthakker's (1955) and shows how aggregation of productivity across heterogeneous production units can determine aggregate total factor productivity in the economy.

The assumption that the elasticity of substitution between capital and labor is less than 1 receives support from a variety of different empirical strategies. The evidence is summarized in Acemoglu (2003a).

The Habakkuk hypothesis has been widely debated in the economic history literature. It was first formulated by Habakkuk (1962), though Rothbarth (1946) had anticipated these ideas almost two decades earlier. David (1975) contains a detailed discussion of the Habakkuk hypothesis and potential theoretical explanations. Recent work by Allen (2005) argues for the importance of the Habakkuk hypothesis for understanding the British Industrial Revolution.

Exercise 15.20 shows how models of directed technological change can clarify the conditions necessary for this hypothesis to apply.

15.11. Exercises

Exercise 15.1. Derive eq. (15.1).

Exercise 15.2. Complete the proof of Proposition 15.1. In particular, verify that in any BGP, (15.27) must hold and derive the equilibrium growth rate as given by (15.29). Also prove that (15.28) ensures that the two free-entry conditions, (15.20) and (15.21), must hold as equalities. Finally, check that this condition is also sufficient to guarantee that the transversality condition is satisfied. [Hint: calculate the equilibrium interest rate and then use (15.22)].

Exercise 15.3. Prove Proposition 15.2. [Hint: use (15.9) to show that when $N_H(0)/N_L(0)$ does not satisfy (15.27), (15.20) and (15.21) cannot both hold as equalities].

Exercise 15.4. Derive equation 15.30.

Exercise 15.5. Explain why in Proposition 15.1 the effect of γ on the BGP growth rate, (15.29), is ambiguous. When is this effect positive? Provide an intuition.

Exercise 15.6. Derive equation 15.31.

Exercise 15.7. Prove Proposition 15.5. [Hint: first substitute for $C(t)$ from the constraint. Then, show that $\mu_H(t)/\mu_L(t) = (\eta_H(t)/\eta_L(t))^{-1}$. Then, use the necessary conditions with $\dot{\mu}_H(t) = \dot{\mu}_L(t)$].

Exercise 15.8. Derive the free-entry conditions (15.34) and (15.35). Provide an intuition for these conditions.

Exercise 15.9. Derive eq. (15.37).

Exercise 15.10. Prove Proposition 15.6. In particular, check that there is a unique BGP and that the BGP growth rate satisfies the transversality condition.

Exercise 15.11. In the model of Section 15.4, show that an increase in η_H will raise the number of scientists working in H-augmenting technologies in the BGP, S_H^*, when $\sigma > 1$ (and $\sigma < 1/\delta$) and reduce it when $\sigma < 1$. Interpret this result.

Exercise 15.12.
(1) Prove Proposition 15.7. In particular, use eq. (15.9) and show that when (15.37) is not satisfied, both free-entry conditions cannot hold simultaneously. Then, show that if $\sigma < 1/\delta$, there will be greater incentives to undertake research for the technology that is relatively scarce, and the opposite holds when $\sigma > 1/\delta$.
(2) Interpret the economic significance of the condition $\sigma < 1/\delta$. [Hint: relate this to the fact that When $\sigma < 1/\delta$, $\partial\left(N_H^{\delta}V_H/N_L^{\delta}V_L\right)/\partial\left(N_H/N_L\right) < 0$, but the inequality is reversed when $\sigma > 1/\delta$].

Exercise 15.13. Prove Proposition 15.8.

Exercise 15.14. Characterize the Pareto optimal allocation in the model with knowledge spillovers and state dependence (Section 15.4). Show that the relative technology ratio in the stationary Pareto optimal allocation no longer coincides with the BGP equilibrium. Explain why this result differs from that in Section 15.3.

Exercise 15.15. Derive eq. (15.42).

Exercise 15.16. Show that in the model of Section 15.5, if $\lambda = 1$, there exists no BGP.

Exercise 15.17. Derive eq.'s (15.43) and (15.44).

Exercise 15.18. Generalize the model of Section 15.4 so that there are no scientists and the R&D sector also uses workers. Thus the labor market clearing condition is

$$L^E(t) + L_L^R(t) + L_H^R(t) \leq 0,$$

where $L^E(t)$ is employment in the final good sector, and $L_L^R(t)$ and $L_H^R(t)$ denote the employment in the two R&D sectors.

(1) Define the equilibrium and BGP allocations.
(2) Specify the free-entry conditions for each machine variety.
(3) Characterize the BGP equilibrium, show that it is uniquely defined and determine conditions such that the growth rate is positive and the transversality condition is satisfied.
(4) Show that the equivalents of Propositions 15.3 and 15.4 hold in this environment.
(5) Characterize the transitional dynamics and show that they are similar to those in Proposition 15.2.
(6) Characterize the Pareto optimal allocation in this economy and show that the Pareto optimal ratio of technologies in the stationary equilibrium are also given by (15.27).

Exercise 15.19. * Consider a version of the baseline directed technological change model introduced above with the only difference that technological change is driven by quality improvements rather than expanding machine varieties. In particular, let us suppose that the intermediate goods are produced with the production functions:

$$\begin{aligned} Y_L(t) &= \frac{1}{1-\beta}\left[\int_0^1 q_L(\nu,t)x_L(\nu,t \mid q)^{1-\beta}d\nu\right]L^\beta, \text{ and} \\ Y_H(t) &= \frac{1}{1-\beta}\left[\int_0^1 q_H(\nu,t)x_H(\nu,t \mid q)^{1-\beta}d\nu\right]H^\beta. \end{aligned}$$

Producing a machine of quality q costs ψq, where we again normalize $\psi \equiv 1-\beta$. R&D of amount $Z_f(\nu,t)$ directed at a particular machine of quality $q_f(\nu,t)$ leads to an innovation at the flow rate $\eta_f Z_f(\nu,t)/q_f(\nu,t)$ and leads to an improved machine of quality $\lambda q_f(\nu,t)$, where $f = L$ or H, and $\lambda \geq (1-\beta)^{-(1-\beta)/\beta}$, so that firms that undertake an innovation can charge the unconstrained monopoly price.

(1) Define the equilibrium and BGP allocations.
(2) Specify the free-entry conditions for each machine variety.
(3) Characterize the BGP equilibrium, show that it is uniquely defined and determine conditions such that the growth rate is positive and the transversality condition is satisfied.
(4) Show that the relative technologies in the BGP equilibrium are given by (15.27).
(5) Show that the equivalents of Propositions 15.3 and 15.4 hold in this environment.

(6) Characterize the transitional dynamics and show that they are similar to those in Proposition 15.2.
(7) Characterize the Pareto optimal allocation in this economy and show that the Pareto optimal ratio of technologies in the stationary equilibrium are also given by (15.27).
(8) What are the pros and cons of this model relative to the baseline model studied in Section 15.3?

Exercise 15.20. As a potential application of the models of directed technological change, consider the famous Habakkuk hypothesis, which claims that technology adoption in the US economy during the 19th century was faster than in Britain because of relative labor scarcity in the former (which increased wages and encouraged technology adoption).

(1) First, consider a neoclassical-type model with two factors, labor and technology, $F(A, L)$, where F exhibits constant returns to scale and A is a technology term, chosen by each firm with costs $\Gamma(A)$ in terms of the final good. Assume that Γ is continuous, differentiable, strictly increasing and convex. Show that an increase in wages (caused by a decline in labor supply or an exogenous increase in wages because of the minimum wage) cannot increase A.
(2) Next, consider the directed technological change model studied in this chapter with H interpreted as land and assume that N_H is fixed (so that there is only R&D for increasing N_L). Show that if $\sigma > 1$, the opposite of the Habakkuk hypothesis obtains. If in contrast, $\sigma < 1$, the model delivers results consistent with the Habakkuk hypothesis. Interpret this result and explain why the implications are different from the neoclassical model considered in part 1 above.

Exercise 15.21. Consider the baseline model of directed technological change in Section 15.3 and assume that it is in steady state.

(1) Show that in steady state the relative price of the two intermediate goods, p, is proportional to $(H/L)^{\beta}$.
(2) Now assume that the economy opens up to world trade, and faces a relative price of intermediate goods $p' < p$. Derive the implications of this for the endogenous changes in technology. Explain why the results are different from those in the text. [Hint: relate your results to the price effect].

Exercise 15.22. (1) Prove Proposition 15.13. In particular, show that in any BGP equilibrium (15.37) must hold, and that this equation is inconsistent with capital accumulation.
(2) Prove that there exists no equilibrium allocation in which consumption grows at the constant rate. [Hint: show that a relationship similar to (15.37) must hold, and this will lead to an increase in $N_K(t)$, which then implies that the interest rate cannot be constant].

Exercise 15.23. Derive eq. (15.47).

Exercise 15.24. Derive eq. (15.49).

EXERCISE 15.25. * Complete the proof of Proposition 15.14 and show that there cannot exist any other CGP equilibrium.

EXERCISE 15.26. * Show that if $\sigma < 1$, the CGP equilibrium in Proposition 15.14 is globally stable. Show that if $\sigma > 1$, it is unstable. Relate your results to Proposition 15.7.

EXERCISE 15.27. Now let us use the results of Proposition 15.14 to revisit the discussion of the experiences of continental European economies provided in Blanchard (1997). Consider the model of Section 15.6. Discuss how a wage push, in the form of a wage floor above the market clearing level, will first cause unemployment and then if $\sigma < 1$, it will cause capital-biased technological change. Can this model shed light on the persistent unemployment dynamics in continental Europe? [Hint: distinguish two cases: (i) the minimum wage floor is constant; (ii) the minimum wage floor increases at the same rate as the growth of the economy].

EXERCISE 15.28. * The analysis in the text has treated the supply of the two factors as exogenous and looked at the impact of relative supplies on factor prices. Clearly, factor prices can also affect relative supplies. This exercise looks at the joint determination of relative supplies and technologies.

Let us focus on a model with the two factors corresponding to skilled and unskilled labor. Suppose a continuum υ of unskilled agents are born every period, and each faces a flow rate of death equal to υ, so that population is constant at 1 (as in Section 9.8 above). Each agent chooses upon birth whether to acquire education to become a skilled worker. For agent x it takes K_x periods to become skilled, and during this time, he earns no labor income. The distribution of K_x is given by the distribution function $\Gamma(K)$ which is the only source of heterogeneity in this economy. The rest of the setup is the same as in the text. Suppose that $\Gamma(K)$ has no mass points. Define a BGP as a situation in which H/L and the skill premium remain constant.

(1) Show first that in BGP, there is a single-crossing property: if an individual with cost of education K_x chooses schooling, another with $K_{x'} < K_x$ must also acquire skills. Conclude from this that there exists a cutoff level of talent, $\bar{K}$, such that all $K_x > \bar{K}$ do not get education.

(2) Show that, along BGP relative supplies take the form:

$$\frac{H}{L} = \frac{\Gamma(\bar{K})}{1-\Gamma(\bar{K})}.$$

Explain why such a simple expression would not hold away from the BGP.

(3) How would you determine $\bar{K}$? [Hint: agent with talent $\bar{K}$ has to be indifferent between acquiring skills and not].

Show that as $\upsilon \to 0$, the relative supply of skills as a function of the skill premium becomes

$$\frac{H}{L} \approx \frac{\Gamma\left(\ln \omega / \left(r^* + \upsilon - g^*\right)\right)}{1-\Gamma\left(\ln \omega / \left(r^* + \upsilon - g^*\right)\right)},$$

where r^* and g^* refer to the BGP interest-rate and growth rate.

(4) Determine the BGP skill premium by combining this equation with (15.27) and (15.30). Can there be multiple equilibria? Explain the intuition.

EXERCISE 15.29. * Consider an economy with a constant population and risk neutral consumers discounting the future at the rate r. The population consists of L unskilled workers, H skilled workers and S scientists. Utility of each agent is defined over the final good, which is produced as

$$Y(t) = \left[\int_0^n y(\nu,t)^{\frac{\varepsilon-1}{\varepsilon}} d\nu\right]^{\frac{\varepsilon}{\varepsilon-1}}$$

where $\varepsilon > 1$ and intermediate $y(\nu,t)$ can be produced using either skilled or unskilled labor. In particular, when a new intermediate is invented, it is first produced using skilled labor only, with the production function $y(\nu,t) = h(\nu,t)$, and eventually, another firm may find a way to produce this good using unskilled labor with the production function $y(\nu,t) = l(\nu,t)$. Assume that new goods are created using scientists according to the innovation possibilities frontier

$$\dot{n}(t) = b_n n(t)^{\delta} m(t)^{1-\delta} S_n(t) \text{ and } \dot{m}(t) = b_m m(t)^{\delta} n(t)^{1-\delta} S_m(t),$$

with $\delta < 1$ and $S_n(t)$ and $S_m(t)$ are the number of scientists allocated to the two types of goods with $S_n(t) + S_m(t) \leq S$. Denote the wage of scientists at time t by $\omega(t)$. A firm that invents a new good becomes the monopolist producer, but can be displaced by a new monopolist who finds a way of producing the good using unskilled labor.

(1) Denote the unskilled wage by $w(t)$ and the skilled wage by $v(t)$. Show that, as long as $v(t)$ is sufficiently larger than $w(t)$, the instantaneous profits of a monopolist producing skill-intensive and labor-intensive goods at time t are

$$\pi_h(t) = \frac{1}{\varepsilon-1}\frac{v(t)H}{n(t)-m(t)} \text{ and } \pi_l(t) = \frac{1}{\varepsilon-1}\frac{w(t)L}{m(t)}.$$

Interpret these equations. Why is the condition that $v(t)$ is sufficiently larger than $w(t)$ necessary?

(2) Assume moreover that a firm that undertakes R&D to replace the skill-intensive good has an equal probability of replacing any of the existing $n-m$ skill-intensive goods. Define a BGP as an allocation where n and m grow at the same rate g. Show that this implies that output and wages of skilled and unskilled workers must grow at the rate $g/(\varepsilon-1)$. [Hint: use the equation for the numeraire setting the price of the final good equal to 1 at each date].

(3) Show that on the BGP the wages of scientists also grow at the same rate as the wages of skilled and unskilled workers.

(4) Show that the BGP must satisfy the following condition

$$b_n \mu^{1-2\delta} \frac{vH}{r-(2-\varepsilon)g/(1-\varepsilon)+\mu g/(1-\mu)} = b_m \frac{wL}{r-(2-\varepsilon)g/(1-\varepsilon)}$$

where $\mu \equiv m/n$. [Hint: note that a monopolist producing a labor-intensive good will never be replaced, but a monopolist producing a skill-intensive good faces a constant

flow rate of being replaced; also use the fact that use the fact that $\dot{m}/(n-m) = g\mu/(1-\mu)$].

(5) Using consumer demands over varieties (i.e., the fact that $y(\nu,t)/y(\nu',t) = (p(\nu,t)/p(\nu',t))^{-\varepsilon}$), characterize the BGP path level of μ. What is the effect of an increase in H/L on μ? Interpret.

(6) Why was it necessary to impose $\delta < 1$ in the innovation possibilities frontier? Briefly discuss how the analysis would change if $\delta = 1$.

Exercise 15.30. * Consider the model presented in Section 15.8.

(1) Show that if capital and labor are allocated in competitive markets, in general more than one technique will be used in equilibrium. [Hint: construct an example in which there are three ideas $i = 1, 2$ and 3, such that when only one can be used, it will be $i = 1$, but output can be increased by allocating some of labor and capital to ideas 2 and 3].

(2) Show that in this case the exact aggregation result used in Section 15.8 does not apply.

Exercise 15.31. * Suppose that y has a Pareto distribution given by $G(y) = 1 - By^{-\alpha}$. Show that the variance of y is infinite.

Exercise 15.32. * Suppose that y has a Pareto distribution given by $G(y) = 1 - By^{-\alpha}$. Show that

$$\mathbb{E}\left[y \mid y \geq y'\right] = \frac{\alpha}{\alpha - 1}\frac{y'}{C^{-\alpha}}.$$

What happens if $\alpha < 1$?