PRAISE FOR *Modeling Monetary Economies: Fourth Edition*

"This is my favorite textbook for my undergraduate course in monetary economics. It requires only a small investment in order to familiarize the students with the overlapping generations model. Thereafter, the book covers a broad set of topics by building simple extensions of the basic model. My students love it and I can also highly recommend it to any reader interested in money, banking, and monetary policy."

– Professor Aleksander Berentsen, *Faculty of Business and Economics, University of Basel, Switzerland*

"In this book, Professor Haslag, a first-rate economist with a theoretical bent, takes the reader into the landscape of modern monetary economics to help us make sense of the deep questions: what does money do, why do we have money, why is banking important, why is inflation bad, and many others. The book uses the overlapping generations model as its workhorse to study a vast range of issues such as money supply, inflation, financial intermediation and banking, and payments systems. The writing, as always, is lucid, precise, and engaging. I can think of no other undergraduate textbook in macro- or monetary economics that infects the reader with as much passion for the craft of economics as this one."

– Professor Joydeep Bhattacharya, *Department of Economics, Iowa State University*

"This book is the most rigorous and accessible treatment of monetary issues based on dynamic, micro-founded models of monetary exchange for advanced undergraduate students. It answers fundamental questions: Why is fiat money valued? Can money coexist with interest-bearing assets? What is the role of banks and central banks? It also addresses topical questions relative to payment systems, liquidity risk, and the effects of the national debt on future growth. It is a must-read for all students eager to learn advanced monetary economics."

– Professor Guillaume Rocheteau, *Department of Economics, University of California, Irvine*

CAMBRIDGE
UNIVERSITY PRESS
www.cambridge.org

ISBN 978-1-316-50867-1
9 781316 508671

Champ
Freeman
Haslag

Modeling Monetary Economies

Fourth Edition

CAMBRIDGE

Modeling Monetary Economies

Fourth Edition

Bruce Champ
Scott Freeman
Joseph Haslag

Modeling Monetary Economies

Too often monetary economics has been taught as a collection of facts about institutions for students to memorize. By teaching from first principles instead, this advanced undergraduate textbook builds on a simple, clear monetary model and applies this framework consistently to a wide variety of monetary questions. Starting with the case in which trade is mutually beneficial, the book demonstrates that money makes people better off, and that government money competes against other means of payments, including other types of government money. After developing each of these topics, the book tackles the issue of money competing against other stores of value, examining issues associated with trade, finance, and modern banking. The book then moves from simple economies to modern economies, addressing the role banks play in making more trades possible, concluding with the information problems plaguing modern banking, which result in financial crises.

Bruce Champ was a Senior Research Economist at the Federal Reserve Bank of Cleveland, and passed away in 2013. Earlier he taught at Virginia Polytechnic Institute, the Universities of Iowa and Western Ontario, and Fordham University. Dr. Champ's research interests focused on monetary economics and his articles have appeared in the *American Economic Review; Journal of Monetary Economics*; *Canadian Journal of Economics*; and *Journal of Money, Credit, and Banking*, among other leading academic publications. He coauthored the first and second editions of *Modeling Monetary Economies* with the late Scott Freeman.

Scott Freeman was a Professor of Economics at the University of Texas, Austin. He taught earlier at Boston College and the University of California, Santa Barbara. Professor Freeman died in 2004 after struggling with Amyotrophic Lateral Sclerosis for several years. Professor Freeman specialized in monetary theory, and his articles appeared in the *Journal of Political Economy*; *American Economic Review*; *Journal of Monetary Economics*; and *Journal of Money, Credit, and Banking*, among other eminent academic journals.

Joe Haslag is Professor and Kenneth Lay Chair in Economics at the University of Missouri. Professor Haslag received his PhD in Economics from Southern Methodist University in 1987. Professor Haslag spent 12 years in the Research Department at the Federal Reserve Bank of Dallas, teaching graduate and undergraduate courses at Southern Methodist University. He visited the Economics Department at Michigan State University in 2000 and the Department of Monetary Economics at Erasmus University in 1994. He has published his research in such prestigious academic journals as the *Journal of Monetary Economics; Journal of Money, Credit, and Banking; The Review of Economics and Statistics; International Economic Review; and Review of Economic Dynamics*, among other leading academic journals.

Modeling Monetary Economies

Fourth Edition

BRUCE CHAMP

SCOTT FREEMAN

JOSEPH HASLAG

University of Missouri-Columbia

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781316508671

First published 2016
3rd printing 2019

Printed in the United Kingdom by TJ International Ltd. Padstow Cornwall

A catalog record for this publication is available from the British Library.

Library of Congress Cataloging in Publication Data
Names: Champ, Bruce, author. | Freeman, Scott, author. | Haslag, Joseph H., author.
Title: Modeling monetary economies / Bruce Champ, Scott Freeman, Joseph Haslag.
Description: Fourth edition. | New York : Cambridge University Press, 2016. | Revised edition of the authors' Modeling monetary economies, 2011. | Includes bibliographical references and index.
Identifiers: LCCN 2016000897| ISBN 9781107145221 (hardback) | ISBN 9781316508671 (paperback)
Subjects: LCSH: Money–Mathematical models.
Classification: LCC HG221 .C447 2016 | DDC 339.5/301–dc23
LC record available at http://lccn.loc.gov/2016000897

ISBN 978-1-107-14522-1 Hardback
ISBN 978-1-316-50867-1 Paperback

Additional resources for this publication at www.cambridge.org/champ4

I dedicate this book to Bruce Champ, a generous friend and a skilled economist. Bruce was loved by many and is missed every day. I am writing this edition to honor his love for economics and his love for friends. He taught me a great deal and I hope to carry on his legacy.

Contents

Preface

Monetary economics is the branch that seeks to explain how people execute trades with one another. In particular, why would a person be willing to accept a colored piece of paper, willingly giving up something valuable? The answer is compelling.

In this fourth edition, we build an undergraduate-level exposition about economies in which these colored pieces of paper are a means of executing trade. The backdrop is the overlapping generations models. Money, with a record-keeping friction, expands the set of allocations that a person can acquire during their lifetime. Once this door is open, the student can begin to dig deeper and deeper into world in which we live. The goal here is to develop a toolkit so that undergraduates can address important questions. After more than 20 years in publication, these models are well within the reach of undergraduates at the intermediate and advanced levels. These elegantly simple models strengthen our fundamental understanding of the most basic questions in monetary economics. How does money promote exchange? What should serve as money? What causes inflation? What are the costs of inflation?

This approach to teaching monetary economics follows the professions general recognition of the need to start building the microeconomic foundations. More directly, our observation is that economists explain aggregate economic phenomena as the implications of the choices of rational people who seek to improve their welfare within their limited means. The use of microeconomic foundations makes macroeconomics easier to understand because the performance of such abstract economic processes as gross domestic product and inflation is linked to something understood by all – rational individual behavior. It brings powerful tools such as indifference curves and budget lines to bear on questions of interest. Finally, the joining of micro- and macroeconomics offers symmetry; instead of studying microeconomics and macroeconomics as independent entities with different tools, there is just economics.

When the first edition of this book was published, inertia and tradition could account for teaching monetary economics as a swamp of institutional details. It was as if monetary economies were only an unchanging set of facts to be memorized. The rapid pace of change in the financial world belies this view. Undergraduates need a way to analyze a wide variety of monetary events and institutional arrangements because the events and institutions of the future will not be the same as those the students learned in the classroom. The teaching of analysis, the heart of a liberal education, is best accomplished by having students learn clear, explicit, and internally consistent models. In this way, students may uncover the links between the assumptions underlying the models and the performance of the model economies and thus apply their lessons to new events or changes in government priorities or policies.

This book implements our goals by starting with the simplest model – the basic overlapping generations model – which we analyze for insights into the most basic questions of monetary economics, including the puzzling demand for intrinsically worthless pieces of paper and the costs of inflation. Of course, such a simple model will not be able to discuss all the issues of monetary economies. Therefore, we proceed in successive chapters by asking which features of actual economies the simple model does not address. We then introduce those neglected features into the model to enable us to discuss the more advanced topics. We believe that this gradual approach allows us to build, step by step, an integrated model of the monetary economy without overwhelming the students.

The book is organized into three parts of increasing complexity. Part I examines money in isolation. Here we take the questions of the demand for fiat money, a comparison of fiat and commodity money, inflation, and exchange rates. In Part II, we add capital, to study money's interaction with other assets, banking, the intermediation of these assets into fiat money, and alternative arrangement of central banking. In Part II, we look at money's effects on saving, investment, output, and non-monetary government debt.

This book is written for undergraduates. Its requirements are no more advanced than the understanding of basic graphs and algebra; calculus is not required. (Those who want to use calculus can find an exposition of this approach in the appendix to Chapter 1.) While the book may prove useful to graduate students as a primer in monetary theory, the main text is pitched at the undergraduate level. This has kept us from a few demanding topics, such as nonstationary equilibria; we hope the reader will be satisfied by the wide range of topics we have been able to discuss within a single, simple framework. Material that is difficult but within the grasp of undergraduates is set apart in appendices and can be easily skipped or inserted. The appendices also have many extensions, such as the model of credit, which instructors may wish to use but are not essential to the main topics.

The references display the most tension between the undergraduates and the technical base in which this approach originated. Whenever possible, we reference material written for undergraduates or general audiences; these references are marked by asterisks. Finally, where undergraduate references were not available, we supply references to academic articles and surveys to offer graduate and advanced undergraduates some places to start with more advance work. This is not intended as a full survey of the advance literature.

The choice of topics to be covered was also difficult. We make no claim to encyclopedic coverage of every topic or opinion related to monetary economics. We limited coverage to the topics most directly linked to money, covering banking (but not finance in general) and government debt (but not macroeconomics in general). We insisted on models with rational agents operating in explicitly specified environments. We also selected topics that could be addressed in the basic framework of the overlapping generations model. In our view, the selected topics are tractably teachable, promoting unity and consistency. We also selected what we best know and understand. We hope that instructors can build on our foundations to fill in any gaps.

To reduce these gaps we added material to examine the 2007 Financial Crisis in the fourth edition. Not since the Great Depression has there been such widespread failure among the set of financial institutions. Liquidity and sudden withdrawals played very big roles during this (hopefully) once-in-our-lifetime event. Monetary economics is uniquely situated to develop models that help us understand financial crises. More important, by building models from first principles, we can examine which policies will help when such events occur. We have greatly expanded our presentations of data and have added new exercises.

In addition, we have updated many of the graphs. We have divided the first chapter into two chapters. By doing so, the student is forced to understand money as a means of overcoming a record-keeping friction that exists in the world. To show how money serves this role, it is important to start with a chapter in which money is not needed in an economy with perfect record keeping. Here, intergenerational credit arrangements develop because trading histories are maintained without using up any resources.

Many have contributed to the development of this book. We owe Neil Wallace a tremendous intellectual debt for impressing upon us the importance of microeconomic theory in monetary economics. Many others have provided helpful suggestions, criticisms, encouragement, and other help during the writing of this book. These include David Andolfatto, Leonardo Auernheimer, Robin Bade, Richard Barnett,Valerie Bencivenga, Joydeep Bhattacharya, Jerry Brozek, Mike Bryan, John Bryant, Douglas Dacy, Siverio Foresi, Greg Hess, Christian Gilles, Paul Gomme, Dennis Jansen, Kam Liu, Mike Loewy, Finn Kydland, Antoine Martin, Helen O'Keefe, John O'Keefe, David Laidler, Michael Parkin, Dan Peled, Pedro Gomis-Porqueras, Guillaume Rocheteau, Steve Russell, Tom Sargent, Pierre Silos, Bruce Smith, Ken Stewart, Dick Tresch, Francois Velde, Paula Hernandez-Verme, Warren Weber, and Steve Williamson. In addition, Rebecca Whitworth, Sumittra Ganguli, Nicholas Pretnar, Lucas Nathe, and Dean Crader were especially helpful on this edition. Overall, we would like to thank the large number of students at Boston College, the University of California at Santa Barbara, the University of Western Ontario, Fordham University, the University of Texas at Austin, Southern Methodist University, Michigan State University, and the University of Missouri-Columbia, who have persevered through the development of this book.

Bruce Champ
Scott Freeman
Joseph Haslag

Part I

Money

Chapter 1

Trade without Money: The Role of Record Keeping

1 Roadmap

In this chapter, the aim is to develop a model of the economy in which trade makes people better off. A model description has four main parts: (i) there is a description of the physical environment, which consists of things like how long does the economy last, who lives there, how long each person lives, what goods are present, and what meetings occur between people; (ii) there is a description of how people get goods, such as things they are endowed with over their lifetime or the ways in which they can produce things; (iii) we need to know what kinds of goods people like and be able to compare different bundles of goods; and (iv) we need a way to combine the different actions that people want to take so that the quantity supplied is equal to quantity demanded. With all four pieces together, we have a model economy.

This is a book that uses a model economy to explain why anyone would be willing to value colored pieces of paper with portraits of famous people. In the remainder of this book, fiat money is what we will call those colored pieces of paper with famous people depicted. Before we offer a view into an economy in which fiat money is valued, we start with a model economy in which information is not hidden and the economy can keep records of every interaction.

There are two main goals. First, we develop the basic framework for studying monetary economics. To analyze problems in which self-interested people trade with one another, it is very useful to build a model of the economy. And in this book you will see that you only need to invest in one framework – the overlapping generations economy – to analyze lots of different problems in monetary economics. Second, we use the model economy to study a case in which money is not present. Here, our aim is to show that trade can occur in economies without money, provided that records of trade histories exist. With a complete history of previous transactions, people can offer and receive gifts. The records keep track of who participates and who should be shunned. Such trades are efficient. This is a good place to start.

2 Beginnings

In this book we will try to learn about monetary economies through the construction of a series of model economies that replicate essential features of actual monetary economies. All such models are simplifications of the complex economic reality in which we live. Model economies are descriptions of the physical environment in which people live, a description of the technology that produces goods and services, what people are endowed with by nature, a description of what people want, and a description of how to solve all these simultaneous problems. The solution is called equilibrium. Model economies are useful because they are able to illustrate key elements of the behavior of people who choose to hold money and to predict the reactions of important economic variables such as output, prices, government revenue, and public welfare to changes in policies that involve money. We start our analysis with a model economy in which information is complete and there is a technology that is costless to operate that records every transaction. If people want to trade with one another – and they do – then it is useful to see how mutually beneficial trade can be accomplished when there are no barriers keeping them from doing so. We will learn what we can from this simple model and then ask how the model fails to adequately represent reality; in particular, what is missing from this idealistic model so that we can match the observation that people value colored pieces of paper. Throughout the book we try to correct the model's oversights by adding, one by one, the features it lacks.

We concentrate on the overlapping generations model. This model, introduced by Paul Samuelson (1958), has been applied to the study of a large number of topics in monetary theory and macroeconomic theory. Among its desirable features are the following:

- Overlapping generations models are easy to solve. Although they can be used to analyze quite complex issues, there are equilibria that are easy to characterize and to find. Many of their predictions may be described on a simple two-dimensional graph.
- Overlapping generations models provide an elegantly parsimonious framework in which to introduce the existence of money. Money in overlapping generations models dramatically facilitates exchange between people who otherwise would be unable to trade.
- Overlapping generations models are dynamic. They demonstrate how behavior in the present can be affected by anticipated future events. They stand in marked contrast to static models, which assume that only current events affect behavior.

We begin this chapter with a very simple version of an overlapping generations model. As we proceed through the book, we introduce extensions to this basic model. These extensions allow us to analyze a variety of interesting issues.

Other model economies share the same three characteristics we identified in the bullet points. Our aim is not to be all encompassing and cover all of these alternatives. Rather, our approach is more topic driven. After building the basic framework, the extensions we introduce are tied to questions. By focusing on the overlapping generations model, we are able to utilize its flexibility. Over time, other model economies with the same three characteristics will likely exhibit the same flexibility, and coverage of the same broad set of topics will be made available.

Therefore, let us turn to the development of the basic overlapping generations model.

3 The Environment

You will quickly see why we call this an overlapping generations economy. Time is divided into equal-sized bits, which we will call periods. For simplicity, we can always call the starting period 1, the next period 2, and so on. When needed, we use the notation t to stand for the time period. People in this economy, however, do not live forever. Indeed, they live for two periods. Anyone born in period $t = 1$ lives in period 1 and period 2, a person born in period $t = 2$, lives in periods 2 and 3, and so on.[1] Generally speaking, anyone born in period $t \geq 1$, is "young" in period t and "old" in period $t + 1$. In each period $t \geq 1$, N_t people are born. Note that we index time with a subscript. For example, in period $t = 2$, N_2 is our notation for the number of people born in period 2. The people born in periods $t = 1, 2, 3, \ldots$ are called the "future generations" of the economy. In addition, in period 1, there are N_0 people that live for just one period. These people are called the "initial old."

Next, we describe the population living in each period. In each period $t \geq 1$, there are N_t young people who were just born and there are N_{t-1} old people. It is the fact that two generations coexist that gives rise to the name overlapping generations. For example, in period t, there are N_{t-1} old people and N_t young people living. Two generations always overlap with each other every period.

For simplicity, there is only one good in this economy. The good is perishable, meaning that it cannot be stored from one period to the next. In this basic setup, each person receives an endowment of the consumption good when young in the first period of life. The amount of this endowment is denoted as y. When old, no one receives any quantity of the consumption good. This pattern of endowments is illustrated in Figure 1.1.

Of course you might think that this simplification is way too costly to help us understand the wide variety of goods and services available in today's economy. But it is also easy to see that if you want to include work effort in this economy, it is easy. Suppose people are endowed with one unit of work time when young. Let

[1] Or, put another way: the number of old period in any date t is the same as the number of young people born at date $t - 1$.

Generation \ Period		1	2	3	4	5	6	7
Initial old	0	**0**						
Future generations	1	**y**	**0**					
	2		**y**	**0**				
	3			**y**	**0**			
	4				**y**	**0**		
	5					**y**	**0**	
							…	…

Figure 1.1. The pattern of endowments. In each period t, generation t is born. Each person lives for two periods. People are endowed with y units of the consumption good when young and 0 units when old. In any given period, one generation of young people and one generation of old people are alive. The name of this model, the overlapping generations model, follows from this generational structure.

people use that time productively, using a technology to transform effort into units of the consumption good. Now, we can interpret the endowment as an endowment of labor – the ability to work. By using this labor endowment (by working), each person is able to obtain a real income of y units of the consumption good. To be even more concrete, consider an economy in which the only consumption good is coconuts. Each young person is capable of climbing the coconut tree and harvesting the edible nut. Old people, however, cannot climb the tree. You could imagine that young people would harvest nuts, storing the harvest in their hut. Unfortunately, coconuts are perishable, going bad before an old person can eat stored nuts.

4 Preferences

People consume the economy's sole commodity and obtain satisfaction – or, in the economist's jargon, utility – from having done so.

4.1 Future Generations

Members of future generations in an overlapping generations model consume both when young and when old. Each person's utility therefore depends on the combination, or bundle, of personal consumption when young and when old. We make the following four assumptions about a person's preferences regarding consumption. The first two assumptions allow you to see that we can assign a consistent numerical value to a bundle of consumption. The second pair of assumptions helps us to draw a picture of a young person's preferences over consumption when young and consumption when old. Based on the consistent numerical value associated with each bundle, the picture is a great device for characterizing the *solution* to each person's lifetime decision problem.

It will be useful to have some notation. We denote the amount of the good that is consumed in the first period of life by a person born in period t with the notation $c_{1,t}$. Similarly, $c_{2,t+1}$ denotes the amount the same person consumes in the second period of life. It is important to note that $c_{2,t+1}$ is consumption that actually occurs in period $t + 1$, when the person born at time t is old. When the time period is not crucial to the discussion, we denote first- and second-period consumption as c_1 and c_2. Let (c_1^a, c_2^a) stand for a bundle of lifetime consumption referred to as Bundle A. Similarly, let (c_1^b, c_2^b) stand for a bundle of lifetime consumption referred to as Bundle B.

Assumption 1 (Completeness) When facing two bundles, a person can provide valid response to two statements. A valid response is either true or false. The two statements are: (1) I get at least as much happiness from Bundle A as I get from Bundle B and (2) I get at least as much happiness from Bundle B as I get from Bundle A.

What do these two true/false answers tell us? If a person says Statement 1 is true and Statement 2 is false, then I can tell that this person gets more happiness from Bundle A than from Bundle B. If the person says Statement 1 is false and Statement 2 is true, then I know that person gets more happiness from Bundle B than from Bundle A. If the person says Statement 1 and Statement 2 are both true, then I know that the person gets the same level of happiness from Bundle A and Bundle B. Thus, Assumption 1 offers a complete description of the happiness obtained from any two bundles. There are three options: Bundle A is preferred to Bundle B, Bundle B is preferred to Bundle A, or the person is indifferent between Bundle A and Bundle B.

Assumption 2 Preferences are transitive.

To illustrate this assumption, I create a third bundle. Let Bundle D be (c_1^d, c_2^d). Transitivity is just an assumption to guarantee consistency. We ask a person to provide valid responses to Statements 1 and 2 for Bundles A, B, and D. Suppose that Bundle A is preferred to Bundle B. Furthermore, suppose Bundle B is preferred to Bundle D. We can ensure that nothing screwy happens insofar as a person satisfying Assumption 2 will prefer Bundle A to Bundle D.

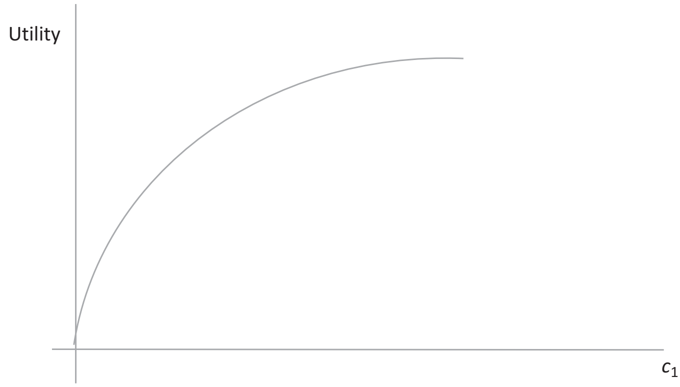

Figure 1.2. Assumptions 3 and 4 are captured. First, the slope of the utility curve is positive, showing that an increase in consumption when young results in greater utility; more is preferred to less. Second, the slope is getting flatter and flatter and the quantity of c_1 increase. Diminishing marginal utility assumes that the marginal utility gain is decreasing as the quantity of the good increases.

Armed with Assumptions 1 and 2, we can define a relationship that assigns a numerical value to each bundle and that numerical value is consistent with the preference ranking obtained from valid responses to Statements 1 and 2. In other words, if Bundle A is preferred to Bundle B, the numerical value assigned to Bundle A – that is, its utility – is greater than the numerical value assigned to Bundle B. If a person is indifferent between Bundle A and Bundle B, for example, the numerical values assigned to each bundle must be equal. The relationship that assigns a numerical value to bundle is called a utility function.

Assumption 3 (More is preferred to less) Suppose Bundle A and Bundle B are constructed so that $c_1^a = c_1^b$ and $c_2^a > c_2^b$. This person is comparing bundles with the same quantity of consumption when young, but when old, Bundle A gives a greater amount of consumption than does Bundle B. According to Assumption 3, this person will always prefer Bundle A to Bundle B.

Assumption 4 (Diminishing marginal utility) The purpose of this assumption is to put some curvature into relationship between bundles. You will know why this is so useful after everything is put together. The simple overview of Assumptions 3 and 4 is that each extra unit you get makes you happier, but extra happiness is getting smaller and smaller with each extra unit. Figure 1.2 graphically depicts the meaning of Assumptions 3 and 4. Figure 1.2 plots the utility value of each extra morsel of consumption when young. Hopefully, from your previous economics classes you remember that marginal utility is defined as the difference between the utility you receive from consuming two different quantities, holding everything else constant. For example, suppose you hold the quantity of consumption when old fixed, call it $\bar{c}_2$ then the marginal utility is the difference in utility value associated with consuming c_1^a and that associated with consuming c_1^b, where $c_1^a > c_1^b$.

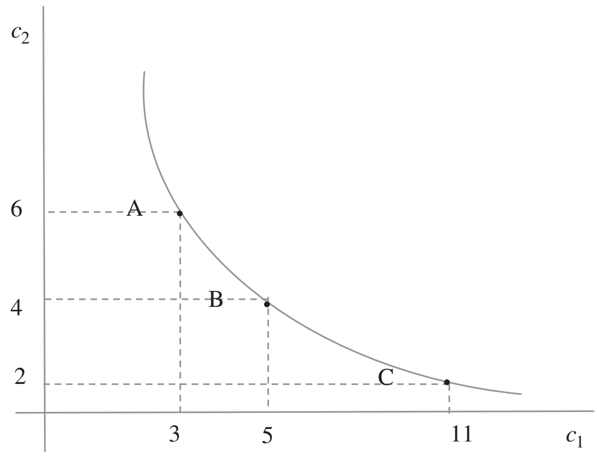

Figure 1.3. An indifference curve. A person's preferences are represented by indifference curves. The figure portrays an indifference curve for a typical person. Along any particular indifference curve, utility is constant. Here, the person is indifferent between Bundles A, B, and C.

Figure 1.2 tells us two things. First, the slope of the utility function is positive, indicating that this person receives greater utility by consuming more of the consumption good when young. Second, the slope is declining, telling us that the marginal value is getting smaller with each additional unit that this person consumes when young.

With Assumptions 1 through 4, we are able to assign a numerical value to every bundle. The utility function is the mathematical representation of a person's preferences over all the bundles. It will be extremely useful to portray a person's preferences graphically. We do this by introducing an indifference curve. An indifference curve connects all the consumption bundles such that there is equal utility. In other words, our young person is saying that for every point on the indifference curve, she responds true to Statements 1 and 2. Figure 1.3 displays a typical indifference curve.

To illustrate the indifference curve, suppose we offer a person the following consumption choices:

- Bundle A, which consists of three units of the consumption good when a person is young and six units of the consumption good when a person is old. We denote this bundle as $c_1 = 3$ and $c_2 = 6$.
- Bundle B, which consists of five units of the consumption good when a person is young and four units of the consumption good when a person is old ($c_1 = 5$ and $c_2 = 4$).

By Assumptions 1 through 4, this person has assigned a numerical value to these bundles by the utility function. On this indifference curve, we show the two points A and B. We also illustrate a third point, C, representing the bundle $c_1 = 11$ and

$c_2 = 2$. Because C lies on the same indifference curve as points A and B, point C yields the same level of utility as points A and B for the person.

Note some features of the indifference curve. The first is that the curve becomes flatter as we move from left to right. This is how indifference curves represent Assumption 4. Note that the slope of the indifference curve is called the marginal rate of substitution. Hence, the curvature of the indifference curve is called the "assumption of diminishing marginal rate of substitution" and diminishing marginal utility can explain this property. To illustrate this assumption, start at point A, where $c_1 = 3$ and $c_2 = 6$. Suppose we reduce the person's second-period consumption by two units. The indifference curve tells us that, to keep the person's utility constant, we must compensate him or her by providing two more units of first-period consumption. This places the person at point B on the indifference curve. Now suppose we reduce second-period consumption by another two units. Our person will remain indifferent if six more units of first-period consumption are provided. In other words, we must compensate a person with ever-increasing amounts of first-period consumption as we successively cut second-period consumption. This should make intuitive sense; people are more reluctant to give up something they do not have much of to begin with.

Consider food and clothing as an example. A person who has a large amount of clothing and very little food would be willing to give up a fairly large amount of clothing for another unit of food. Conversely, this person would be willing to give up only a small amount of food to obtain another unit of clothing.

We demonstrate this assumption of diminishing marginal rate of substitution by drawing an indifference curve that becomes flatter as we move downward and to the right along the curve.

We also assume that the indifference curves become infinitely steep as we approach the vertical axis and perfectly flat as we approach the horizontal axis. The curves never cross either axis. This might be justified by saying that consuming nothing in any one period would mean horrible starvation, to which consuming even a small amount is preferable. This is Assumption 3.

It is also important to keep in mind that the indifference curves are dense in the (c_1, c_2) space. This means that if you pick a combination of first- and second-period consumption, there is an indifference curve running through that point. However, to avoid clutter, we normally show only a few of these indifference curves. A group of indifference curves shown on one graph is often called an "indifference map." Figure 1.4 illustrates an indifference map that obeys our assumptions.

Note that utility is increasing in the direction of the arrow. How do we know this? Compare points A, B, and C. Each of these bundles gives the person the same amount of second-period consumption. However, moving from point A to B to C, the person receives more and more first-period consumption. Hence, the person will prefer point B to point A. Likewise, point C will be preferable to points A and B. This is Assumption 2.

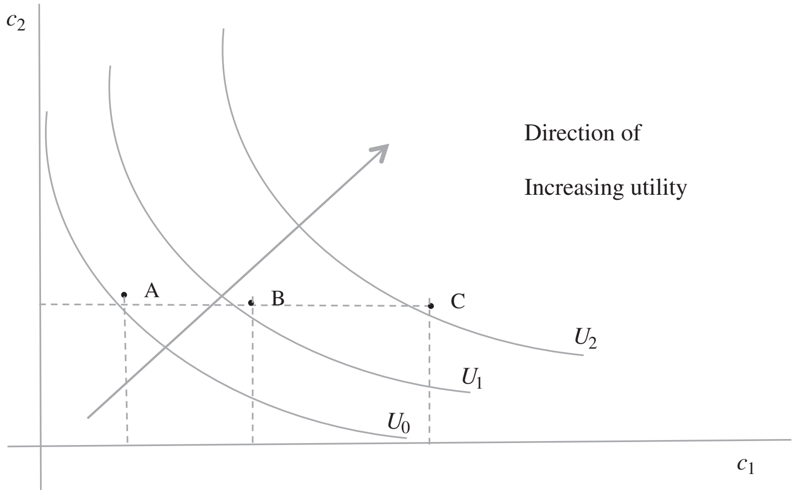

Figure 1.4. An indifference map. An indifference map consists of a collection of indifference curves. For a constant amount of consumption in one period, people prefer a greater amount of consumption in the other period. For this reason, people prefer Bundle C to Bundle B and Bundle B to Bundle A. Utility increase in the general direction of the arrow.

It is often useful to draw an analogy between an indifference map and a contour map that shows elevation. On an indifference map, the curves represent points of constant utility; on a contour map, the curves represent points of constant elevation. Extending the analogy, if we think of traversing the indifference map in a northeasterly direction, we would be going uphill. In other words, utility would be increasing. In fact, an indifference map, like a contour map, is merely a handy way to illustrate a three-dimensional concept on a two-dimensional drawing. The three dimensions here are first-period consumption, second-period consumption, and utility.

One other important concept is that our person's rankings of preferences are transitive. If a person prefers bundle B to bundle A and bundle C to bundle B, then that person must also prefer bundle C to bundle A. Graphically, this implies that indifference curves cannot cross. To do so would violate this property of transitivity and Assumption 2 (see Figure 1.5). This figure portrays two indifference curves that cross at point A. We know that indifference curves represent bundles that give a person the same level of utility. In other words, the person whose preferences are represented by Figure 1.5 is indifferent between Bundles A and B because they lie on the same indifference curve U_0. Similarly, the person must be indifferent between Bundles A and C on indifference curve U_1. We see, then, that the person is indifferent between all three bundles. However, if we compare Bundles B and C, we also observe that they consist of the same amount of second-period consumption but that C contains more first-period consumption than B. According to Assumption 3, the person must prefer C to B. But this contradicts our earlier statement about

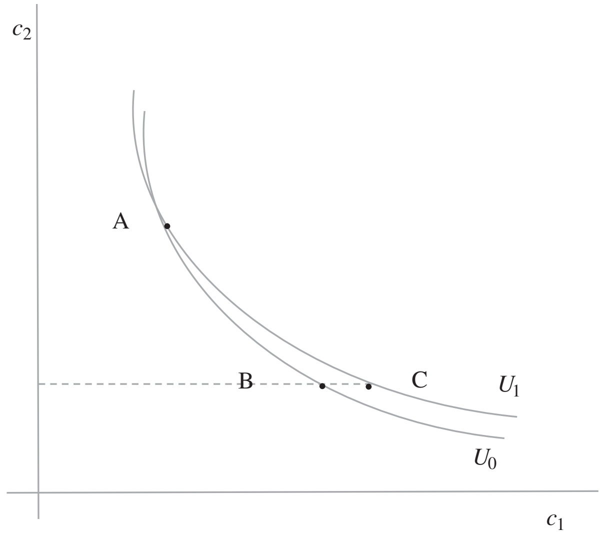

Figure 1.5. Indifference curves cannot cross. By Assumption 3 about preferences, the person preferences are represented by these indifference curves prefers Bundle C over Bundle B because Bundle C consists of more consumption when young and the same amount of consumption when old compared with Bundle B. However, because the person must be indifferent between all three bundles, A, B, and C, a contradiction arises. Our assumptions rule out the possibility of indifference curves that cross.

indifference when comparing the three bundles. For this reason, indifference curves that cross violate our assumptions about preferences.

4.2 The Initial Old

The preferences of the initial old are much easier to describe than those of future generations. The initial old live and consume only in the initial period and thus simply want to maximize their consumption in that period.

4.3 The Never-Ending Economy

You may wonder why time never ends in the physical environment. Actually, inifinity plays an important economic role. Because young people have goods and old people do not, there is a possibility that trade across generations could result in higher lifetime welfare. Indeed, we will show that is true later in the chapter. In this model economy, meetings are one-time interactions between the young and the old. Therefore, any mutually beneficial trade actually requires a never-ending future. Suppose, for example, the economy has a known end date. The young born in the last period will never trade with the old. The reason is simple, the young will never get anything in return. There is nothing mutually beneficial to the young born in the last period to induce them to give up any of their goods. Let's work backward one period. The young born in the next-to-last period know that no young person born next period will trade any goods with them. It follows that young people born in the next-to-last period will not trade with any old people because no trades will occur in the last period. The same decision problem faces people born in the second-to-last period. With knowledge that no young person will trade with them in the next-to-last period, the young born in the second-to-last period will not trade with the old. By applying backward induction over and over, we can see that trade between young and old ceases in *every* period.

Though it may a technical device, the economics of infinity actually ensure that a future of one-time meetings will result in trade.

5 The Economic Problem

The problem facing future generations of this economy is very simple. They want to acquire goods they do not have. Each has access to the nonstorable consumption good only when young but wants to consume in both periods of life. They must therefore find a way to acquire consumption in the second period of life and then decide how much they will consume in each period of life.

We examine, in turn, two solutions to this economic problem. The first, a centralized solution, proposes that an all-knowing, benevolent planner will allocate the economy's resources between consumption by the young and by the old.[2] In the second, decentralized solution, we allow people to use money to trade for what they want. We then compare the two solutions and ask which is more likely to offer people the highest utility. The answer helps provide a first illustration of the economic usefulness of participating in the transfers across generations.

5.1 Feasible Allocations

Imagine for a moment that we are central planners with complete knowledge of and total control over the economy. Our job is to allocate the available goods among the young and old people alive in the economy at each point in time.

As central planners, under what constraint would we operate? Put simply, at any given time, we cannot allocate more goods than are available in the economy. Recall that only the young people are endowed with the consumption good at time t. There are N_t of these young people at time t. We have

$$(\text{total amount of consumption good})_t = N_t y. \tag{1.1}$$

[2] No one believes that such a benevolent central planner exists. For one thing, it is costly to redistribute goods among people. Economists use the "central planner" device to understand what allocations are economically efficient in our model economies. Under the best circumstances, we can gauge how well an economy is doing by comparing the equilibrium in a decentralized economy with the efficient allocation chosen by the fictitious central planner.

Suppose that every member of generation t is given that same lifetime allocation $(c_{1,t}, c_{2,t+1})$ of the consumption good (our society's view of equity). In this case, total consumption by the young people in period t is

$$\text{(total young consumption)}_t = N_t c_{1,t}. \tag{1.2}$$

Furthermore, total old consumption in period t is

$$\text{(total old consumption)}_t = N_{t-1} c_{2,t}. \tag{1.3}$$

Let us make sure the notation is clear. Recall that the old people in time t are those who were born at time $t - 1$. There were N_{t-1} of these people born at time $t - 1$. Furthermore, recall that $c_{2,t}$ denotes the second-period (time t) consumption by someone who was born at time $t - 1$. This implies that total consumption by the old at time t must be $N_{t-1} c_{2,t}$.

Total consumption by young and old is the sum of the amounts in Equations 1.2 and 1.3. We are now ready to state the constraint facing us as central planners: Total consumption by young and old cannot exceed the total amount of available goods (Equation 1.1). In other words,

$$N_t c_{1,t} + N_{t-1} c_{2,t} \leq N_t y. \tag{1.4}$$

For simplicity, we assume for now that the population is constant ($N_t = N$ for all t). In this case, we rewrite Equation 1.4 as

$$N c_{1,t} + N c_{2,t} \leq N y.$$

Dividing through by N, we obtain the per-capita form of the constraint facing us as central planners:

$$c_{1,t} + c_{2,t} \leq y. \tag{1.5}$$

For now, we are also concerned with a stationary allocation.[3] A stationary allocation is one that gives the members of every generation the same lifetime consumption pattern. In other words, in a stationary allocation, $c_{1,t} = c_1$ and $c_{2,t} = c_2$ for every period $t = 1, 2, 3$, and so on. However, it is important to realize that a stationary allocation does not necessarily imply that $c_1 = c_2$. With a stationary allocation, the per-capita constraint becomes

$$c_1 + c_2 \leq y. \tag{1.6}$$

This represents a very simple linear equation in c_1 and c_2, which is illustrated in Figure 1.6.

The set of stationary, feasible, per-capita allocations – the "feasible set" – is bounded by the triangle in the diagram. We refer to the triangular region as the

[3] Nonstationary equilibria have been studied by Azariadas (1981) and by Cass and Shell (1983).

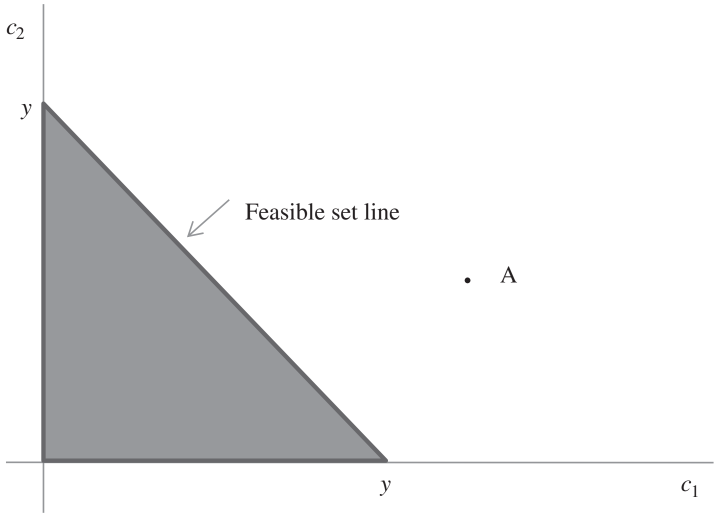

Figure 1.6. The feasible set. The feasible set, the triangle, represents the set of possible allocation that can be attained given the resources available in the economy. Points outside the feasible set, such as Bundle A, are unattainable given the resources of the economy.

feasible set. The thick diagonal line on the boundary of the feasible set is called the "feasible set line." The feasible set line represents Equation 1.6, evaluated at equality.

5.2 The Golden Rule Allocation

If we now superimpose a typical person's indifference map on this diagram, we can identify the preferences of future generations among feasible stationary allocations. This is shown in Figure 1.7.

The feasible allocation a central planner selects depends on the objective. One reasonable and benevolent objective is the maximization of the utility of future generations, an objective we call the "golden rule." The golden rule in Figure 1.7 is represented by point E, which offers each person the consumption bundle (c_1^*, c_2^*). This combination of c_1 and c_2 yields the highest feasible level of utility during a person's entire lifetime. Note that the golden rule occurs at the unique point of tangency between the feasible set boundary and an indifference curve. Any other point that lies within the feasible set yields a lower level of utility. For example, points B and C are feasible because they lie on the boundary of the feasible set. However, they lie on an indifference curve that represents a lower level of utility than the one on which point A lies. Point D is preferable to point A, but it is unattainable. The endowments of the economy simply are not large enough to support the allocation implied by point D.

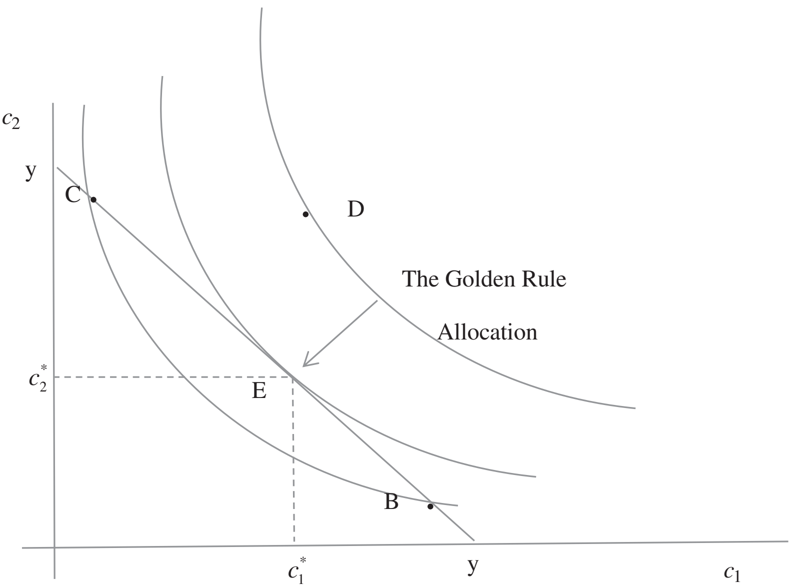

Figure 1.7. The golden rule allocation. The golden rule allocation is the stationary, feasible allocation of consumption that maximizes the welfare of future generations. It is located at a point of tangency between the feasible set line and an indifference curve (Bundle E). This is the highest indifference curve in contact with the feasible set. As drawn, the golden rule allocation E allocates more goods to people when young than when old ($c_1^* > c_2^*$) but this is arbitrary. The tangency point can just as easily have been drawn at a point where $c_2^* > c_1^*$.

5.2.1 The Initial Old

It is important to consider the welfare of all participants in the economy – including the initial old – when considering the effects of any policy. Although the golden rule allocation maximizes the utility of future generations, it does not maximize the utility of the initial old. Recall that the initial old's utility depends solely and directly on the amount of the good they consume in their second period of life. The goal of the initial old is to get as much consumption as possible in period 1, the only period in which they live. (You may want to imagine that the initial old also lived in period 0, however, because this period is in the past, it cannot be altered by the central planner, who assumes control of the economy in period 1.) If the central planner's goal were to maximize the welfare of the initial old, the planner would want to give as much of the consumption good as possible to the initial old. This would be accomplished among stationary feasible allocations at the vertical intercept of the feasible set line in Figure 1.7, which allocates y units of the good for consumption by the old (including consumption by the initial old) and nothing for consumption by the young.

This stationary allocation, which implies that people consume nothing when young, would not maximize the utility of the future generations. They prefer the more balanced combination of consumption when young and old, represented by

(c_1^*, c_2^*). Faced with this conflict in the interests of the initial old and future generations, an economist cannot choose among them on purely objective grounds. Nevertheless, the reader will find that, on subjective grounds (influenced by the fact that there are an infinite number of future generations and only a single generation of initial old), we tend to pay particular attention to the golden rule in this book.

6 Decentralized Solutions

In the previous section, we found the feasible allocation that maximizes the utility of the future generations. However, to achieve this allocation, in each period the central planner would have to take away c_2^* from each young person and give this amount to each old person. Such redistribution requires that the central planner have the ability to reallocate endowments costlessly between the generations. Furthermore, to determine c_1^* and c_2^*, this central planner also must know the exact utility function of the subjects.

These are strong assumptions about the power and wisdom of central planners. This leads us to ask if there is some way we can achieve this optimal allocation in a more decentralized manner, one in which economy reaches the optimal allocation through mutually beneficial trades conducted by the people themselves. In other words, can we let a market do the work of the central planner?

Before we answer this question, we need to define some terms that are used throughout the book. First, we discuss the notion of a competitive equilibrium. A "competitive equilibrium" has the following properties:

1. Each person makes mutually beneficial trades with other people.
2. People act as if their actions have no effect on prices (rates of exchange).
3. Supply equals demand in all markets. In other words, markets clear.

In a sense, the definition of competitive equilibrium tells us how to solve for the equilibrium prices and quantities. Each person maximizes lifetime welfare subject to their budget constraint when determining whether to trade with another person. There is no collusion in the sense that people do not get together to set prices. Rather, each person takes the price as given and maximizes lifetime utility. So a utility-maximizing person chooses the quantity of consumption when young and consumption when old that maximizes lifetime utility, treating the price as given. This solves for the quantities of lifetime consumption as a function of the price. To pin down the price, the definition of competitive equilibrium tells us it is the price that equates the demand for the consumption good with the supply.

We assume there are no frictions in this economy. This means that every young person can observe what trades each old person conducted when they were young. Because no resources are used to see trading histories, record keeping is referred to as perfect. Because young people have goods and old people do not, the key question is whether there is a pattern of mutually beneficial trade that can be sustained. In other words, what kind of trades, if any, will occur between young and old in an economy with perfect record keeping.

6.1 A Record-Keeping Equilibrium

Let us now examine how a person will decide how much to consume when young and how much to consume when old. To answer, we must first establish the constraints on the choices of the person – why he cannot simply enjoy infinite consumption both when young and when old. As was the case for the entire society, the constraints on each person are that he cannot consume more goods than he has. We will refer to the limitations on a person's consumption as his "budget constraints." In this section, we treat the trades "as if" they occur in a market. Unfortunately, there is no mutually beneficial trade between people. Literally, a young person is giving up goods to an old and receiving a record of this transfer. Next period, the now old person has a record of the gift provided when young that everyone can see. Does this record provide any consumption when the person is old?[4] For now, we will pretend that such a sequence of trades over a person's lifetime occurs and solves for the competitive equilibrium.

When young, each person has an endowment of y goods. The person can do two things with these goods – consume them or transfer them to an old person. By perfect record keeping, any goods transferred to an old person are tallied. The quantity of the consumption given to an old person at time t is denoted by Φ_t. We can therefore write the budget constraint facing the person in the first period of life as

$$c_{1,t} + \Phi_t \leq y. \tag{1.7}$$

The left-hand side of Equation 1.7 is the person's total uses of goods (consumption and transfer). The right-hand side of Equation 1.7 represents the total sources of goods (the person's endowment).

When old, no person receives an endowment. Hence, when old, a person only acquires consumption goods by receiving a transfer. This means that the constraint facing the person in the second period of life is

$$c_{2,t+1} \leq \Phi^R_{t+1} \tag{1.8}$$

where Φ^R_{t+1} denotes the goods received as transfer when old. Note that when you are young, the person decides how to transfer to an old person. In contrast, the quantity transferred to an old person is taken as given; that is why we need two separate

[4] The record is not a tangible item, like an IOU. Throughout this book, an IOU is a piece of paper that says Henry, for example, will receive goods from Charles today. The IOU says when Henry will pay Charles back. In the overlapping generations economy, Henry does not meet with Charles in the future because Henry is dead.

pieces of notation. In this model economy, Φ_t is something that each young person decides while Φ^R_{t+1} is, strictly speaking, something that is outside the person's lifetime decision. We will deal with this complication in the next section. But, for now, it is easy to think about the record as a measure of the quantity a young person gave up. With that quantity in the record, the old person can use that as if were something they can trade for consumption goods. Thus, perfect record keeping offers a connections between two people in two different generations who will never meet again.

Suppose a young person would pay Φ_t goods when young in order to obtain Φ^R_{t+1} goods when old. In this hypothetical market, the relationship between the transfer when young and the transfer received when old is represented by an equation that converts the period $t + 1$ goods into what they are worth as period t goods. We could formalize this conversion rate, denoted x_{t+1}, as the rate at which a person can exchange units of the consumption good at period t for units of the consumption good in period $t + 1$. In equation form, $\Phi_t x_{t+1} = \Phi^R_{t+1}$. Armed with this conversion rate and the market, we can rewrite Equation 1.8 as

$$c_{2,t+1} \leq x_{t+1}\Phi_t. \tag{1.9}$$

We are creating a relationship between the gift given when young and the gift received when old. Each person takes the price, x_{t+1}, as given. This trick allows us to write down each young person's lifetime budget constraint. By definition, $x_{t+1} > 0$ for all t, so that we can rewrite the old-age constraint as $\Phi_t \geq (c_{2,t+1})/(x_{t+1})$ and substitute it into the first-period constraint (Equation 1.7) to obtain

$$c_{1,t} + \frac{c_{2,t+1}}{x_{t+1}} \leq y. \tag{1.10}$$

Equation 1.10 expresses the various combinations of first- and second-period consumption that a person can afford over a lifetime. In other words, it is the person's "lifetime budget constraint."

We can graph this budget constraint as shown in Figure 1.8. We can easily verify that the intercepts of the budget line are as illustrated.

The budget line represents Equation 1.10 at equality. If nothing is consumed in the second period of life ($c_{2,t+1} = 0$), then the constraint implies that $c_{1,t} = y$. This is the horizontal intercept of the budget line. On the other hand, if nothing is consumed in the first period of life ($c_{1,t} = 0$), so that the entire endowment of y is used to provide gifts to old people, the constraint implies that $c_{2,t+1}/x_{t+1} = y$ or $c_{2,t+1} = yx_{t+1}$. This represents the vertical intercept of the budget line.

Note that x_{t+1} can be considered as the "(real) rate of return to gifting" because it expresses how many goods can be obtained in period $t + 1$ if one unit of the good is presented as a gift to an old person in period t.

For a given rate of return of money, x_{t+1}, we can find the $(c^*_{1,t}, c^*_{2,t+1})$ combination that will be chosen by people who are seeking to maximize their utility.

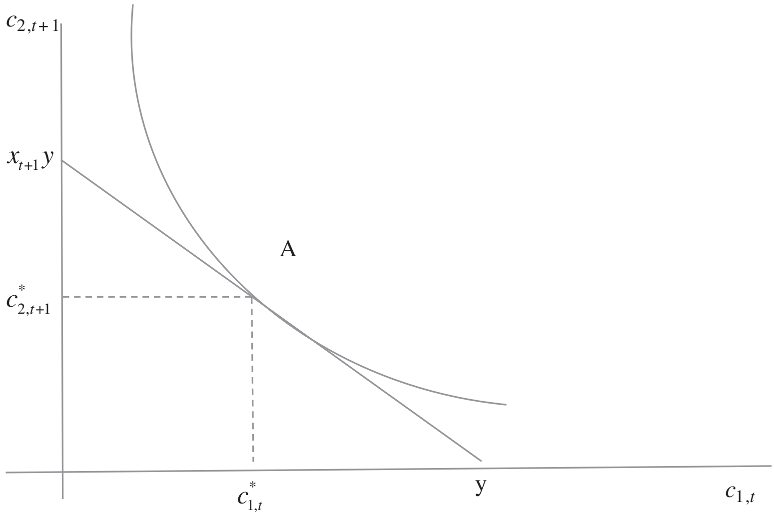

Figure 1.8. The choice of consumption with perfect record keeping. At Bundle A, a person maximizes utility given their lifetime budget constraint in a perfect record keeping equilibrium. Bundle A is found by locating a point of tangency between an indifference curve and the person's lifetime budget set line. The conversion rate on transfers determines the slope of the budget set line.

This point is shown in Figure 1.8. It is the point along the budget line that touches the highest indifference curve. This must occur at a point where the budget line is tangent to an indifference curve.[5]

6.2 Finding the Conversion Rate

How can we determine the rate at which transfers by young people can be converted into transfers received when old? In this setup, the hypothetical market for transfers clears when the supply of transfers by young people is equal to the demand for transfers by old people. In this economy, the conversion rate serves as a two-piece transaction. The young person gives up so many goods to an old person, receives a tally in the record-keeping system, and trades that record in for a transfer when old.

We start by focusing on a case in which the conversion is the same for every generation. This assumption is reasonable because in this basic model, every generation faces the same problem: endowments, preferences and population are the same for every generation. If views about the future are also the same across generations, then each person will react in the same manner each period, choosing $c_{1,t} = c_1$ and

[5] There are other ways for a decentralized economy to build credit arrangements similar to how our old people take their "participation record" to young people in order to get goods. One could imagine a social norm in which people are shunned – not allowed to consume – if they do not participate. In this model economy, the key distinction is that this is not a social norm imposed on people but a market for trade that requires a record of past performance.

$c_{2,t+1} = c_2$ for each period t. We call such equilibria "stationary equilibria." Notice that because each person faces different circumstances, depending on whether they are young or old, we are not imposing c_1 equal to c_2 in a stationary equilibrium. People may choose to consume more when young than when old or vice versa. It turns out that the relative mix of consumption when young and consumption when old depends on people's preferences and on the conversion rate.

We also assume that people in our economy form their expectations of the future rationally. In this nonrandom environment, where there are no surprises, "rational expectations" means that the person's expected values of future variables will be equal to the actual values of these future variables. In this special case, we say that people have perfect foresight. With perfect foresight, there are no errors in a person's forecast of important economic variables that affect their decisions. In the context of our model, this assumption means that a person born in period t will perfectly forecast the conversion rate in the next period, x_{t+1}. The person's expectation of this conversion rate will be exactly realized. This assumption would be less credible in an economy buffeted by random shocks than in our model economy, where preferences and the environment are unchanging and therefore are perfectly predictable.

To see the importance of perfect foresight, consider the alternative in a nonrandom economy – that people always expect a conversion rate greater or less than the conversion rate that actually occurs. People with wrong beliefs about the conversion rate will not choose the transfers that maximize their utility. They therefore have an incentive to figure out the conversion rate that actually will occur.

Let us now employ the assumptions of stationarity and perfect foresight to find an equilibrium time path for the conversion rate. In perfectly competitive markets, the price (or value) of an object is determined as the price at which the supply of the object equals its demand. This applies to the determination of the price (value) of money as well as the price of any good.

The supply of transfers by people is the number of goods each young person offers, which equals the goods of the endowment that the person does not consume when young, $y - c_{1,t}$. The total transfer supply by all people in the economy in period time t is therefore $N_t(y - c_{1,t})$.

The total demand for transfer, measured in units of the young person's consumption good is $N_{t-1}x_t\Phi_{t-1}$, implying that the total supply of transfers is the product of the number of current old people and the value of goods transferred when they were young. Equality of supply and demand therefore requires that

$$N_{t-1}x_t\Phi_{t-1} = N_t(y - c_{1,t}). \tag{1.11}$$

This, in turn, implies that

$$x_t = \frac{N_t(y - c_{1,t})}{N_{t-1}\Phi_{t-1}}, \tag{1.12}$$

which states that the conversion rate for transfers is given by the ratio of the real supply of transfers to the aggregate transfers from the previous period. Since $\Phi_{t-1} = y - c_{1,t-1}$, can write Equation 1.12 as

$$x_t = \frac{N_t(y - c_{1,t})}{N_{t-1}(y - c_{1,t-1})}. \tag{1.13}$$

To simplify this, we look for a stationary solution, where $c_{1,t} = c_1$ and $c_{2,t} = c_2$ for all t. Because all generations have the same endowments and preferences and anticipate the same future pattern of endowments and preferences, it seems quite reasonable to look for a stationary equilibrium. Then, after some cancellation, Equation 1.13 becomes

$$x_t = \frac{N_t}{N_{t-1}}. \tag{1.14}$$

Because we are assuming a constant population ($N_{t+1} = N_t$), the terms in Equation 1.14 cancel out and we find that

$$x_t = 1 \tag{1.15}$$

implying a one-for-one conversion rate.

Notice that the conversion rate is also a constant (1) in the stationary record-keeping equilibrium. Identical people who face the same conversion rate will choose the same consumption and transfers over time, a stationary equilibrium.

Using the information that $x_t = 1$ and recalling that the budget line in a stationary record-keeping equilibrium is represented by $c_1 + \frac{c_2}{x_t} = y$, we determine that $c_1 + c_2 = y$. Our graph of the budget line therefore becomes the one depicted in Figure 1.9.

Be aware that the stationary record-keeping equilibrium may not be a unique equilibrium. There also may exist more complicated nonstationary equilibria. In this text, however, we confine our attention to stationary equilibria because there is much that can be learned from these easy-to-study cases.

6.3 The Game and the Enforcement

In the decentralized solution, the mechanism used was the competitive market. To characterize this equilibrium, we had to pretend that a young person received a marker that was maintained in the record keeping system. By trading with an old person, the record-keeping system costlessly kept information on the quantity traded by each young person. In doing so, the young person could then trade that information for consumption goods next period when they are old.

In this section, we relax the competitive market mechanism. Instead, suppose there is a game young people decide to play, taking what future young people will do as given. In this way, the game does not require an old person to physically

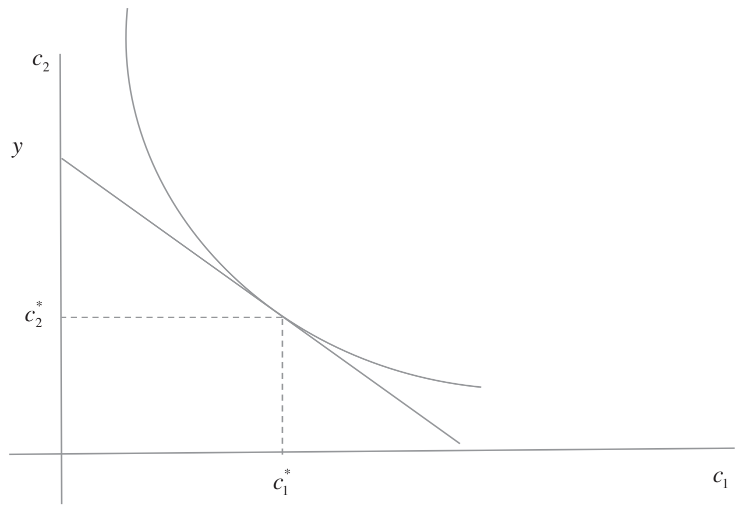

Figure 1.9. A person's choice of consumption when the population is constant. With perfect-record keeping, the conversion rate is 1, implying the lifetime budget constraint of the diagram.

exchange anything with a young person that is mutually beneficial. Here, a young person takes an action, which is a trade with an old person, and then waits, taking the action by next period's young person as given. If you are familiar with the prisoner's dilemma game, you are familiar with the description of a person's best response to the rules of the game. When each person's action is a best response, and there is no unilateral incentive for a person to deviate, we have a Nash equilibrium.

So our goal in this section is to characterize a young person's best response in a world in which perfect record keeping is present. To keep matters simple, suppose we consider a case in which there are two actions; a young person can either not transfer any goods to an old person or transfer $\Phi = c_2^*$ to an old person. Such an action depends on the payoff to the young person next period. If, for example, the young person takes $\Phi^R = c_2^*$ as the action chosen by next period's young person, then the transfer when young is justified as a best response. Indeed, the lifetime welfare for the young person would be strictly lower if any other action were taken.[6]

What role does perfect record keeping play in this game? Here, perfect record keeping is a means of holding a young person accountable. There is no way to hide your actions when young, so the actions each generation takes are costlessly verifiable. Consider a case in which there is a friction and perfect record keeping is absent. Now, suppose a young person decides not to transfer any goods to an old person. In this case, future generations cannot observe this person's actions, lessening the chance that next period's young will offer any transfers to them. Through the perfect record-keeping device, future generations can punish a person who does not

[6] People cannot trade directly in this model because they are separated in time. The same absence of trade would result if they were separated by space, as in the models of Robert Townsend (1980).

participate in the intergenerational transfer program. Now, from the person's utility function, we know they like to consume when old. So the perfect record keeping acts as a kind of contract; basically, if a young person transfers, they will receive some transfer when old. In this way, the game that a person plays requires a best response, understanding that choosing not to transfer results in a punishment in which no transfers are provided when old.

Thus, young people are held accountable by the social accounting system. The records show that you participated when young so that you get goods when old. If you deviate, the records are free for everyone to see and you are excluded from any old-age transfer. We refer to the costless record keeping as being consistent with a frictionless world. Therefore, nothing impedes the smooth process of intergenerational transfers. And it rests squarely on the idea that perfect record keeping means there are no resources needed to enforce the intergenerational transfers.

7 Is the Record-Keeping Equilibrium the Golden Rule?

We have seen that record keeping can provide for old-age consumption, improving the welfare of people otherwise unable to trade. We would like to make the people in our economy not just better off but as well off as possible. It remains to ask, therefore, whether the record-keeping equilibrium results in the best possible allocation of goods. In particular, we would like to see whether the stationary record-keeping equilibrium we have just found maximizes the welfare of future generations. In other words, does the record-keeping equilibrium reach the golden rule?

Compare the budget line of Figure 1.9 with the feasible set line of Figure 1.7. They are identical. The choice of consumption in this record-keeping equilibrium will be identical to the one we found when we were looking at the stationary allocation that was dictated by a central planner who wanted to maximize the utility of the future generations. This implies that the stationary record-keeping equilibrium obeys the golden rule. The introduction of decentralized trade with perfect record keeping not only allows them to reach their maximum feasible utility through trade but, in this case, also allows them to reach their maximum feasible utility. This will not always be the case. The budget set and the feasible set answer different economic questions. The budget set depicts the constraint on a person, whereas the feasible set describes the constraint on the society as a whole.

The initial old are also better off in the record-keeping equilibrium than they were with the autarkic equilibrium. In the record-keeping equilibrium, everyone was among the initial old will receive Φ^R units of the consumption good. This means their consumption will be positive. In the autarkic equilibrium, their consumption would be zero. They are certainly better off in record-keeping equilibrium.

Because we concentrate on stationary record-keeping equilibria in this book, it may be useful to summarize the features of such equilibria. A stationary consumption bundle of a record-keeping equilibrium satisfies two basic properties:

- It provides the maximum level of utility given the person's budget set. It is found where an indifference curve lies tangent to the person's budget set.
- It lies on the feasible set line, with the boundary of the set representing all feasible per-capita allocations.

8 Summary

In this overlapping generations economy, each person has a pair of one-time meetings over their lifetime. When young, a person meets with an old person. Next period, our old person meets with a young person. Each meeting offers an opportunity for trade. The sticking point is that young people have goods but old people do not. Why would a young person give up any goods to someone they will never meet again?

In a world where record-keeping does not cost anything, trading history is a record that can provide an incentive to trade with an old person. The quantity you offer this old person is recorded so that other people, like next period's young, can see whether you participated or not. The key is that there is punishment; no young person will trade with an old person that does not have a record of offering goods when they were young. Each generation compares lifetime welfare when they offer a quantity of goods to each old person with lifetime welfare when they do not. Because giving up goods when young is best for every current and future generation, people get to consume when young and when old. And the absence of any friction associated with keeping these records enables punishment to be applied, thus creating a kind of discipline that keeps the meet-and-trade sequence going forever between the young and old at these one-time meetings.

9 Exercises

1.1. Consider an economy with a constant population in which each person is endowed with y_1 when young and y_2 when old. Assume that y_2 is sufficiently small so that everyone wants to consume more that y_2 in the second period of life. Bear in mind that under the new assumptions, the equations and graphs you use may differ from the ones in this chapter.

a. Apply the Equations 1.1 through 1.6 to find the feasible set.

b. Assume that all people within a generation will be treated alike and graph the set of stationary per-capita feasible allocations. Draw arbitrarily located, but correctly shaped, indifference curves on your graph and point out the allocation that maximizes the utility of the future generations.

1.2. Suppose a person faces the following two bundles: Bundle A, which consists of 6 units of the consumption good when a person is young and 12 units of the consumption good when a person is old ($c_1 = 6$ and $c_2 = 12$); and Bundle B, which consists of 4 units of the consumption good when a person is young and 10 units of the consumption good when a person is old ($c_1 = 4$ and $c_2 = 10$). Which bundle would this person prefer? Which assumption on preferences did you use to draw this conclusion?

1.3. Consider an economy in which the population follows the rule $N_t = 1.1N_{t-1}$. In addition, suppose that endowments per young person grow each period according to $y_t = 1.05y_{t-1}$. Assume old people do not receive any endowments. Assume that a young person's preferences are such that they want to consume one-half of their endowment so that $c_{1,t} = 0.5y_t$. Compute the rate at which transfers by young people can be converted into transfers received when old, that is, the conversion rate for this economy.

1.4. Consider two economies, labelled A and B. In each one, let every two-period-lived person be endowed with 20 units of the consumption when young and nothing when old. In Economy A, each young person chooses to consume 10 units of the consumption good. In Economy B, each young person chooses to consume 8 units of the consumption good. In each economy, the young person's choice is the one that maximizes lifetime welfare.

a. What, if anything, can you infer about the welfare level of the current and future generations from this information? Specifically, is one on an indifference curve representing greater welfare than the other?

b. What, if anything can you infer about the welfare of the initial old from the description given for Economies A and B?

1.5. Suppose a person has constant marginal utility over both goods instead of diminishing marginal utility for consumption when young and consumption when old.

a. Draw an indifference curve for constant-marginal-utility preferences.

b. If the marginal utility of consumption when young were greater than the marginal utility of consumption when old, how would this affect the equilibrium level of consumption over a person's lifetime?

c. What if the marginal utility of consumption when old were greater than the marginal utility of consumption when young?

Chapter 2

A Simple Model of Money

1 Roadmap

Now that we know how people execute trades in a pair of one-time meetings without money, we add a friction into the model that can account for valued fiat money. The key friction is that record keeping is costly. Suppose, for example, that the cost of maintaining such a record-keeping device is prohibitively high. Now young people can fake trading with an old person in order to get goods when they are old. Faking means that there is no way to punish non-participating young people. What should people do in the economy with such a friction?

There are three goals for this chapter. First, we want to demonstrate that the only equilibrium in the decentralized economy results in no trade between the young and the old. We call this non-participating equilibrium autarky. In other words, old people want to use the credit they earned by trading when they were young. If there is no way to verify that an old person did indeed trade when young, such credit arrangements will not be granted. Not surprisingly, autarky is not the efficient equilibrium. Second, we propose valued fiat money as a government policy that can attempt to deal with the friction. Third, it important to show that the stationary equilibrium is efficient in an economy with a constant money stock over time.

Thus, the chief purpose of this chapter is to create a friction and incorporate that into the overlapping generations economy. As we did in the previous chapter, we compare the equilibrium quantities with what a planner would choose for the current and future generations. Even with perfect social memory absent, and the externality associated with this friction, the equilibrium quantities are efficient in an economy with valued fiat money. Because the equilibrium in the monetary economy and in the perfect record-keeping economy is identical, we use this equivalence to say that money is memory.[1]

[1] Memory is an important device in trade. To illustrate the importance, there is a story about two people who commit to having dinner with each other every Saturday night. Neither person can remember whose turn it is

2 The Environment

To arrive at the simplest possible model of money, we must ask ourselves which features are essential to monetary economies. The demand for money is distinct from the demand for the goods studied elsewhere in economics. People want goods for the utility received from their consumption. In contrast, people do not want money in order to consume it; they want money because money helps them get the things they want to consume. In this way, money is a medium of exchange – something acquired to make it easier to trade for the goods whose consumption is desired.

A model of this distinction in the demand for money therefore requires two special features. First, there must be some "friction" to trade that inhibits people from directly acquiring the goods they desire in the absence of money. One of the main messages from Chapter 1 is that people want to smooth their consumption over lifetime. If there are no frictions, they can accomplish this feat through the discipline imposed by social memory; that is, if, when young, Andy participates by giving goods to an old person, then next period, a young person will redeem that memory by giving goods to Andy. Moreover, the level of giving is part of that memory and determines how much is given when young and received when old. Social memory is necessary because trading partners only meet once in their lifetime.[2] So what happens if there is no social memory. To make this point more concrete, let Ψ_t represent the cost of keeping societal records at any date t. We assume that $\Psi_t > N_t y$ so that keeping perfect memory is not feasible in this economy. With the desire to trade, but without memory, the question is whether some other means can be developed that can support trade.

Second, we propose a model in which fiat money is used to deal with social amnesia. Someone must be willing to hold money from one period to the next for this to be even possible. This is necessary because money is an asset held over some period of time, however short, before it is spent. The overlapping generations model makes it possible for people to acquire money when young and use when they are old.

The model economy has the same physical environment as the model we studied in Chapter 1 with the addition of no record keeping. Each period there are young people born. Each person lives for two periods, except for a group that is alive in the first period of the economy, whom we refer to as the initial old. There is a single, perishable consumption good. And each young person receives an amount of that good. Each young person wants to consume when young and when old.

to cook the meal. So they devise a plan to let a stone serve as the memory device. Whoever has the stone is the cook and the other person takes the stone with them at the end of the evening.

[2] The modifier "social" applied to memory is very important. We are not talking about people suffering amnesia. What is crucial to social memory – and therefore, to credit – is that the set of all possible interactions between people can be observed. Individually, you are not an amnesiac, but socially, you are incapable of keeping track of every person's trading history, especially the times they reneged.

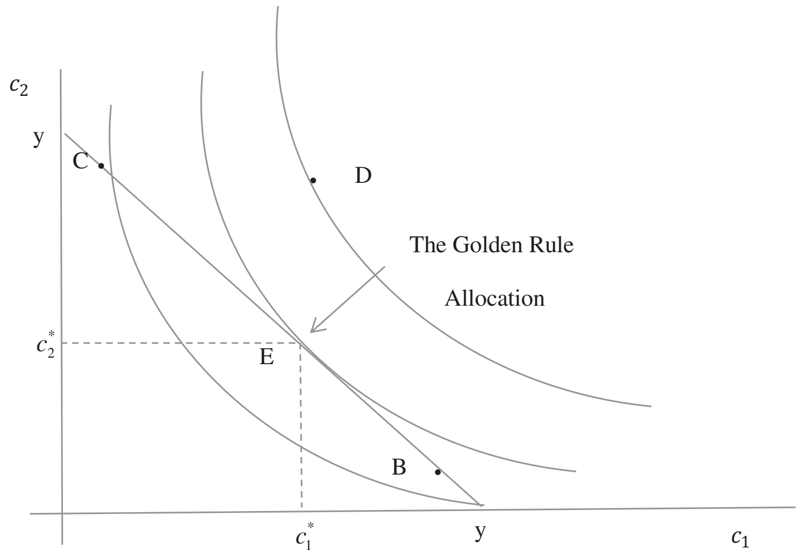

Figure 2.1. The golden rule allocation. The golden rule allocation is the stationary, feasible allocation of consumption that maximizes the welfare of future generations. It is located at a point of tangency between the feasible set line and an indifference curve (Bundle E). This is the highest indifference curve in contact with the feasible set. As drawn, the golden rule allocation E allocates more goods to people when young than when old ($c_1^* > c_2^*$) but this is arbitrary. The tangency point can just as easily have been drawn at a point where $c_2^* > c_1^*$.

3 The Economic Problem

The problem facing future generations of this economy is very simple. They want to acquire goods they do not have. Each has access to the nonstorable consumption good only when young but wants to consume in both periods of life. They must therefore find a way to acquire consumption in the second period of life and then decide how much they will consume in each period of life.

We examine, in turn, two solutions to this economic problem. The first, a centralized solution, proposes that an all-knowing, benevolent planner will allocate the economy's resources between consumption by the young and by the old. In the second, decentralized solution, we allow people to use money to trade for what they want. We then compare the two solutions and ask which is more likely to offer people the highest utility. The answer helps provide a first illustration of the economic usefulness of money.

3.1 Feasible Allocations

Imagine for a moment that we are central planners with complete knowledge of and total control over the economy. We studied this case in Chapter 1 and remind you what the Golden Rule allocation is for comparisons that are useful in this chapter. The Golden Rule allocation is represented in Figure 2.1 for a case in which population is constant over time.

The feasible allocation that a central planner selects depends on the objective. One reasonable and benevolent objective is the maximization of the utility of future generations, an objective we call the "golden rule." The golden rule in Figure 2.1 is represented by point E, which offers each person the consumption bundle (c_1^*, c_2^*). This combination of c_1 and c_2 yields the highest feasible level of utility during a person's entire lifetime.

What if the planner's objective was to maximize the consumption by the initial old? With this objective, the planner would allocate y units of the consumption good to the old and nothing for the consumption by the young. This stationary allocation would not maximize the utility of the future generations. All two-period-lived people would prefer the more balanced combination of consumption when young and old, represented by (c_1^*, c_2^*). Faced with this conflict in the interests of the initial old and future generations, an economist cannot choose among them on purely objective grounds. Nevertheless, the reader will find that, on subjective grounds (influenced by the fact that there are an infinite number of future generations and only a single generation of initial old), we tend to pay particular attention to the golden rule in this book.

4 Decentralized Solutions

In the previous section, we found the feasible allocation that maximizes the utility of the future generations. However, to achieve this allocation, in each period the central planner would have to take away c_2^* from each young person and give this amount to each old person. Such redistribution requires that the central planner have the ability to reallocate endowments costlessly between the generations. Furthermore, to determine c_1^* and c_2^*, this central planner also must know the exact utility function of the subjects.

These are strong assumptions about the power and wisdom of central planners. This leads us to ask if there is some way we can achieve this optimal allocation in a more decentralized manner, one in which economy reaches the optimal allocation through mutually beneficial trades conducted by each person. In other words, can we let a market do the work of the central planner?

We consider an economy in which perfect record keeping is not feasible. So we apply the notion of a competitive equilibrium to such an economy. Before offering money as way to achieve the optimal allocation, we consider the equilibrium outcome without any store of value.

4.1 Equilibrium without Money

Let us consider the nature of the competitive equilibrium when there is no money in our economy of overlapping generations. Recall that agents are endowed with some of the consumption good when young. Their endowment is zero when old. Their utility can be increased if they give up some of their endowment when they are young in exchange for some of the goods when they are old. Without the presence of an all-powerful central planner, we must ask ourselves if there are trades between people in the economy that could achieve this result.

No such trades are possible. Remember in Figure 1.1, there is a pattern of endowments described by young having goods and old people having nothing. A young person at period t has two types of people with whom to trade potentially in period t – other young people of the same generation or old people of the previous generation. However, trade with fellow young people would be of no benefit to the young person under consideration. They, like him, have none of the consumption good when they are old. Trade with the old would also be fruitless; the old want the good the young have, but they do not have what the young want (because they will not be alive in the next period). The source of the consumption good at time $t + 1$ is from the people who are born in that period. However, in period t, these people have not yet come into the world and so do not want what young people have to trade. This lack of possible trades is the manner in which the basic overlapping generations model captures the "absences of double coincidence of wants" (a term introduced by the 19th century economist W. S. Jevons [1875] to explain the need for money). Each generation wants what the next generation has but does not have what the next generation wants.

The resulting equilibrium is "autarkic" – people have no economic interaction with others. Unable to make mutually beneficial trades, each person consumes his entire endowment when young and nothing when old. In this autarkic equilibrium, utility is low. Both the future generations and the initial old are worse off than they would be with almost any other feasible consumption bundle. A member of the future generations would gladly give up some of his endowment when young in order to consume something when old. A member of the initial old would also like to consume something when old.

Figure 2.2 depicts the autarkic equilibrium in an economy with no record keeping and no money. In autarky, all the consumption is done when young and nothing is consumed by old people. It is easy to see how the record-keeping friction has bite when we see the autarkic equilibrium. When young, no one will unilaterally give up goods to an old person. Every old person will say that they gave goods when they were young. Social amnesia, however, keeps a young person from verifying the old person did participate. Such one-time meetings, therefore, result in no trade since a young person cannot establish a record of giving when young and will not trade with when they are old.

4.2 Equilibrium with Money

To open up a trading opportunity that might permit an exit from this grim autarkic equilibrium, we now introduce fiat money into our simple economy. "Fiat money"

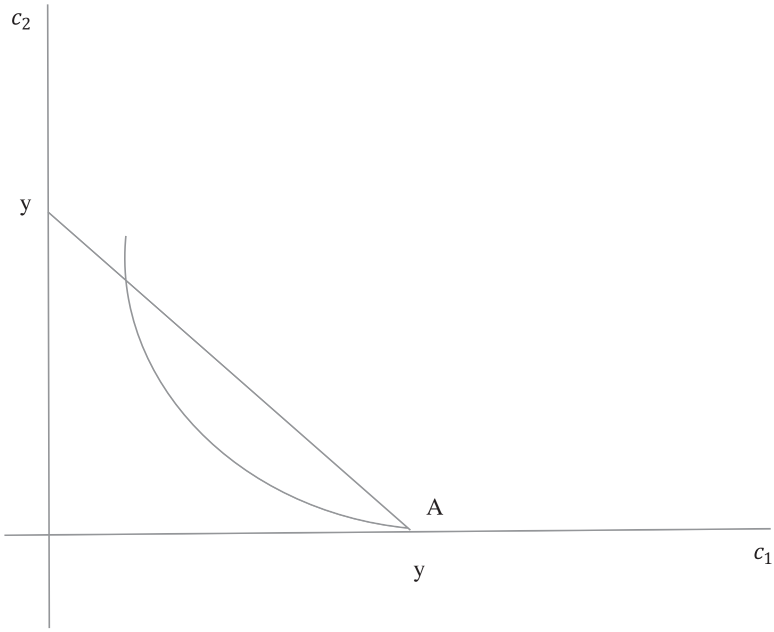

Figure 2.2. The autarkic equilibrium. It is located at a point where the indifference curve touches the budget line on the c_1-*axis* (Bundle A). The autarkic equilibrium results in all the consumption when young and nothing when old.

is a nearly costlessly produced commodity that cannot itself be used in consumption or production and is not a promise for anything that can be used in consumption or production.

For the purposes of our model, we assume the government can produce fiat money costlessly but that it cannot be produced or counterfeited by anyone else. Fiat money can be costlessly stored (held) from one period to the next and is costless to exchange. Pieces of paper distinctively marked by the government serve as fiat money.

Because people derive no direct utility from holding or consuming money, fiat money is valuable only if it enables people to trade for something they want to consume.

A "monetary equilibrium" is a competitive equilibrium in which there is a valued supply of fiat money. By valued, we mean that the fiat money can be traded for some of the consumption good. For fiat money to have value, its supply must be limited, and it must be impossible (or very costly) to counterfeit. Obviously, if everyone has the ability to print money costlessly, its supply will rapidly approach infinity, driving the value of any one unit to zero.

We began our analysis of monetary economies with an economy with a fixed stock of M perfectly divisible units of fiat money. We assume that each of the initial old begins with an equal number, M/N, of these units.

The presence of fiat money opens up a trading possibility. A young person can sell some of his endowment of goods (to old persons) for fiat money, hold the money until the next period, and then trade the fiat money for goods (with the young of that period).

5 Finding the Demand for Fiat Money

Of course, this new trading possibility exists only if fiat money is valued – in other words, if people are willing to give up some of the consumption good in trade for fiat money and vice versa. Because fiat money is intrinsically useless, its value depends on one's view of its value in the future, when it will be exchanged for the goods that do increase a person's utility.

If it is believed that fiat money will not be valued in the next period, then fiat money will have no value in this period. No one will be willing to give up some of the consumption good in exchange for it. That would be tantamount to trading something for nothing.

Extending this logic, we can predict that fiat money will have no value today if it is known with complete certainty that fiat money will be valueless at any future date T. To see this, first ask what the value of fiat money will be at time $T-1$; in other words, ask how many goods you would be willing to give for money at $T-1$ if it is known that it will be worthless at time T. The answer, of course, is that you would not be willing to give up any goods at time $T-1$ for money. In other words, fiat money would have no value at time $T-1$. Then what must its value be at time $T-2$? By similar reasoning, we see that it will also be valueless at time $T-2$. Working backward in this manner, we can see that fiat money will have no value today if it will be valueless at some point in the future.

Now let us consider a more interesting equilibrium in which money has a positive value in all future periods. We define v_t as the value of 1 unit of fiat money (let us call the unit a dollar) in terms of goods; that is, it is the number of goods that one must give up to obtain one dollar. It is the inverse of the dollar price of the consumption good, which we write as p_t. For example, if a banana costs 20 cents, $p_t = 1/5$ dollars and the value of a dollar, v_t, is five bananas. Note also that because our economy has only one good, the price of that good p_t can be viewed as the price level in this economy.

5.1 A Person's Budget

Let us now examine how people will decide how much money to acquire (assuming that fiat money will have a positive value in the future). To answer, we must first establish the constraints on the choices of the person – why he cannot simply enjoy infinite consumption both when young and when old. As was the case for the entire society, the constraints on a person are that he cannot give up more goods than he has. We will refer to the limitations on a person's consumption as his "budget constraints."

In the first period of life, a person has an endowment of y goods. The person can do two things with these goods – consume them and/or sell them for money. Notice that no one in the future generations is born with fiat money. To acquire fiat money, a person must trade. If the number of dollars acquired by a person (by giving up some of the consumption good) at time t is denoted by m_t, then the total number of goods sold for money is $v_t m_t$. We can therefore write the budget constraint facing each person in the first period of life as

$$c_{1,t} + v_t m_t \leq y. \tag{2.1}$$

The left-hand side of Equation 2.1 is the person's total uses of goods (consumption and acquisition of money). The right-hand side of Equation 2.1 represents the total sources of goods (the person's endowment).

In the second period of life, the person receives no endowment. Hence, when old, a person can acquire goods for consumption only by spending the money acquired in the previous period. In the second period of life (period $t + 1$), this money will purchase $v_{t+1} m_t$ units of the consumption good. The only use for these goods is second-period consumption. This means that the constraint facing the person in the second period of life is

$$c_{2,t+1} \leq v_{t+1} m_t. \tag{2.2}$$

In a monetary equilibrium in which, by definition, $v_t > 0$ for all t, we can rewrite this constraint as $m_t \geq (c_{2,t+1})/(v_{t+1})$ and substitute it into the first-period constraint (Equation 2.1) to obtain

$$c_{1,t} + \frac{v_t c_{2,t+1}}{v_{t+1}} \leq y \tag{2.3}$$

or

$$c_{1,t} + \left[\frac{v_t}{v_{t+1}}\right] c_{2,t+1} \leq y. \tag{2.4}$$

Equation 2.4 expresses the various combinations of first- and second-period consumption that a person can afford over a lifetime. In other words, it is the person's "lifetime budget constraint."

We can graph this budget constraint as shown in Figure 2.2. We can easily verify that the intercepts of the budget line are as illustrated. The budget line represents Equation 2.4 at equality. If nothing is consumed in the second period of life ($c_{2,t+1} = 0$), then the constraint implies that $c_{1,t} = y$. This is the horizontal intercept of the budget line. On the other hand, if nothing is consumed in the first period of life ($c_{1,t} = 0$), so that the entire endowment of y is used to purchase money, the constraint implies that $[(v_t)/(v_{t+1})]c_{2,t+1} = y$ or $c_{2,t+1} = [(v_{t+1})/(v_t)]y$. This represents the vertical intercept of the budget line.

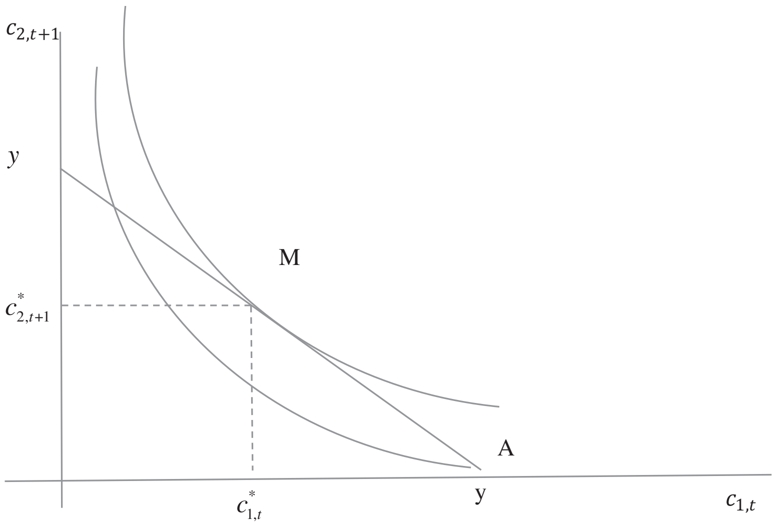

Figure 2.3. The choice of consumption with perfect record keeping. At Bundle M, a person maximizes utility given their lifetime budget constraint in a monetary equilibrium. Bundle M is found by locating a point of tangency between an indifference curve and the person's lifetime budget set line. The rate of return on money determines the slope of the budget set line.

Note that $(v_{t+1})/(v_t)$ can be considered as the "(real) rate of return of fiat money" because it expresses how many goods can be obtained in period $t+1$ if one unit of the good is sold for money in period t.

For a given rate of return of money, $(v_{t+1})/(v_t)$, we can find the $(c^*_{1,t}, c^*_{2,t+1})$ combination that will be chosen by a person who is seeking to maximize their utility. This point is shown in Figure 2.3. It is the point along the budget line that touches the highest indifference curve. This must occur at a point where the budget line is tangent to an indifference curve.

5.2 Finding Fiat Money's Rate of Return

But how can we determine the rate of return on intrinsically useless fiat money? The value that a person places on a unit of fiat money at time t, v_t, depends on what that person believes will be the value of one unit of money at $t+1$, v_{t+1}. By similar logic, the value of a unit of fiat money at time $t+1$ depends on a person's beliefs about the value of money in period $t+2$, v_{t+2}. And so on. We see that the value of fiat money at any point in time depends on an infinite chain of expectations about its future values. This indefiniteness is not due to any peculiarity in our model but rather to the nature of fiat money, which, because it has no intrinsic value, has a value that is determined by views about the future.

Whatever the views of the future value of money, a reasonable benchmark is the case in which these views are the same for every generation. This is plausible because in our basic model, every generation faces the same problem; endowments, preferences, and population are the same for every generation. If views about the future are also the same across generations, then people will react in the same manner in each period, choosing $c_{1,t} = c_1$ and $c_{2,t} = c_2$ for each period t. We call such equilibria "stationary equilibria." Notice that because each person faces different circumstances, depending on whether they are young or old, c_1 will not in general be equal to c_2 in a stationary equilibrium. People may choose to consume more when young or more when old. It turns out that the relative mix of first- and second-period consumption depends on preferences and on the rate of return of fiat money.

We also assume that people in our economy form their expectations of the future rationally. In this nonrandom economy, where there are no surprises, "rational expectations" means that a person's expectations of future variables equal the actual values of these future variables. In this special case, we say that people have perfect foresight. With perfect foresight, there are no errors in people's forecast of the important economic variables that affect their decisions. In the context of our model, this assumption means that a person born in period t will perfectly forecast the value of money in the next period, v_{t+1}. The person's expectation of this value will be exactly realized. This assumption would be less credible in an economy buffeted by random shocks than in our model economy, where preferences and the environment are unchanging and therefore are perfectly predictable.

To see the importance of perfect foresight, consider the alternative in a nonrandom economy – that people always expect a value of money greater or less than the value of money that actually occurs. A person with wrong beliefs about the future value of money will not choose the money balances that maximize their utility. They therefore have an incentive to figure out the value of money that actually will occur.

Let us now employ the assumptions of stationarity and perfect foresight to find an equilibrium time path of the value of money. In perfectly competitive markets, the price (or value) of an object is determined as the price at which the supply of the object equals its demand. This applies to the determination of the price (value) of money as well as the price of any good.

The demand for fiat money of each person is the number of goods each chooses to sell for fiat money, which equals the goods of the endowment that the person does not consume when young, $y - c_{1,t}$. The total money demand by all people in the economy at time t is therefore $N_t(y - c_{1,t})$.

The total supply of fiat money, measured in units of the consumption goods, is $v_t M_t$, implying that the total supply of fiat money measured in goods is the number of dollars multiplied by the value of each dollar, or $v_t M_t$. Equality of supply and demand therefore requires that

$$v_t M_t = N_t(y - c_{1,t}). \tag{2.5}$$