Macroeconomics

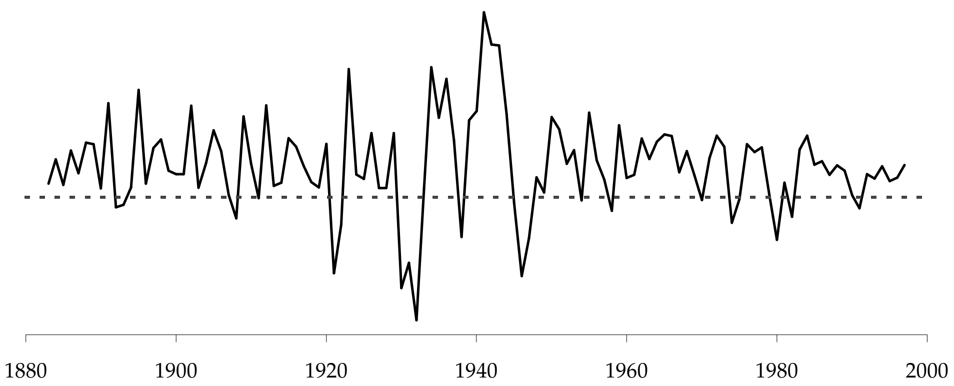

Matthias Doepke
University of Chicago

Andreas Lehnert
Board of Governors of the
Federal Reserve System

Andrew W. Sellgren
George Mason University

This book was typeset in Palatino and Computer Modern using LaTeX.

Draft of 1 September 1999.

Preface

We have designed this book to be a supplement to Robert J. Barro's *Macroeconomics*, which is the textbook that is used in introductory macroeconomics courses at the University of Chicago. In teaching these courses, we have found that Barro's treatment of the subject does not make use of the mathematical skills of our students. In particular, Barro relies almost exclusively on economic intuition and graphs to elucidate his subject. Since our students are familiar with calculus, we are able to work out formal models. This almost always allows greater concreteness and concision.

We have attempted to align our chapters with those in Barro's textbook. Sometimes our chapters present mathematical versions of the models that Barro introduces in his corresponding chapters (as in Chapters 2 and 19). Other times, our chapters contain material that extends his work (as in Chapters 5 and 17). Throughout, we have tried to add value to the treatment in Barro's book and to minimize redundancy. For example, we have nothing to add to Barro's Chapters 7, 16, and 20, so we have not covered those chapters. Three chapters deviate from this plan. Chapter 1 develops the mathematics of interest rates and growth rates; Barro does not cover these topics, but they are behind the scenes in his Chapter 1 and throughout his book. Chapter 10, which covers unemployment, is completely unrelated to Barro's Chapter 10. It is intended as a companion to the book *Job Creation and Destruction* by Davis, Haltiwanger, and Schuh. Chapter 18 covers the relationship between the government budget constraint and inflation along the lines of the "Unpleasant Monetarist Arithmetic" of Sargent and Wallace. Although Barro has a sidebar on this topic in his Chapter 14, we feel that it is important enough to merit a chapter of its own. We chose Chapter 18 since it is a natural point between fiscal policy (Chapters 12, 13, and 14) and monetary policy (Chapter 19). Barro's Chapter 18 is a review of the empirical evidence on the effect of monetary shocks on the real economy, and is well worth covering.

There are exercises after each chapter, and we have provided complete solutions at the end of this book. We believe that exercises are essential for students to learn this material. They give students a sense of what they ought to know, since these exercises have been drawn from several years of exams. Also, we often use exercises to introduce extensions to the material in the text. We have attempted to estimate the difficulty of these exercises, labeling them as "Easy," "Moderate", or "Hard". An exercise with a "Hard" rating may require a lot of algebra, or it may use unfamiliar concepts. Most other questions are rated as "Moderate", unless they have one-line solutions, in which case we usually rated them as "Easy".

We teach this material in two ten-week courses. In the first course we cover Chapters 1, 2, 3, 6, 4, 5, 7, 8, 9, and 11, in that order. This allows us to keep together all the material on monetary economics (Chapters 4, 5, 7, and 8). In the second course, we cover Chapter 10 (unemployment); Chapters 12, 13, and 14 (fiscal policy); Chapters 15 and 16 (international macro); and Chapters 17, 18 and 19 (money and banking). Since this is quite a lot to cover in ten weeks, instructors of the second course have traditionally touched only briefly on unemployment and international macro and concentrated instead on monetary and fiscal policy. The second course can benefit substantially from outside readings, such as: *Rational Expectations and Inflation* by Thomas Sargent; *A Monetary History of the United States* by Milton Friedman and Anna Schwartz; and *Job Creation and Destruction* by Davis, Haltiwanger, and Schuh.

This book would not have been possible without the support of the Department of Economics at the University of Chicago and the encouragement of Grace Tsiang. We would also like to thank the many students and faculty who have helped us to develop this material. A number of exercises in the first half of the book were based on questions written by Robert E. Lucas, Jr. The material in the second half of this book has benefited from several generations of instructors of Economics 203. In particular, Alexander Reyfman wrote a series of lectures which were the genesis of Chapters 12 through 19. Reyfman's teaching assistant Bill Dupor, and Lehnert's teaching assistants Jerry Cubbin and Tom Miles, all contributed valuable suggestions. During Cubbin's tenure as TA, he wrote most of the solutions to the problem sets, and several of these have found their way into this book. All students subjected to early drafts of this material contributed to the book's current form; Shannon Thaden, Ben Ruff, and Calvin Chan deserve special mention.

In spite of all the comments and suggestions we have received, this book inevitably contains errors and omissions. We would be grateful if you would bring these to our attention. The authors can be reached by e-mail at:

`m-doepke@uchicago.edu`

There is also a tear-out feedback form at the end of the book, along with a tear-out midterm-evaluation form for Economics 202 and 203.

Finally, some of the material in this book involves policy prescriptions. At some level, policy is a matter of opinion. The opinions expressed herein are not necessarily those of the Board of Governors of the Federal Reserve System.

Chicago, Illinois
September 1999

Contents

Chapter 1

Preliminaries

This chapter introduces interest rates and growth rates. The two topics are closely related, so we treat them together. The concepts discussed here are not in Barro, but they will help you understand the graphs and statistics that he uses throughout his book.

1.1 Compound Interest

We begin with some common terms and calculations from the realm of fixed-income investments. The amount of the investment is called the *principal*. The "fixed-income" from the investments is called *interest*. The interest per unit of principal per unit of time is called the *interest rate*. Most commonly, interest rates are quoted in dollars per year per dollar of principal. These units can be written: $\$/(y\$)$. The dollar units cancel, so this interest rate has units of one over years. Similarly, if the interest rate is apples per day per apple borrowed, the apple units will cancel, and the units of the interest rate will be one over days. In general, the units of an interest rate are one over some unit of time.

When the unit of time is a year, we say that an interest rate is an *annual* interest rate. If the unit of time is not mentioned, then it will almost always be an annual interest rate. Interest rates that are quoted in some specific unit of time can be converted to any other unit of time via a simple linear transformation. For example, a daily interest rate of x% corresponds to an annual interest rate of (365)(x)%.[1] (See Exercise 1.1 for an example.)

We use P for the principal of a fixed-income investment and R for the annual interest rate. Under *simple interest* the interest is earned on the amount of the principal only. In this case,

[1] You may be wondering about leap years. These are handled according to any of a number of conventions. For example, some interest rates are quoted using 360 days as a year; others use 365; still others use 365.25.

after n years the value of the investment will be:

$$V_s(n) = RPn + P. \tag{1.1}$$

For example, suppose you invest $5,000 at a 4.5% simple annual interest rate. After two years the value of your investment will be:

$$V_s(2) = (0.045)(\$5,000)(2) + \$5,000 = \$5,450.$$

It is much more common for interest to be compounded annually. In this case, at the end of each year, that year's interest will be added to the principal, so the investment will earn interest on the interest. The first year will be just like simple interest, since none of the interest will yet be compounded. Accordingly, the the value after the first year will be:

$$V_a(1) = RP + P = (1 + R)P.$$

After the second year, the value will be:

$$V_a(2) = RV_a(1) + V_a(1) = R(1 + R)P + (1 + R)P = (1 + R)^2P.$$

Similarly, after n years, the value will be:

$$V_a(n) = (1 + R)^n P. \tag{1.2}$$

Of course, this formula works only an integral numbers of years. For non-integral numbers, you round down to the nearest integral year n, compute $V_a(n)$, and use that in the simple-interest formula (1.1) for the fraction of the last year. (See Exercise 1.6 for an example.)

Let's revisit our previous example. Once again, you invest $5,000 at a 4.5% annual interest rate, but this time interest compounds annually. After two years the value of your investment will be:

$$V_a(2) = (1 + 0.045)^2(\$5,000) = \$5,460.13.$$

(Here and throughout, dollar amounts are rounded to the nearest cent.) Notice that the investment is worth less under simple interest than under compound interest, since under compounding you earn about $10 of interest on the first year's interest.

The above reasoning for compounding annually applies to compounding more frequently. The only catch is that the interest rate needs to be quoted in terms of the same time interval as the compounding. If R is an annual interest rate, and interest is to compound t times per year, then the value of an investment after n years will be:

$$V_t(n) = \left(1 + \frac{R}{t}\right)^{tn} P.$$

We return to our example again, this time supposing that interest compounds daily. After two years, the value will be:

$$V_{365}(2) = \left(1 + \frac{0.045}{365}\right)^{(365)(2)} (\$5,000) = \$5,470.84.$$

As we compound more and more frequently, we arrive at the expression for continuous compounding:

$$V_c(n) = \lim_{t\to\infty} \left(1 + \frac{R}{t}\right)^{tn} P.$$

We can make this much more tractable by using the fact that:

$$e = \lim_{x\to\infty} \left(1 + \frac{1}{x}\right)^{x},$$

where e is Euler's constant. This gives us the following formula for continuous discounting:

$$V_c(n) = \lim_{t\to\infty} \left(1 + \frac{R}{t}\right)^{tn} P = \left[\lim_{(t/R)\to\infty} \left(1 + \frac{1}{(t/R)}\right)^{(t/R)}\right]^{Rn} P = e^{Rn}P. \tag{1.3}$$

We return to our example one last time, this time assuming continuous compounding. After two years, the value of the investment will be:

$$V_c(2) = e^{(0.045)(2)}(\$5,000) = \$5,470.87.$$

Again, notice how throughout these examples the value of the investment is greater the more often the interest compounds. Continuous compounding results in the highest value, but the returns to more-frequent compounding fall off fairly quickly. For example, the value is almost the same under daily versus continuous discounting.

1.2 Growth Rates

Economists are often interested in the *growth rates* of economic variables. You might read, "Real Gross Domestic Product grew at a 2.3% annual rate this quarter" or "Inflation is 4%" or "The world's population is growing 20% every decade." Each of these statements deals with a growth rate.

An interest rate is just the growth rate of the value of an asset, and all the terminology and formulae from the previous section apply to growth rates generally. For example, we can calculate simple annual growth rates and annual growth rates that are compounded annually or continuously.

Consider the following values for the Gross Domestic Product (GDP) of a hypothetical country:

Year	GDP
1991	$100,000,000
1992	$130,000,000
1993	$135,000,000

The growth rate of GDP is just the interest rate that GDP would have had to earn if it were a fixed-income investment.

For example, the simple rate of growth of GDP between 1992 and 1993 is given by R in equation (1.1). Starting GDP is P, ending GDP is $V_s(n)$, and n is one year. Plugging all the numbers in, we get:

$$\$135K = (R)(\$130K)(1) + \$130K, \text{ so:}$$
$$R = \$135K/\$130K - 1 \approx 1.03846154 - 1 = 3.846154\%.$$

As another example, to calculate the annual rate of growth of GDP, compounded annually, between 1991 and 1993, we use equation (1.2). Starting GDP is P, ending GDP is $V_a(n)$, and n is two years. This gives us:

$$\$135K = (1 + R)^2(\$100K), \text{ so:}$$
$$R = (\$135K/\$100K)^{(0.5)} - 1 \approx 1.16189500 - 1 = 16.189500\%.$$

As a final example, we do the same calculation, but using continuous compounding. We just solve equation (1.3) for R. Starting GDP is P, ending GDP is $V_c(n)$, and n is two years.

$$\$135K = e^{2R}(\$100K), \text{ so:}$$
$$R = [\ln(\$135K) - \ln(\$100K)](0.5) \approx 0.15005230 = 15.15005230\%.$$

Economists generally prefer to use continuous compounding, for two reasons. First, under continuous compounding, computing the growth rate between two values of a series requires nothing more than taking the difference of their natural logarithms, as above.

This property is useful when graphing series. For example, consider some series that is given by $V(n) = V_0 e^{0.08n}$, which is depicted in Figure 1.1. By the equations above, we know that this series grows at an 8% continuous rate. Figure 1.2 depicts the natural logarithm of the same series, i.e., $\ln[V(n)] = \ln(V_0) + 0.08n$. From the equation, you can see that this new series is linear in n, and the slope (0.08) gives the growth rate. Whenever Barro labels the vertical axis of a graph with "Proportionate scale", he has graphed the natural logarithm of the underlying series. For an example, see Barro's Figure 1.1.

The second reason economists prefer continuous growth rates is that they have the following desirable property: if you compute the year-by-year continuous growth rates of a series and then take the average of those rates, the result is equal to the continuous growth rate over the entire interval.

For example, consider the hypothetical GDP numbers from above: $\$100K$, $\$130K$, and $\$135K$. The continuous growth rate between the first two is: $\ln(\$130K) - \ln(\$100K)$. The continuous growth rate between the second two is: $\ln(\$135K) - \ln(\$130K)$. The average of these two is:

$$\frac{\left[\ln(\$135K) - \ln(\$130K)\right] + \left[\ln(\$130K) - \ln(\$100K)\right]}{2}.$$

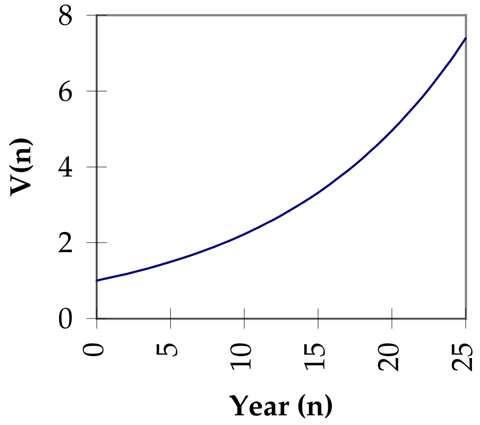

Figure 1.1: Exponential Growth

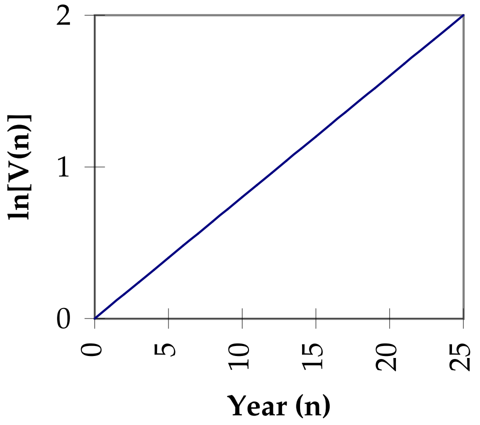

Figure 1.2: Log of Exp Growth

The two ln(\$130K) terms cancel, leaving exactly the formula for the continuous growth rate between the first and third values, as we derived above.

If we carry out the same exercise under simple growth or annually compounded growth, we will find that the average of the individual growth rates will not equal the overall growth rate. For example, if GDP grows by 8% this year and 4% next year, both calculated using annual compounding, then the two-year growth rate will *not* be 6%. (You should verify that it will actually be 5.98%.) On the other hand, if the 8% and 4% numbers were calculated using continuous compounding, then the continuous growth rate over the two-year period would be 6%.

Exercises

Exercise 1.1 (Easy)
My credit card has an APR (annualized percentage rate) of 16.8%. What is the daily interest rate?

Exercise 1.2 (Easy)
My loan shark is asking for \$25 in interest for a one-week loan of \$1,000. What is that, as an annual interest rate? (Use 52 weeks per year.)

Exercise 1.3 (Moderate)
The Consumer Price Index (CPI) is a measure of the prices of goods that people buy. Bigger numbers for the index mean that things are more expensive. Here are the CPI numbers for four months of 1996 and 1997:

Variable	Definition
P	Principal (amount invested)
R	Nominal interest rate
n	Number of years invested
$V_s(n)$	Value after n years under simple interest
$V_a(n)$	Value after n years under annual compounding
$V_t(n)$	Value after n years when compounded t times per year
$V_c(n)$	Value after n years under continuous compounding
V_0	Initial value of the investment

Table 1.1: Notation for Chapter 1

Year	Mar	Jun	Sep	Dec
1996	155.7	156.7	157.8	158.6
1997	160.0	160.3	161.2	161.3

What is the growth rate of the CPI between June 1996 and September 1996? (Use a continuous growth rate and annualize your answer.)

Exercise 1.4 (Moderate)
Use the CPI data from the previous exercise to compute the growth rates in the CPI in the four quarters starting in March 1996 (i.e, Mar-Jun 1996, Jun-Sep 1996, etc.). (Use a continuous growth rate but do not annualize your answer.) Show that the sum of these four rates equals the (continuous) growth rate from March 1996 to March 1997.

Exercise 1.5 (Easy)
Real output of the United States will likely grow by about 2% over the first half of the next century. At that rate (of continuous growth), how long will it take for real output to double? Compare your exact answer with the approximation given by the "Rule of 72."[2]

Exercise 1.6 (Hard)
This morning you invest $10,000 at 6.5% interest that compounds annually. What is the first date on which you would have at least $15,000? (Quote the answer in terms of years + days from today. Interest accrues each night, but compounds only annually.)

Exercise 1.7 (Easy)
Suppose that 4.6 percent of the earth's forests are cleared each year. How long will it take

[2]The "Rule of 72" is as follows. If the interest rate on an investment is x percent, then it takes about $72/x$ years for the value of the investment to double.

for half our current forests to be cleared? (Use annual compounding and solve for the fewest number of whole years.)

Exercise 1.8 (Moderate)
World population was about 679 million in the year 1700 and about 954 million in 1800.

1. What was the annual growth rate of population between 1700 and 1800? (Use continuous compounding.)
2. Suppose that the human race began with Adam and Eve and that the annual growth rate between 1700 and 1800 prevailed in all years prior to 1700. About when must it have been that Adam and Eve were evicted from the Garden of Eden? (Hint: What was the population in that year?)

Exercise 1.9 (Moderate)
According to figures compiled by the World Bank, per capita real income in the U.S. was $15,400 in 1984, while the corresponding figure for Japan was $10,600. Between 1965 and 1984, per capita real income in the U.S. grew at an annual rate of 1.7 percent (using annual compounding), while the corresponding figure for Japan was 4.7 percent.

1. If these two growth rates remain constant at their 1965-84 levels, in what year will per capita real income be the same in these two countries? (Again, use annual compounding, and use hundredths of a year.)
2. What will be the common per capita real income of these two countries at that date?

Chapter 2

Work Effort, Production, and Consumption

Robinson Crusoe is alone on an island, so he is an economy unto himself. He has preferences over consumption and leisure and can produce consumption goods by using labor and capital. We examine production first. Then we turn to preferences. Putting these two pieces together yields Crusoe's optimal choices of labor, leisure, and consumption.

2.1 Crusoe's Production Possibilities

Crusoe uses factors of production in order to make output y. We can think of this output as being coconuts. Two common factors of production, and those we consider here, are capital k and labor l. Capital might be coconut trees, and labor is the amount of time Crusoe works, measured as a fraction of a day. How much Crusoe produces with given resources depends on the type of technology A that he employs. We formalize this production process via a production function.

We often simplify our problems by assuming that the production function takes some particular functional form. As a first step, we often assume that it can be written: $y = Af(k, l)$, for some function $f(\cdot)$. This means that as technology A increases, Crusoe can get more output for any given inputs. It is reasonable to require the function $f(\cdot)$ to be increasing in each argument. This implies that increasing either input k or l will increase production. Another common assumption is that output is zero if either input is zero: $f(0, l) = 0$ and $f(k, 0) = 0$, for all k and l.

One functional form that has these properties is the Cobb–Douglas function, for example:

$y = Ak^{1-\alpha}l^{\alpha}$, for some α between zero and one. This particular Cobb-Douglas function exhibits constant returns to scale, since $(1-\alpha)+(\alpha)=1$. Figure 2.1 is a three-dimensional rendering of this function for particular values of A and α.

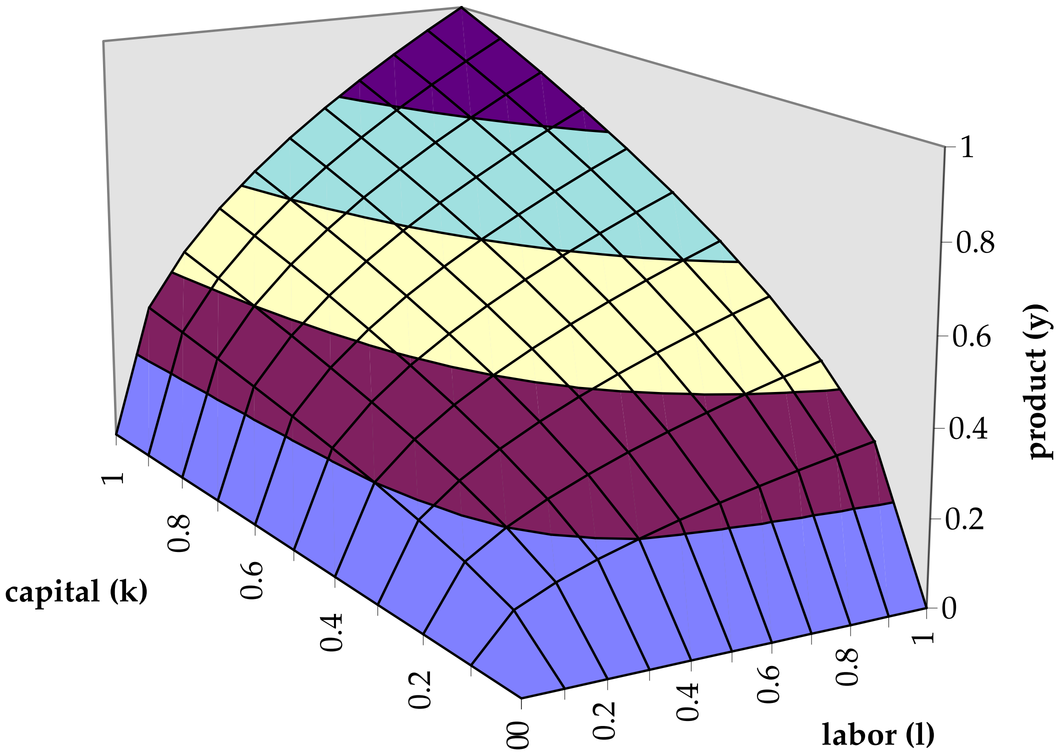

Figure 2.1: Cobb-Douglas Production

We will not be dealing with capital k until Chapter 9, so for now we assume that capital is fixed, say, at $k = 1$. This simplifies the production function. With a slight abuse of notation, we redefine $f(\cdot)$ and write production as $y = f(l)$. This is like what Barro uses in Chapter 2.

If the original production function was Cobb–Douglas, $y = Ak^{1-\alpha}l^{\alpha}$, then under $k = 1$ the production function becomes: $y = Al^{\alpha}$. The graph of this curve is just a slice through the surface depicted in Figure 2.1. It looks like Barro's Figure 2.1.

As you know, the marginal product of some factor of production (e.g., labor l) is the additional output, or "product", that results from increasing the input of that factor. Formally, the marginal product of an input is the derivative of the production function with respect to that input. For example, the marginal product of labor is: $dy/dl = f'(l)$.[1] Since the marginal product is the derivative of the the production function, and the derivative gives the slope, we can read the marginal product as the slope of the production function, as Barro does in his Figure 2.1.

In the particular case where production is Cobb–Douglas (and capital is fixed), the production function is: $y = Al^{\alpha}$, so the marginal product of labor is: $dy/dl = A\alpha l^{\alpha-1}$. This is always positive, as we require, and it decreases as we increase l. Accordingly, this production function exhibits diminishing marginal product: the first unit of labor is more productive than the tenth unit of labor. Graphing this marginal product equation gives us something like Barro's Figure 2.2.

Barro talks about improvements in technology and argues how both the production function and the marginal-product schedule shift as a result. The effects of such a change in technology are clearer when we examine a particular production function. For example, consider our production function: $y = Al^{\alpha}$. The improvement in technology means that A goes up. Accordingly, whatever production was before, it undergoes the same percentage increase as the increase in A. For example, if A doubles, then output at each l will be double what it used to be. Notably, when l is zero, output is zero just as before, since twice zero is still zero. The result is that the production function undergoes a kind of upward rotation, pivoting about the anchored origin, $l = 0$. That is precisely what Barro depicts in his Figure 2.3.

We can examine the marginal-product schedule as well. Under the particular functional form we are using, the marginal product of labor (MPL) is: $dy/dl = A\alpha l^{\alpha-1}$. Accordingly, the marginal product at each l undergoes the same percentage change as does A. Since the MPL is higher at low levels of l, the marginal-product curve shifts up more at those levels of l. Refer to Barro's Figure 2.4.

2.2 Crusoe's Preferences

Crusoe cares about his consumption c and his leisure. Since we are measuring labor l as the fraction of the day that Crusoe works, the remainder is leisure. Specifically, leisure is $1 - l$. We represent his preferences with a utility function $u(c, l)$. Take note, the second argument is not a "good" good, since Crusoe does not enjoy working. Accordingly, it might have been less confusing if Barro had written utility as $v(c, 1 - l)$, for some utility function $v(\cdot)$. We assume that Crusoe's preferences satisfy standard properties: they are increasing in each "good" good, they are convex, etc.

We will often simplify the analysis by assuming a particular functional form for Crusoe's

[1] Barro uses primes to denote shifted curves rather than derivatives. For example, when Barro shifts the $f(l)$ curve, he labels the new curve $f(l)'$. This is not a derivative. Barro's notation is unfortunate, but we are stuck with it.

preferences. For example, we might have: $u(c,l) = \ln(c) + \ln(1-l)$. With such a function in hand, we can trace out indifference curves. To do so, we set $u(c,l)$ to some fixed number $\bar{u}$, and solve for c as a function of l. Under these preferences, we get:

$$c = \frac{e^{\bar{u}}}{1-l}.$$

As we change $\bar{u}$, we get different indifference curves, and the set of those looks like Barro's Figure 2.6. These should look strange to you because they are increasing as we move to the right. This is because we are graphing a "bad" good (labor l) on the horizontal axis. If we graph leisure $(1-l)$ instead, then we will get indifference curves that look like what you saw in your microeconomics courses.

2.3 Crusoe's Choices

When we put preferences and technology together, we get Crusoe's optimal choices of labor l, leisure $1-l$, and consumption c. Formally, Crusoe's problem is:

$$\max_{c,l} \; u(c,l), \text{ such that:} \tag{2.1}$$

$$c \leq y, \text{ and:} \tag{2.2}$$

$$y = f(l). \tag{2.3}$$

There are two elements of equation (2.1). First, under the max, we indicate the variables that Crusoe gets to choose; in this case, he chooses c and l. Second, after the word "max" we place the maximand, which is the thing that Crusoe is trying to maximize; in this case, he cares about his utility.

Equation (2.2) says that Crusoe cannot consume more than he produces. We can use simple deduction to prove that we can replace the "$\leq$" symbol with "=". Suppose Crusoe chooses c and l such that $c < y$. This cannot be optimal because he could increase the maximand a little bit if he raised c, since $u(c,l)$ is increasing in c. Simply put: it will never be optimal for Crusoe to waste output y, so we know that $c = y$.

Finally, equation (2.3) simply codifies the production technology that is available to Crusoe.

With all this in mind, we can simplify the way we write Crusoe's problem as follows:

$$\max_{c,l} \; u(c,l), \text{ such that:}$$

$$c = f(l).$$

Here, we are making use of the fact that $c = y$, and we are substituting the second constraint into the first.

There are two principal ways to solve such a problem. The first is to substitute any constraints into the objective. The second is to use Lagrange multipliers. We consider these two methods in turn.

Substituting Constraints into the Objective

In the maximization problem we are considering, we have c in the objective, but we know that $c = f(l)$, so we can write the max problem as:

$$\max_{l} \; u[f(l), l].$$

We no longer have c in the maximand or in the constraints, so c is no longer a choice variable. Essentially, the $c = f(l)$ constraint tacks down c, so it is not a free choice. We exploit that fact when we substitute c out.

At this point, we have a problem of maximizing some function with respect to one variable, and we have no remaining constraints. To obtain the optimal choices, we take the derivative with respect to each choice variable, in this case l alone, and set that derivative equal to zero.[2] When we take a derivative and set it equal to zero, we call the resulting equation a *first-order condition*, which we often abbreviate as "FOC".

In our example, we get only one first-order condition:

$$\text{(FOC } l\text{)} \qquad \frac{d}{dl}\left\{u[f(l^\star), l^\star]\right\} = u_1[f(l^\star), l^\star]f'(l^\star) + u_2[f(l^\star), l^\star] = 0.$$

(See the Appendix for an explanation of the notation for calculus, and note how we had to use the chain rule for the first part.) We use $l^\star$ because the l that satisfies this equation will be Crusoe's optimal choice of labor.[3] We can then plug that choice back into $c = f(l)$ to get Crusoe's optimal consumption: $c^\star = f(l^\star)$. Obviously, his optimal choice of leisure will be $1 - l^\star$.

Under the particular functional forms for utility and consumption that we have been considering, we can get explicit answers for Crusoe's optimal choices. Recall, we have been using $u(c, l) = \ln(c) + \ln(1 - l)$ and $y = f(l) = Al^\alpha$. When we plug these functions into the first-order condition in equation (FOC l), we get:

$$\left(\frac{1}{A(l^\star)^\alpha}\right)\left(A\alpha(l^\star)^{\alpha-1}\right) + \frac{-1}{1 - l^\star} = 0. \tag{2.4}$$

[2] The reason we set the derivative equal to zero is as follows. The maximand is some hump-shaped object. The derivative of the maximand gives the slope of that hump at each point. At the top of the hump, the slope will be zero, so we are solving for the point at which the slope is zero.

[3] Strictly speaking, we also need to check the second-order condition in order to make sure that we have solved for a maximum instead of a minimum. In this text we will ignore second-order conditions because they will always be satisfied in the sorts of problems we will be doing.

The first term in parentheses is from $u_1(c, l) = 1/c$, using the fact that $c = Al^\alpha$. The second term in parentheses is from the chain rule; it is the $f'(l)$ term. The final term is $u_2(c, l)$. We can cancel terms in equation (2.4) and rearrange to get:

$$\frac{\alpha}{l^\star} = \frac{1}{1 - l^\star}.$$

Cross multiplying and solving yields:

$$l^\star = \frac{\alpha}{1 + \alpha}.$$

When we plug this value of $l^\star$ into $c^\star = f(l^\star)$, we get:

$$c^\star = A\left(\frac{\alpha}{1 + \alpha}\right)^\alpha.$$

These are Crusoe's optimal choices of labor and consumption.

Using Lagrange Multipliers

In many problems, the technique of substituting the constraints into the objective is the quickest and easiest method of carrying out the constrained maximization. However, sometimes it is difficult to solve the constraints for a particular variable. For example, suppose you have a constraint like:

$$c + \ln(c) = 10 + l + \ln(1 - l).$$

You cannot solve for either c or l, so the solution method described above is not applicable.

Accordingly, we describe how to use Lagrange multipliers to tackle problems of constrained maximization when it is either difficult or impossible to solve the constraints for individual variables. At first we treat the method as a cook-book recipe. After we are done, we will try to develop intuition for why the technique works.

Recall, we are working with the following problem:

$$\max_{c,l} \ u(c, l), \text{ such that:}$$
$$c = f(l).$$

The first step in using Lagrange multipliers is to solve the constraint so that everything is on one side, leaving a zero on the other side. In that regard, we have either:

$$f(l) - c = 0, \text{ or:}$$
$$c - f(l) = 0.$$

Either of those two will work, but we want to choose the first one, for reasons that are described below. The general heuristic is to choose the one that has a minus sign in front of the variable that makes the maximand larger. In this case, more c makes utility higher, so we want the equation with $-c$.

The second step is to write down a function called the *Lagrangean*, defined as follows:

$$\mathcal{L}(c, l, \lambda) = u(c, l) + \lambda[f(l) - c].$$

As you can see, the Lagrangean is defined to be the original objective, $u(c, l)$, plus some variable λ times our constraint. The Lagrangean is a function; in this case its arguments are the three variables c, l, and λ. Sometimes we will write it simply as $\mathcal{L}$, suppressing the arguments. The variable λ is called the *Lagrange multiplier*; it is just some number that we will calculate. If there is more than one constraint, then each one is: (i) solved for zero; (ii) multiplied by its own Lagrange multiplier, e.g., λ_1, λ_2, etc.; and (iii) added to the Lagrangean. (See Chapter 3 for an example.)

Before we used calculus to maximize our objective directly. Now, we work instead with the Lagrangean. The standard approach is to set to zero the derivatives of the Lagrangian with respect to the choice variables and the Lagrange multiplier λ. The relevant first-order conditions are:

$$\frac{\partial}{\partial c}[\mathcal{L}(c^\star, l^\star, \lambda^\star)] = u_1(c^\star, l^\star) + \lambda^\star[-1] = 0; \tag{FOC c}$$

$$\frac{\partial}{\partial l}[\mathcal{L}(c^\star, l^\star, \lambda^\star)] = u_2(c^\star, l^\star) + \lambda^\star[f'(l^\star)] = 0; \text{ and:} \tag{FOC l}$$

$$\frac{\partial}{\partial \lambda}[\mathcal{L}(c^\star, l^\star, \lambda^\star)] = f(l^\star) - c^\star = 0. \tag{FOC λ}$$

Again, we use starred variables in these first-order conditions to denote that it is only for the optimal values of c, l, and λ that these derivatives will be zero. Notice that differentiating the Lagrangian with respect to λ simply gives us back our budget equation. Now we have three equations in three unknowns ($c^\star$, $l^\star$, and $\lambda^\star$) and can solve for a solution. Typically, the first step is to use equations (FOC c) and (FOC l) to eliminate $\lambda^\star$. From (FOC c) we have:

$$u_1(c^\star, l^\star) = \lambda^\star,$$

and from (FOC l) we have:

$$-\frac{u_2(c^\star, l^\star)}{f'(l^\star)} = \lambda^\star.$$

Combining the two gives us:

$$u_1(c^\star, l^\star) = -\frac{u_2(c^\star, l^\star)}{f'(l^\star)}. \tag{2.5}$$

When we are given particular functional forms for $u(\cdot)$ and $f(\cdot)$, then equation (2.5) gives us a relationship between $c^\star$ and $l^\star$ that we can plug into the budget equation and solve further. For example, under $u(c, l) = \ln(c) + \ln(1 - l)$ and $f(l) = Al^\alpha$, equation (2.5) becomes:

$$\frac{1}{c^\star} = -\left(\frac{-1}{1 - l^\star}\right)\left(\frac{1}{A\alpha(l^\star)^{\alpha-1}}\right),$$

or equivalently:

$$c^\star = A\alpha(1 - l^\star)(l^\star)^{\alpha-1}.$$

Now we plug in the budget equation $c = Al^\alpha$ to get:

$$A(l^\star)^\alpha = A\alpha(1 - l^\star)(l^\star)^{\alpha-1}.$$

After some canceling and algebraic manipulation, this reduces to:

$$l^\star = \frac{\alpha}{1 + \alpha}.$$

Finally, we plug this answer for the optimal labor $l^\star$ back into the budget equation to get:

$$c^\star = A\left(\frac{\alpha}{1 + \alpha}\right)^\alpha.$$

Notice that these are the same answers for $c^\star$ and $l^\star$ that we derived in the previous subsection, when we plugged the constraint into the objective instead of using a Lagrange multiplier.

Now let's try to figure out why the technique of Lagrange multipliers works. First, we want to understand better what the Lagrange multiplier λ is. Our first-order condition with respect to c gave us:

$$u_1(c^\star, l^\star) = \lambda^\star, \tag{2.6}$$

This tells us that, at the optimum, $\lambda^\star$ is the marginal utility of an extra unit of consumption, given by the left-hand side. It is this interpretation of λ that motivated our choice of $f(l) - c = 0$ rather than $c - f(l) = 0$ when we attached the constraint term to the Lagrangean. If we had used the latter version of the constraint, then the right-hand side of equation (2.6) would have been $-\lambda$, which would have been minus the marginal utility of income.

Now look at the terms in the Lagrangian:

$$\mathcal{L}(c, l, \lambda) = u(c, l) + \lambda[f(l) - c].$$

It contains our objective $u(\cdot)$ and then the Lagrange multiplier times the constraint. Remember, λ is the marginal utility of an additional unit of consumption. Notice that if the budget equation is satisfied, then $f(l) = c$, so the constraint term is zero, and the Lagrangian $\mathcal{L}$ and the objective $u(\cdot)$ are equal. Ceteris paribus, the Lagrangian will be big whenever the objective is.

Now, think about the contributions from the constraint term. Suppose Crusoe is at some choice of c and l such that the budget is exactly met. If he wants to decrease labor l by a little bit, then he will have to cut back on his consumption c. The constraint term in the Lagrangean is: $\lambda[f(l) - c]$. The Lagrangean, our new objective, goes down by the required cut in c times λ, which is the marginal utility of consumption. Essentially, the Lagrangean subtracts off the utility cost of reducing consumption to make up for shortfalls in budget balance. That way, the Lagrangean is an objective that incorporates costs from failing to meet the constraint.

2.4 Income and Substitution Effects

Barro uses graphs to examine how Crusoe's optimal choices of consumption and labor change when his production function shifts and rotates. He calls the changes in Crusoe's choices "wealth and substitution effects". That discussion is vaguely reminiscent of your study of income and substitution effects from microeconomics. In that context, you considered shifts and rotations of linear budget lines. Crusoe's "budget line" is his production function, which is not linear.

This difference turns out to make mathematical calculation of income and substitution effects impractical. Furthermore, the "wealth effects" that Barro considers violate our assumption that production is zero when labor l is zero. Such a wealth effect is depicted as an upward shift of the production function in Barro's Figure 2.8. This corresponds to adding a constant to Crusoe's production function, which means that production is not zero when l is.

Barro's Figure 2.10 depicts a pivot of the production about the origin. This type of change to production is much more common in macroeconomics, since it is how we typically represent technological improvements. If Crusoe's production function is $y = Al^{\alpha}$, then an increase in A will look exactly like this. Given a specific functional form for $u(\cdot)$ as well, it is straightforward to compute how Crusoe's choices of consumption c and labor l change for any given change in A.

For example, suppose $u(c, l) = \ln(c) + \ln(1 - l)$ as before. Above we showed that:

$$c^{\star} = A\left(\frac{\alpha}{1+\alpha}\right)^{\alpha}.$$

Determining how $c^{\star}$ changes when A changes is called *comparative statics*. The typical exercise is to take the equation giving the optimal choice and to differentiate it with respect to the variable that is to change. In this case, we have an equation for Crusoe's optimal choice of $c^{\star}$, and we are interested in how that choice will change as A changes. That gives us:

$$\frac{\partial c^{\star}}{\partial A} = \left(\frac{\alpha}{1+\alpha}\right)^{\alpha}. \tag{2.7}$$

The derivative in equation (2.7) is positive, so Crusoe's optimal choice of consumption will increase when A increases.

The comparative statics exercise for Crusoe's optimal labor choice $l^\star$ is even easier. Above we derived:

$$l^\star = \frac{\alpha}{1+\alpha}.$$

There is no A on the right-hand side, so when we take the partial derivative with respect to A, the right-hand side is just a constant. Accordingly, $\partial l^\star/\partial A = 0$, i.e., Crusoe's choice of labor effort does not depend on his technology. This is precisely what Barro depicts in his Figure 2.10.

The intuition of this result is as follows. When A goes up, the marginal product of labor goes up, since the slope of the production function goes up. This encourages Crusoe to work harder. On the other hand, the increase in A means that for any l Crusoe has more output, so he is wealthier. As a result, Crusoe will try to consume more of any normal goods. To the extent that leisure $1 - l$ is a normal good, Crusoe will actually work less. Under these preferences and this production function, these two effects happen to cancel out precisely. In general, this will not be the case.

Variable	Definition
y	Income, in units of consumption
k	Capital
l	Labor, fraction of time spent on production
$f(l)$	Production function
α	A parameter of the production function
A	Technology of production
c	Consumption
$1-l$	Leisure, fraction of time spent recreating
$\mathcal{L}(\cdot)$	Lagrangean function
λ	Lagrange multiplier

Table 2.1: Notation for Chapter 2

Exercises

Exercise 2.1 (Easy)
An agent cares about consumption and leisure. Specifically, the agent's preferences are: $U = \ln(c) + \ln(l)$, where c is the agent's consumption, and l is the number of hours the agent

spends per day on leisure. When the agent isn't enjoying leisure time, the agent works, either for herself or for someone else. If she works n_s hours for herself, then she produces $y = 4n_s^{0.5}$ units of consumption. For each hour that she works for someone else, she gets paid a competitive wage w, in units of consumption.

Write out the agent's optimization problem.

Exercise 2.2 (Moderate)
Suppose Crusoe's preferences are given by: $u(c, l) = c^{\gamma}(1 - l)^{1-\gamma}$, for some γ between zero and one. His technology is: $y = f(l) = Al^{\alpha}$, just like before. Solve for Crusoe's optimal choices of consumption c and labor l. (You can use either substitution or a Lagrangean, but the former is easier in this sort of problem.)

Appendix: Calculus Notation

Suppose we have a function: $y = f(x)$. We can think of differentiation as an operator that acts on objects. Write $\frac{d}{dx}$ as the operator that differentiates with respect to x. We can apply the operator to both sides of any equation. Namely,

$$\frac{d}{dx}\big(y\big) = \frac{d}{dx}\big(f(x)\big).$$

We often write the left-hand side as $\frac{dy}{dx}$, and the right-hand side as $f'(x)$. These are just notational conventions.

When we have functions of more than one variable, we are in the realm of multivariate calculus and require more notation. Suppose we have $z = f(x, y)$. When we differentiate such a function, we will take *partial derivatives* that tell us the change in the function from changing only one of the arguments, while holding any other arguments fixed. Partial derivatives are denoted with curly dees (i.e., with ∂) to distinguish them from normal derivatives. We can think of partial differentiation as an operator as before:

$$\frac{\partial}{\partial x}\big(z\big) = \frac{\partial}{\partial x}\big(f(x, y)\big).$$

The left-hand side is often written as $\frac{\partial z}{\partial x}$, and the right-hand side as $f_1(x, y)$. The subscript 1 on f indicates a partial derivative with respect to the first argument of f. The derivative of f with respect to its second argument, y, can similarly be written: $f_2(x, y)$.

The things to remember about this are:

- Primes (f') and straight dees (df) are for functions of only one variable.
- Subscripts (f_1) and curly dees (∂f) are for functions of more than one variable.

Chapter 3

The Behavior of Households with Markets for Commodities and Credit

In this chapter we move from the world in which Robinson Crusoe is alone on his island to a world of many identical households that interact. To begin, we consider one particular representative household. When we add together the behaviors of many households, we get a macroeconomy.

Whereas in Chapter 2 we looked at Crusoe's choices between consumption and leisure at one point in time, now we consider households' choices of consumption over multiple periods, abstracting from the labor decisions of households. Section 3.1 introduces the basic setup of the chapter. In Section 3.2 we work out a model in which households live for only two periods. Households live indefinitely in the model presented in Section 3.3. Both these models follow Barro fairly closely, but of course in greater mathematical detail. The primary difference is that Barro has households carry around money, while we do not.

3.1 The General Setup

The representative household cares about consumption c_t in each period. This is formalized by some utility function $U(c_1, c_2, c_3, \dots)$. Economists almost always simplify intertemporal problems by assuming that preferences are additively separable. Such preferences look like: $U(c_1, c_2, c_3, \dots) = u(c_1) + \beta u(c_2) + \beta^2 u(c_3) + \cdots$. The $u(\cdot)$ function is called the *period utility*. It satisfies standard properties of utility functions. The variable β is called the *discount factor*. It is just a number, say 0.95. The fact that it is less than 1 means that the household cares a little more about current consumption than it cares about future consumption.

The household gets exogenous income y_t in each period. This income is in terms of consumption goods. We say that it is exogenous because it is independent of anything that the household does. Think of this income as some bequest from God or goods that fall from the sky.

At time t, the household can buy or sell consumption goods c_t at a price of P per unit. (As in Barro, the price level P does not change over time.) For example, if the household sells 4 units of consumption goods to someone else, then the seller receives $\$4P$ for those goods.

The household is able to save money by buying bonds that bear interest. We use b_t to denote the number of dollars of bonds that the household buys at period t, for which it will collect principal and interest in period $t + 1$. If the household invests \$1 this period, then next period it gets back its \$1 of principal plus $\$R$ in interest. Hence, if the household buys b_t in bonds this period, then next period the principal plus interest will be $b_t(1 + R)$. The household comes into the world with no bonds, i.e., $b_0 = 0$.

Since each \$1 of investment in bonds pays $\$R$ of interest, R is the simple rate of interest on the bonds. If the bonds pay R "next period", then whether the interest rate is daily, monthly, annual, etc., is determined by what the length of a "period" is. If the "period" is a year, then the interest rate R is an annual rate.

The household can either borrow or lend, i.e., the household can issue or buy bonds, whatever makes it happier. If b_t is negative, then the household is a net borrower.

At period t the household's resources include its income y_t and any bonds that it carries from last period, with interest. The dollar value of these resources is:

$$Py_t + b_{t-1}(1 + R).$$

At period t the household allocates its resources to its current consumption and to investment in bonds that it will carry forward to the next period. The dollar cost of these uses is:

$$Pc_t + b_t.$$

Putting these together gives us the household's period-t budget equation:

$$Py_t + b_{t-1}(1 + R) = Pc_t + b_t.$$

In a general setup, we would have one such budget equation for every period, and there could be arbitrarily many periods. For example, if a period were a year, and the household "lived" for 40 years, then we would have forty budget constraints. On the other hand, a period could be a day, and then we would have many more budget constraints.

3.2 A Two-Period Model

We begin this section with a discussion of the choices of a representative household. Then we put a bunch of these households together and discuss the resulting macroeconomic equilibrium.

Choices of the Representative Household

In this model the household lives for two time periods, $t = 1, 2$. In this case, the household's preferences reduce to:

$$U(c_1, c_2) = u(c_1) + \beta u(c_2). \tag{3.1}$$

Given that the household will not be around to enjoy consumption in period 3, we know that it will not be optimal for the household to buy any bonds in period 2, since those bonds would take away from period-2 consumption c_2 and provide income only in period 3, at which time the household will no longer be around. Accordingly, $b_2 = 0$. That leaves only b_1 in this model.

The household's budget constraints simplify as well. In period 1 the household's budget equation is:

$$Py_1 = Pc_1 + b_1, \tag{3.2}$$

and in period $t = 2$ it is:

$$Py_2 + b_1(1 + R) = Pc_2. \tag{3.3}$$

The household's problem is to choose consumptions c_1 and c_2 and first-period bond holdings b_1 so as to maximize utility (3.1) subject to the budget equations (3.2) and (3.3). The household takes the price level P and the interest rate R as given.

We write out the household's problem:

$$\max_{c_1, c_2, b_1} \{u(c_1) + \beta u(c_2)\}, \text{ subject to:} \tag{3.4}$$

$$Py_1 = Pc_1 + b_1, \text{ and:} \tag{3.5}$$

$$Py_2 + b_1(1 + R) = Pc_2. \tag{3.6}$$

We solve this problem by using the method of Lagrange multipliers. The Lagrangean is:

$$\mathcal{L} = u(c_1) + \beta u(c_2) + \lambda_1[Py_1 - Pc_1 - b_1] + \lambda_2[Py_2 + b_1(1 + R) - Pc_2],$$

where λ_1 and λ_2 are our two Lagrange multipliers. The first-order conditions are:

$$u'(c_1^\star) + \lambda_1^\star[-P] = 0; \tag{FOC c_1}$$

$$\beta u'(c_2^\star) + \lambda_2^\star[-P] = 0; \text{ and:} \tag{FOC c_2}$$

$$\lambda_1^\star[-1] + \lambda_2^\star[(1 + R)] = 0. \tag{FOC b_1}$$

(Again, stars denote that only the optimal choices will satisfy these first-order conditions.) We leave off the first-order conditions with respect to the Lagrange multipliers λ_1 and λ_2, since we know that they will give us back the two budget constraints.

Rewriting the first two FOCs gives us:

$$\frac{u'(c_1^\star)}{P} = \lambda_1^\star, \quad \text{and:} \quad \frac{\beta u'(c_2^\star)}{P} = \lambda_2^\star.$$

We can plug these into the FOC with respect to b_1 to get:

$$-\frac{u'(c_1^\star)}{P} + \frac{\beta u'(c_2^\star)}{P}(1+R) = 0,$$

which we can rewrite as:

$$\frac{u'(c_1^\star)}{u'(c_2^\star)} = \beta(1+R). \tag{3.7}$$

Equation (3.7) is called an *Euler equation* (pronounced: OIL-er). It relates the marginal utility of consumption in the two periods. Given a functional form for $u(\cdot)$, we can use this equation and the two budget equations to solve for the household's choices $c_1^\star$, $c_2^\star$, and $b_1^\star$.

It is possible to use the Euler equation to make deductions about these choices even without knowing the particular functional form of the period utility function $u(\cdot)$, but this analysis is much more tractable when the form of $u(\cdot)$ is given. Accordingly, we assume $u(c_t) = \ln(c_t)$. Then $u'(c_t) = 1/c_t$, and equation (3.7) becomes:

$$\frac{c_2^\star}{c_1^\star} = \beta(1+R). \tag{3.8}$$

Before we solve for $c_1^\star$, $c_2^\star$, and $b_1^\star$, let us think about this equation. Recall, preferences are: $u(c_1)+\beta u(c_2)$. Intuitively, if β goes up, then the household cares more about the future than it used to, so we expect the household to consume more c_2 and less c_1.

This is borne out graphically in Barro's Figure 3.4. Larger β corresponds to smaller slopes in the household's indifference curves, which rotate downward, counter-clockwise. Accordingly, the household's choice of c_2 will go up and that of c_1 will go down, like we expect.

We can show the result mathematically as well. An increase in β causes an increase in right-hand side of the Euler equation (3.8), so $c_2^\star$ goes up relative to $c_1^\star$, just like we expect.

Now we consider changes on the budget side. Suppose R goes up. Then the opportunity cost of consumption c_1 in the first period goes up, since the household can forego c_1 and earn a higher return on investing in bonds. By the same reasoning, the opportunity cost of c_2 goes down, since the household can forego less c_1 to get a given amount of c_2. Accordingly, if R goes up, we expect the household to substitute away from c_1 and toward c_2.

Refer to Barro's Figure 3.4. If R goes up, then the budget line rotates clockwise, i.e., it gets steeper. This indicates that the household chooses larger c_2 and smaller c_1 (subject to being on any given indifference curve), just like our intuition suggests.

Mathematically, we refer once again to the Euler equation. If R goes up, then the right-hand side is larger, so $c_2^\star/c_1^\star$ goes up, again confirming our intuition.

Given $u(c_t) = \ln(c_t)$, we can actually solve for the household's optimal choices. The Euler equation and equations (3.2) and (3.3) give us three equations in the three unknowns, $c_1^\star$, $c_2^\star$, and $b_1^\star$. Solving yields:

$$\begin{aligned} c_1^\star &= \frac{y_2 + y_1(1+R)}{(1+\beta)(1+R)}, \\ c_2^\star &= \big[y_2 + y_1(1+R)\big]\left[\frac{\beta}{1+\beta}\right], \text{ and:} \\ b_1^\star &= Py_1 - \frac{P[y_2 + y_1(1+R)]}{(1+\beta)(1+R)}. \end{aligned}$$

You can verify these if you like. Doing so is nothing more than an exercise in algebra.

If we tell the household what the interest rate R is, the household performs its own maximization to get its choices of c_1, c_2, and b_1, as above. We can write these choices as functions of R, i.e., $c_1^\star(R)$, $c_2^\star(R)$, and $b_1^\star(R)$, and we can ask what happens to these choices as the interest rate R changes. Again, this exercise is called "comparative statics". All we do is take the derivative of the choices with respect to R. For example:

$$\frac{\partial c_2^\star}{\partial R} = \frac{y_1\beta}{1+\beta} > 0,$$

so $c_2^\star$ goes up as the interest rate goes up, like our intuition suggests.

Market Equilibrium

So far we have restricted attention to one household. A macroeconomy would be composed of a number of these households, say N of them, so we stick these households together and consider what happens. In this model, that turns out to be trivial, since all households are identical, but the exercise will give you practice for more-difficult settings to come.

The basic exercise is to close our model by having the interest rate R determined endogenously. Recall, we said that households can be either lenders or borrowers, depending on whether b_1 is positive or negative, respectively. Well, the only borrowers and lenders in this economy are the N households, and all of them are alike. If they all want to borrow, there will be no one willing to lend, and there will be an excess demand for loans. On the other hand, if they all want to lend, there will be an excess supply of loans. More formally, we can write the aggregate demand for bonds as: $Nb_1^\star$. Market clearing requires:

$$Nb_1^\star = 0. \tag{3.9}$$

Of course, you can see that this requires that each household neither borrows nor lends, since all households are alike.

Now we turn to a formal definition of equilibrium. In general, a *competitive equilibrium* is a solution for all the variables of the economy such that: (i) all economic actors take prices as given; (ii) subject to those prices, all economic actors behave rationally; and (iii) all markets clear. When asked to define a competitive equilibrium for a specific economy, your task is to translate these three conditions into the specifics of the problem.

For the economy we are considering here, there are two kinds of prices: the price of consumption P and the price of borrowing R. The actors in the economy are the N households. There are two markets that must clear. First, in the goods market, we have:

$$Ny_t = Nc_t^\star,\ t = 1, 2. \tag{3.10}$$

Second, the bond market must clear, as given in equation (3.9) above. With all this written down, we now turn to defining a competitive equilibrium for this economy.

A competitive equilibrium in this setting is: a price of consumption $P^\star$; an interest rate $R^\star$; and values for $c_1^\star$, $c_2^\star$, and $b_1^\star$, such that:

- Taking $P^\star$ and $R^\star$ as given, all N households choose $c_1^\star$, $c_2^\star$, and $b_1^\star$ according to the maximization problem given in equations (3.4)-(3.6);
- Given these choices of $c_t^\star$, the goods market clears in each period, as given in equation (3.10); and
- Given these choices of $b_1^\star$, the bond market clears, as given in equation (3.9).

Economists are often pedantic about all the detail in their definitions of competitive equilibria, but providing the detail makes it very clear how the economy operates.

We now turn to computing the competitive equilibrium, starting with the credit market. Recall, we can write $b_1^\star$ as a function of the interest rate R, since the lending decision of each household hinges on the interest rate. We are interested in finding the interest rate that clears the bond market, i.e., the $R^\star$ such that $b_1^\star(R^\star) = 0$. We had:

$$b_1^\star(R) = Py_1 - \frac{P[y_2 + y_1(1 + R)]}{(1 + \beta)(1 + R)},$$

so we set the left-hand side to zero and solve for $R^\star$:

$$Py_1 = \frac{P[y_2 + y_1(1 + R^\star)]}{(1 + \beta)(1 + R^\star)}. \tag{3.11}$$

After some algebra, we get:

$$R^\star = \frac{y_2}{\beta y_1} - 1. \tag{3.12}$$

This equation makes clear that the equilibrium interest rate is determined by the incomes (y_1 and y_2) of the households in each period and by how impatient the households are (β). We can perform comparative statics here just like anywhere else. For example:

$$\frac{\partial R^\star}{\partial y_2} = \frac{1}{\beta y_1} > 0,$$

so if second-period income increases, then $R^\star$ does too. Conversely, if second-period income decreases, then $R^\star$ does too. This makes intuitive sense. If y_2 goes down, households will try to invest first-period income in bonds in order to smooth consumption between the two periods. In equilibrium this cannot happen, since net bond holdings must be zero, so the equilibrium interest rate must fall in order to provide a disincentive to investment, exactly counteracting households' desire to smooth consumption.

You can work through similar comparative statics and intuition to examine how the equilibrium interest rate changes in response to changes in y_1 and β. (See Exercise 3.2.)

Take note that in this model and with these preferences, only relative incomes matter. For example, if both y_1 and y_2 shrink by 50%, then y_2/y_1 does not change, so the equilibrium interest rate does not change. This has testable implications. Namely, we can test the reaction to a temporary versus a permanent decrease in income.

For example, suppose there is a temporary shock to the economy such that y_1 goes down by 10% today but y_2 is unchanged. The comparative statics indicate that the equilibrium interest rate must increase. This means that temporary negative shocks to income induce a higher interest rate. Now suppose that the negative shock is permanent. Then both y_1 and y_2 fall by 10%. This model implies that $R^\star$ does not change. This means that permanent shocks to not affect the interest rate.

The other price that is a part of the competitive equilibrium is $P^\star$, the price of a unit of consumption. It turns out that this price is not unique, since there is nothing in our economy to pin down what $P^\star$ is. The variable P does not even appear in the equations for $c_1^\star$ and $c_2^\star$. It does appear in the equation for $b_1^\star$, but P falls out when we impose the fact that $b_1^\star = 0$ in equilibrium; see equation (3.11). The intuition is that raising P has counteracting effects: it raises the value of a household's income but it raises the price of its consumption in exactly the same way, so raising P has no real effect. Since we cannot tack down $P^\star$, any number will work, and we have an infinite number of competitive equilibria. This will become clearer in Chapter 5.

3.3 An Infinite-Period Model

The version of the model in which the representative household lives for an infinite number of periods is similar to the two-period model from the previous section. The utility of the household is now:

$$U(c_1, c_2, \dots) = u(c_1) + \beta u(c_2) + \beta^2 u(c_3) + \cdots .$$

In each period t, the household faces a budget constraint:

$$Py_t + b_{t-1}(1 + R) = Pc_t + b_t.$$

Since the household lives for all $t = 1, 2, \dots,$ there are infinitely many of these budget constraints. The household chooses c_t and b_t in each period, so there are infinitely many choice variables and infinitely many first-order conditions. This may seem disconcerting, but don't let it intimidate you. It all works out rather nicely. We write out the maximization problem in condensed form as follows:

$$\max_{\{c_t, b_t\}_{t=1}^{\infty}} \sum_{t=1}^{\infty} \beta^{t-1} u(c_t), \text{ such that:}$$
$$Py_t + b_{t-1}(1 + R) = Pc_t + b_t, \ \forall\, t \in \{1, 2, \dots\}.$$

The "$\forall$" symbol means "for all", so the last part of the constraint line reads as "for all t in the set of positive integers".

To make the Lagrangean, we follow the rules outlined on page 15. In each time period t, the household has a budget constraint that gets a Lagrange multiplier λ_t. The only trick is that we use summation notation to handle all the constraints:

$$\mathcal{L} = \sum_{t=1}^{\infty} \beta^{t-1} u(c_t) + \sum_{t=1}^{\infty} \lambda_t \left[Py_t + b_{t-1}(1 + R) - Pc_t - b_t\right].$$

Now we are ready to take first-order conditions. Since there are infinitely many of them, we have no hope of writing them all out one by one. Instead, we just write the FOCs for period-t variables. The c_t FOC is pretty easy:

(FOC c_t) $$\frac{\partial \mathcal{L}}{\partial c_t} = \beta^{t-1} u'(c_t^\star) + \lambda_t^\star[-P] = 0.$$

Again, we use starred variables in first-order conditions because these equations hold only for the optimal values of these variables.

The first-order condition for b_t is harder because there are two terms in the summation that have b_t in them. Consider b_2. It appears in the $t = 2$ budget constraint as b_t, but it also appears in the $t = 3$ budget constraint as b_{t-1}. This leads to the $t + 1$ term below:

(FOC b_t) $$\frac{\partial \mathcal{L}}{\partial b_t} = \lambda_t^\star[-1] + \lambda_{t+1}^\star[(1 + R)] = 0.$$

Simple manipulation of this equation leads to:

$$\frac{\lambda_t^\star}{\lambda_{t+1}^\star} = 1 + R. \tag{3.13}$$

Rewriting equation (FOC c_t) gives us:

$$\beta^{t-1} u'(c_t^\star) = \lambda_t^\star P. \tag{3.14}$$

We can rotate this equation forward one period (i.e., replace t with $t+1$) to get the version for the next period:

$$\beta^t u'(c_{t+1}^\star) = \lambda_{t+1}^\star P. \tag{3.15}$$

Dividing equation (3.14) by equation (3.15) yields:

$$\frac{\beta^{t-1} u'(c_t^\star)}{\beta^t u'(c_{t+1}^\star)} = \frac{\lambda_t^\star P}{\lambda_{t+1}^\star P}, \text{ or:}$$
$$\frac{u'(c_t^\star)}{\beta u'(c_{t+1}^\star)} = \frac{\lambda_t^\star}{\lambda_{t+1}^\star}.$$

Finally, we multiply both sides by β and use equation (3.13) to get rid of the lambda terms on the right-hand side:

$$\frac{u'(c_t^\star)}{u'(c_{t+1}^\star)} = \beta(1 + R). \tag{3.16}$$

If you compare equation (3.16) to equation (3.7), you will find the Euler equations are the same in the two-period and infinite-period models. This is because the intertemporal trade-offs faced by the household are the same in the two models.

Just like in the previous model, we can analyze consumption patterns using the Euler equation. For example, if $\beta = 1/(1 + R)$, then the household's impatience exactly cancels with the incentives to invest, and consumption is constant over time. If the interest rate R is relatively high, then the right-hand side of equation (3.16) will be greater than one, and consumption will be rising over time.

A Present-Value Budget Constraint

Now we turn to a slightly different formulation of the model with the infinitely-lived representative household. Instead of forcing the household to balance its budget each period, now the household must merely balance the present value of all its budgets. (See Barro's page 71 for a discussion of present values.) We compute the present value of all the household's income:

$$\sum_{t=1}^{\infty} \frac{P y_t}{(1 + R)^{t-1}}.$$

This gives us the amount of dollars that the household could get in period 1 if it sold the rights to all its future income. On the other side, the present value of all the household's consumption is:

$$\sum_{t=1}^{\infty} \frac{Pc_t}{(1+R)^{t-1}}.$$

Putting these two present values together gives us the household's single present-value budget constraint. The household's maximization problem is:

$$\max_{\{c_t\}_{t=1}^{\infty}} \sum_{t=1}^{\infty} \beta^{t-1} u(c_t), \text{ such that:}$$
$$\sum_{t=1}^{\infty} \frac{P(y_t - c_t)}{(1+R)^{t-1}} = 0.$$

We use λ as the multiplier on the constraint, so the Lagrangean is:

$$\mathcal{L} = \sum_{t=1}^{\infty} \beta^{t-1} u(c_t) + \lambda \left[\sum_{t=1}^{\infty} \frac{P(y_t - c_t)}{(1+R)^{t-1}} \right].$$

The first-order condition with respect to c_t is:

$$\text{(FOC } c_t\text{)} \qquad \beta^{t-1} u'(c_t^\star) + \lambda^\star \left[\frac{P(-1)}{(1+R)^{t-1}} \right] = 0.$$

Rotating this forward and dividing the c_t FOC by the c_{t+1} FOC yields:

$$\frac{\beta^{t-1} u'(c_t^\star)}{\beta^t u'(c_{t+1}^\star)} = \frac{\lambda^\star \left[\frac{P}{(1+R)^{t-1}} \right]}{\lambda^\star \left[\frac{P}{(1+R)^t} \right]},$$

which reduces to:

$$\frac{u'(c_t^\star)}{u'(c_{t+1}^\star)} = \beta(1+R),$$

so we get the same Euler equation once again. It turns out that the problem faced by the household under the present-value budget constraint is equivalent to that in which there is a constraint for each period. Hidden in the present-value version are implied bond holdings. We could deduce these holdings by looking at the sequence of incomes y_t and chosen consumptions $c_t^\star$.

Exercises

Exercise 3.1 (Hard)
Consider the two-period model from Section 3.2, and suppose the period utility is:

$$u(c_t) = c_t^{\frac{1}{2}}.$$

Variable	Definition
$U(\cdot)$	Overall utility
t	Time
c_t	Consumption at period t
$u(\cdot)$	Period utility
β	Household's discount factor
y_t	Household's income in period t, in units of consumption
P	Cost of a unit of consumption
R	Nominal interest rate
b_t	Number of dollars of bonds bought at period t
λ_t	Lagrange multiplier in period t
N	Number of households

Table 3.1: Notation for Chapter 3

1. Determine the Euler equation in this case.
2. Determine the representative household's optimal choices: $c_1^\star$, $c_2^\star$, and $b_1^\star$.
3. Determine the equilibrium interest rate $R^\star$.
4. Determine the effect on the equilibrium interest rate $R^\star$ of a permanent negative shock to the income of the representative household. (I.e., both y_1 and y_2 go down by an equal amount.) How does this relate to the case in which $u(c_t) = \ln(c_t)$?

Exercise 3.2 (Easy)
Refer to equation (3.12), which gives the equilibrium interest rate $R^\star$ in the two-period model.

1. Suppose the representative household becomes more impatient. Determine the direction of the change in the equilibrium interest rate. (Patience is measured by β. You should use calculus.)
2. Suppose the representative household gets a temporary negative shock to its period-1 income y_1. Determine the direction of the change in the equilibrium interest rate. (Again, use calculus.)

Exercise 3.3 (Moderate)
Maxine lives for two periods. Each period, she receives an endowment of consumption goods: e_1 in the first, e_2 in the second. She doesn't have to work for this output. Her preferences for consumption in the two periods are given by: $u(c_1, c_2) = \ln(c_1) + \beta \ln(c_2)$, where

c_1 and c_2 are her consumptions in periods 1 and 2, respectively, and β is some discount factor between zero and one. She is able to save some of her endowment in period 1 for consumption in period 2. Call the amount she saves s. Maxine's savings get invaded by rats, so if she saves s units of consumption in period 1, she will have only $(1 - \delta)s$ units of consumption saved in period 2, where δ is some number between zero and one.

1. Write down Maxine's maximization problem. (You should show her choice variables, her objective, and her constraints.)
2. Solve Maxine's maximization problem. (This will give you her choices for given values of e_1, e_2, β, and δ.)
3. How do Maxine's choices change if she finds a way reduce the damage done by the rats? (You should use calculus to do comparative statics for changes in δ.)

Exercise 3.4 (Moderate)
An agent lives for five periods and has an edible tree. The agent comes into the world at time $t = 0$, at which time the tree is of size x_0. Let c_t be the agent's consumption at time t. If the agent eats the whole tree at time t, then $c_t = x_t$ and there will be nothing left to eat in subsequent periods. If the agent does not eat the whole tree, then the remainder grows at the simple growth rate α between periods. If at time t the agent saves $100s_t$ percent of the tree for the future, then $x_{t+1} = (1 + \alpha)s_t x_t$. All the agent cares about is consumption during the five periods. Specifically, the agent's preferences are: $U = \sum_{t=0}^{4} \beta^t \ln(c_t)$. The tree is the only resource available to the agent.

Write out the agent's optimization problem.

Chapter 4

The Demand for Money

This chapter seeks to explain one stark fact: the authors used to withdraw \$20 when they went to the ATM, whereas now they tend to withdraw \$300. We are going to make a model to examine this question. In our model, a consumer chooses how often to go to the bank and how much money to withdraw once there.

Let T be the amount of time (in fractions of a year) between a consumer's trips to the bank to get money. If T is $1/3$, then the consumer goes to the bank every 4 months, or three times a year. For arbitrary T, the consumer makes $1/T$ trips to the bank in a year.

Going to the bank is a pain. It takes time and effort, and the bank may charge for each withdrawal. We accumulate all such expenses into some dollar cost γ. We could derive γ by: (i) calculating the consumer's opportunity cost of time; (ii) multiplying that by the amount of time required to go to the bank; and (iii) adding any fees charged by the bank.

The cost per year of this consumer's trips to the bank is just the number of trips times the cost per trip, so the consumer's annual transactions costs are: $(1/T)(\gamma)$. If all the prices in the economy double, then these costs double, since both bank fees and the opportunity cost of the consumer's time double.[1] Accordingly, in order to get the real impact on the consumer of these annual costs, we need to adjust them by the price level P, so the consumers real[2] Here we see that if prices double, then both P and γ double, and those extra factors of two cancel, so real costs do not change, as we require.

[1] If all prices in the economy double, then the prices of anything the consumer produces double. Put differently, the consumer's wage doubles. Either of these implies that the opportunity cost of the consumer's time doubles.

[2] The distinction between "real" and "nominal" values means the same thing here as in Barro's discussion about real versus nominal GDP. (See his Chapter 1.) "Nominal" values are actual dollars. "Real" dollars are scaled so that their purchasing power is constant. In this model, a unit of consumption costs P dollars. This is the observed or "nominal" price. If prices double, each dollar has half the purchasing power, so any nominal amount of dollars goes down in value by a factor of two. In general, we convert from nominal dollar amounts to real dollar amounts by dividing the nominal amounts by the price level P.

Now, going to the bank is costly, but the consumer still does it because the consumer needs to withdraw money in order to buy things. Assume that our consumer spends Pc dollars on consumption each year, where this spending is smooth from day to day. (This gives the consumer c real dollars of consumption each year.) In order to pay for all this consumption, the consumer needs enough money on hand at any given instant to make the purchases.

We can calculate how much money the consumer spends between trips to the bank. (Recall, T measures time between trips, in fractions of a year.) If T is 1, then the consumer spends Pc. If T is 1/2, then the consumer spends $(Pc)/2$. In general, the consumer spends PcT dollars between trips to the bank. That is the amount the consumer must withdraw on each trip. The consumer can choose to go less often (T bigger), but only if the consumer is willing to withdraw more on each trip.

Barro's Figure 4.1 gives a graphical illustration of how the consumer's money holdings evolve over time. After going to the bank, the consumer's money holdings decline linearly, so the consumer's average money holdings are:

$$\overline{m} = \frac{PcT}{2}.$$

(This uses the fact that the area of a triangle is one half the base times the height.) The consumer's average real money holdings are:

$$\frac{\overline{m}}{P} = \frac{cT}{2}.$$

Notice that the consumer's average money holdings are increasing in the amount of time between bank visits, i.e., the longer between visits, the more money the consumer holds on average.

Since there are transactions costs involved in each trip to the bank, we might wonder why the consumer does not go once, withdraw a ton of money, and get it all over with. All the money that the consumer holds onto between trips to the bank does not earn interest, but money left in the bank does earn interest. This foregone interest is the opportunity cost of holding money. If the annual nominal interest rate is R, then each year the consumer loses out on about:

$$R\overline{m} = \frac{RPcT}{2}$$

dollars of interest. Notice that higher average money holdings result in larger amounts of foregone interest.

We can state this dollar amount of interest in real terms:

$$\text{real interest foregone annually} = \frac{R\overline{m}}{P} = \frac{RcT}{2}.$$

We are now ready to put all this together. The consumer chooses T, the time between bank visits. We have calculated the annual cost of the consumer's bank visits and the annual cost in foregone interest from the consumer's money holdings, both in real terms. Adding these two costs together gives us:

$$\text{total annual real costs} = \frac{\gamma}{PT} + \frac{RcT}{2}. \tag{4.1}$$

This equation is graphed in Barro's Figure 4.2.

We now use calculus to calculate the consumer's optimal behavior. Namely, we want to derive the consumer's cost-minimizing choice of the time T between visits to the bank to withdraw money. We are interested in the minimum costs, so we take the first-order condition of equation (4.1) with respect to T:

$$\frac{\partial}{\partial T}\left(\frac{\gamma}{PT^\star} + \frac{RcT^\star}{2}\right) = 0, \text{ or :}$$
$$-\frac{\gamma}{P(T^\star)^2} + \frac{Rc}{2} = 0.$$

Solving this expression for $T^\star$ yields:

$$T^\star = \sqrt{\frac{2\gamma}{PRc}}.$$

With this answer, we can now write down the algebraic expression for the consumer's average holdings of real money $\overline{m}/P$, which Barro calls $\phi(R, c, \gamma/P)$. The consumer's average money holdings are:

$$\frac{\overline{m}}{P} = \left(\frac{1}{2}\right) cT.$$

When we plug in our expression for $T^\star$, we get:

$$\phi(R, c, \gamma/P) = \left(\frac{1}{2}\right) cT^\star = \left(\frac{1}{2}\right) c\sqrt{\frac{2\gamma}{PRc}} = \sqrt{\frac{\gamma c}{2PR}}. \tag{4.2}$$

We can do comparative statics to examine how these money holdings are affected by changes in the underlying parameters. See the exercises for examples. The solutions to these exercises provide the answer to question posed at the beginning of this chapter: Why do the authors now withdraw \$300 from the ATM, whereas they used to withdraw only \$20? Well, today they spend more money, the opportunity cost of their time is higher, the transactions costs at the ATM are higher, and interest rates are lower.

Presumably, the consumer that underlies this model of money demand also makes a choice of how much to consume c each year. We now briefly discuss whether it makes sense to have the consumer choose c and T separately.

When a consumer chooses how much to consume c, she considers the price of the goods she would be buying. Higher prices generally mean the consumer chooses to consume less. Now, the costs of a good as rung up at the cash register are not the full costs to the consumer of acquiring the good. The consumer might have to: expend effort to get to the store; spend valuable time waiting in line; or spend time and money to have the cash on hand to make the purchase. All of these things add to the cost that the consumer faces when making a purchase decision. If those additional costs change, then the consumer's consumption will change. This implies that the same things that the consumer considers when choosing T will affect the consumer's optimal choice of c. Since c was one of the things that went into our determination of $T^\star$, it is a shortcoming of our model that we assumed that we could separate these decisions.

Think about the following example. Suppose ATM fees go up temporarily to $100 per transaction. In our model, this implies that γ increases, so $T^\star$ goes up, since people want to go to the bank less often. Our model assumes that c is fixed, but in reality c will fall because of the new ATM fees, since consumption is now more expensive (especially consumption of goods that have to be purchased with cash). Hence, our solution for $T^\star$ (which assumes a fixed c) is liable differ from that implied by a more sophisticated model. Since c goes down as γ goes up, and $\partial T^\star/\partial c < 0$, T would go up by more in a model that took the relationship between c and T into account.

Variable	Definition
T	Time (in years) between trips to the bank
γ	Opportunity cost of a trip to the bank
P	Price of consumption
c	Consumption per year
$\bar{m}$	Consumer's average money holdings
R	Nominal interest rate
$\phi(\cdot)$	Real money demand

Table 4.1: Notation for Chapter 4

Exercises

Exercise 4.1 (Easy)

1. Determine the effect of an increase in the interest rate R on the consumer's money demand $\phi(\cdot)$, as given by equation (4.2).
2. Determine the effect of an increase in the consumer's consumption c on the consumer's money demand $\phi(\cdot)$, as given by equation (4.2).
3. Determine the effect of an increase in the consumer's real transactions costs γ/P on the consumer's money demand $\phi(\cdot)$, as given by equation (4.2).

(Use calculus for all three parts. The way you do the last one is to replace γ/P with some other variable, say α, and differentiate with respect to the new variable.)

Chapter 5

The Market-Clearing Model

Most of the models that we use in this book build on two common assumptions. First, we assume that there exist markets for all goods present in the economy, and that all markets clear. Second, we assume that all agents behave competitively, which means that they take prices as given. Models that satisfy these assumptions are called *general equilibrium* models. There are a number of important results that apply to all general equilibrium models, regardless of what kind of goods, agents, or technologies are used. In this chapter, we will demonstrate three of these results within a general setting. Many of the models that we use throughout the rest of the book will be special cases of the general model presented here. Since we omit most of the simplifying assumptions that we make in other chapters, the treatment is more formal and mathematical than usual.

Section 5.1 introduces our general equilibrium framework. In Section 5.2 we show that within this framework the general price level is undetermined. This implies that prices can be normalized without loss of generality. For example, in many models we set the price of the consumption good to be one. In Section 5.3 we show that in a general equilibrium model one market clearing constraint is redundant, a fact known as *Walras' Law*. Section 5.4 presents the *First Welfare Theorem*, which states that under certain conditions equilibria in a general equilibrium model are efficient.

5.1 A General Pure-Exchange Economy

We will consider an economy with many different goods and consumers. Instead of having a representative consumer, we allow for the possibility that each consumer has a different utility function. However, we make one simplification: there is no production in the economy. The consumers have endowments of goods and can trade their endowments in markets, but there is no possibility of producing any goods in excess of the endowments.[1]

There are N different goods in the economy, where N is any positive integer. For each good there is a market, and the price of good n is denoted p_n. There are I different consumers. Each consumer has a utility function over her consumption of the N goods in the economy. Consumption of good n by consumer i is denoted as c_n^i, and the utility function for consumer i is $u_i(c_1^i, c_2^i, \ldots, c_N^i)$. Notice that the utility function is indexed by i, so that it can be different for each consumer. The consumers also have endowments of the N goods, where e_t^i is the endowment of consumer i of good n.

All consumers meet at the beginning of time in a central marketplace. Here the consumers can sell their endowments and buy consumption goods. If consumer i sells all her endowments, her total income is $\sum_{n=1}^N p_n e_n^i$. Similarly, total expenditure on consumption goods is $\sum_{n=1}^N p_n c_n^i$. Consumer i maximizes utility subject to her budget constraint, which states that total expenditure on consumption has to equal total income from selling the endowment. Mathematically, the problem of consumer i is:

$$\max_{\{c_n^i\}_{n=1}^N} \ u_i(c_1^i, c_2^i, \ldots, c_N^i) \ \text{ subject to:} \tag{5.1}$$
$$\sum_{n=1}^N p_n c_n^i = \sum_{n=1}^N p_n e_n^i.$$

We will also need a market-clearing constraint for each of the goods. The market-clearing condition for good n is:

$$\sum_{i=1}^I c_n^i = \sum_{i=1}^I e_n^i. \tag{5.2}$$

Note that in the budget constraint we sum over all goods for one consumer, while in the market-clearing conditions we sum over all consumers for one good. The only additional assumptions that we will make throughout this chapter are: that I and N are positive integers, that all endowments e_n^i are positive and that all utility functions are strictly increasing in all arguments. The assumption of increasing utility functions is important because it implies that all prices are positive in equilibrium. We will use this fact below. Notice that we do not make any further assumptions like differentiability or concavity, and that we do not restrict attention to specific functional forms for utility. The results in this chapter rest solely on the general structure of the market-clearing model. We are now ready to define an equilibrium for this economy along the lines developed in Chapter 3.

An *allocation* is a set of values for consumption for each good and each consumer. A *competitive equilibrium* is an allocation $\{c_1^i, c_2^i, \ldots, c_N^i\}_{i=1}^I$ and a set of prices $\{p_1, p_2, \ldots, p_N\}$ such that:

[1]While this assumption may seem restrictive, in fact all results of this chapter can be shown for production economies as well. However, notation and algebra are more complicated with production, so we concentrate on the pure-exchange case.

- Taking prices as given, each consumer i chooses $\{c_1^i, c_2^i, \dots, c_N^i\}$ as a solution to the maximization problem in equation (5.1); and

- Given the allocation, all market-clearing constraints in equation (5.2) are satisfied.

The model is far more general than it looks. For example, different goods could correspond to different points of time. In that case, the budget constraint would be interpreted as a present-value budget constraint, as introduced in Chapter 3. We can also incorporate uncertainty, in which case different goods would correspond to different states of the world. Good 1 could be consumption of sun-tan lotion in case it rains tomorrow, while good 2 could be sun-tan lotion in case it's sunny. Presumably, the consumer would want to consume different amounts of these goods, depending on the state of the world. By using such time- and state-contingent goods, we can adapt the model to almost any situation.

5.2 Normalization of Prices

In our model, the general level of prices is undetermined. For example, given any equilibrium, we can double all prices and get another equilibrium. We first ran into this phenomenon in the credit-market economy of Section 3.2, where it turned out that the price level P was arbitrary. An important application is the possibility of normalizing prices. Since it is possible to multiply prices by a positive constant and still have an equilibrium, the constant can be chosen such that one price is set to one. For example, if we want to normalize the price of the first good, we can choose the constant to be $1/p_1$. Then, when we multiply all prices by this constant, the normalized price of the first good becomes $(p_1)(1/p_1) = 1$. If for every equilibrium there is another one in which the price of the first good is one, there is no loss in generality in assuming that the price is one right away. Without always mentioning it explicitly, we make use of this fact in a number of places throughout this book. Normally the price of the consumption good is set to one, so that all prices can be interpreted in terms of the consumption good.[2] The good whose price is set to one is often called the *numéraire*.

In order to show that the price level is indeterminate, we are going to assume that we have already found an allocation $\{c_1^i, c_2^i, \dots, c_N^i\}_{i=1}^I$ and a price system $\{p_1, p_2, \dots, p_N\}$ that satisfy all the conditions for an equilibrium. We now want to show that if we multiply all prices by a constant $\gamma > 0$ we will still have an equilibrium. That is, the allocation $\{c_1^i, c_2^i, \dots, c_N^i\}_{i=1}^I$ will still satisfy market-clearing, and the values for consumption will still be optimal choices for the consumers given the new price system $\{\gamma p_1, \gamma p_2, \dots, \gamma p_N\}$.

It is obvious that the market-clearing constraints will continue to hold, since we have not changed the allocation and the prices do not enter in the market-clearing constraints. Therefore we only need to show that the allocation will still be optimal, given the new price

[2]Examples are the labor market model in Section 6.1 and the business-cycle model in Chapter 9. In both cases, the price of consumption is set to one.

system. We know already that the allocation is an optimal choice for the consumers given the old price system. If we can show that the new price system does not change the budget constraint of the consumer, then the consumer's problem with the new prices will be equivalent to the original problem, so it will have the same solution. The budget constraint with the new prices is:

$$\sum_{n=1}^{N}(\gamma p_n)c_n^i = \sum_{n=1}^{N}(\gamma p_n)e_n^i.$$

We can pull the common γ terms outside the summations, so we can divide each side by γ to yield:

$$\sum_{n=1}^{N} p_n c_n^i = \sum_{n=1}^{N} p_n e_n^i,$$

which is equal to the budget constraint under the original price system. The consumer's problem does not change, so it still has the same solution. This shows that the allocation $\{c_1^i, c_2^i, \ldots, c_N^i\}_{i=1}^I$ and prices $\{\gamma p_1, \gamma p_2, \ldots, \gamma p_N\}$ form an equilibrium as well.

The basic idea is that only relative prices matter, not absolute prices. If all prices are multiplied by a constant, income from selling the endowment will increase in the same proportion as the cost of consumption goods. As long as the relative prices are constant, such a change will not influence the decisions of consumers. Note that we did not need to look at any first-order conditions to prove our point. The possibility of normalizing prices derives from the basic structure of this market economy, not from specific assumptions about utility or technology.

5.3 Walras' Law

In a general equilibrium model, one market-clearing constraint is redundant. This means that if each consumer's budget constraint is satisfied and all but one market-clearing conditions hold, then the last market-clearing condition is satisfied automatically. This fact is of practical value, because it implies that we can omit one market-clearing constraint right away when computing an equilibrium. Without mentioning it, we made already use of this in Section 3.2. While the definition of equilibrium required the goods market to clear, the market-clearing constraints for goods were not actually used afterwards. This was possible because they were implied by the budget constraints and the fact that the bond market cleared. This feature of general equilibrium models is known as *Walras' Law*.

To see that Walras' law holds in our general pure-exchange economy, assume that the budget constraints for each of the I consumers and the market-clearing constraints for the first $N-1$ goods are satisfied. We want to show that the last market-clearing constraint for good N is also satisfied. Summing the budget constraints over all consumers yields:

$$\sum_{i=1}^{I}\sum_{n=1}^{N} p_n c_n^i = \sum_{i=1}^{I}\sum_{n=1}^{N} p_n e_n^i.$$

Rearranging gives:

$$\sum_{n=1}^{N}\sum_{i=1}^{I} p_n c_n^i = \sum_{n=1}^{N}\sum_{i=1}^{I} p_n e_n^i,$$
$$\sum_{n=1}^{N} p_n \sum_{i=1}^{I} c_n^i = \sum_{n=1}^{N} p_n \sum_{i=1}^{I} e_n^i, \text{ or:}$$
$$\sum_{n=1}^{N} p_n \left[\sum_{i=1}^{I} c_n^i - \sum_{i=1}^{I} e_n^i\right] = 0. \tag{5.3}$$

Inside the brackets we have the difference between the total consumption and the total endowment of good n. If the market for good n clears, this difference is zero. Since we assume that the first $N-1$ markets clear, equation (5.3) becomes:

$$p_N \left[\sum_{i=1}^{I} c_N^i - \sum_{i=1}^{I} e_N^i\right] = 0.$$

Since $p_N > 0$, this implies:

$$\sum_{i=1}^{I} c_N^i - \sum_{i=1}^{I} e_N^i = 0, \text{ or:}$$
$$\sum_{i=1}^{I} c_N^i = \sum_{i=1}^{I} e_N^i.$$

Thus the Nth market will clear as well.

The intuition behind this result is easiest to see when the number of markets is small. If there is only one good, say apples, the budget constraints of the consumers imply that each consumer eats as many apples as she is endowed with. Then the market-clearing constraint has to be satisfied as well, since it is already satisfied on the level of each consumer. Now assume there is one more good, say oranges, and the market-clearing constraint for apples is satisfied. That implies that total expenditures on apples equal total income from selling apples to other consumers. Since each consumer balances spending with income, expenditures have to equal income for oranges as well, so the market for oranges clears.

5.4 The First Welfare Theorem

The first two features of general equilibrium models that we presented in this chapter were technical. They are of some help in computing equilibria, but taken for themselves they do not provide any deep new insights that could be applied to the real world. The situation is different with the last feature that we are going to address, the efficiency of outcomes in general equilibrium economies. This result has important implications for the welfare properties of economic models, and it plays a key role in the theory of comparative economic systems.

Before we can show that equilibria in our model are efficient, we have to make precise what exactly is meant by efficiency. In economics, we usually use the concept of *Pareto efficiency*. Another term for Pareto efficiency is Pareto optimality, and we will use both versions interchangeably. An allocation is Pareto efficient if it satisfies the market-clearing conditions and if there is no other allocation that: (1) also satisfies the market-clearing conditions; and (2) makes everyone better off. In our model, an allocation $\{c_1^i, c_2^i, \dots, c_N^i\}_{i=1}^I$ is therefore Pareto efficient if the market-clearing constraint in equation (5.2) holds for each of the N goods and if there is no other allocation $\{\bar{c}_1^i, \bar{c}_2^i, \dots, \bar{c}_N^i\}_{i=1}^I$ that also satisfies market-clearing and such that:

$$u(\bar{c}_1^i, \bar{c}_2^i, \dots, \bar{c}_N^i) > u(c_1^i, c_2^i, \dots, c_N^i)$$

for every consumer i.[3] Notice that the concept of Pareto optimality does not require us to take any stand on the issue of distribution. For example, if utility functions are strictly increasing, one Pareto-optimal allocation is to have one consumer consume all the resources in the economy. Such an allocation is clearly feasible, and every alternative allocation makes this one consumer worse off. A Pareto-efficient allocation is therefore not necessarily one that many people would consider "fair" or even "optimal". On the other hand, many people would agree that it is better to make everyone better off as long as it is possible to do so. Therefore we can interpret Pareto efficiency as a minimal standard for a "good" allocation, rather than as a criterion for the "best" one.

We now want to show that any equilibrium allocation in our economy is necessarily Pareto optimal. The equilibrium consists of an allocation $\{c_1^i, c_2^i, \dots, c_N^i\}_{i=1}^I$ and a price system $\{p_1, p_2, \dots, p_N\}$. Since market-clearing conditions hold for any equilibrium allocation, the first requirement for Pareto optimality is automatically satisfied. The second part takes a little more work. We want to show that there is no other allocation that also satisfies market-clearing and that makes everyone better off. We are going to prove this by contradiction. That is, we will assume that such a better allocation actually exists, and then we will show that this leads us to a contradiction. Let us therefore assume that there is another allocation $\{\bar{c}_1^i, \bar{c}_2^i, \dots, \bar{c}_N^i\}_{i=1}^I$ that satisfies market-clearing and such that:

$$u(\bar{c}_1^i, \bar{c}_2^i, \dots, \bar{c}_N^i) > u(c_1^i, c_2^i, \dots, c_N^i)$$

[3] A weaker notion of Pareto efficiency replaces the strict inequality with weak inequalities plus the requirement that at least one person is strictly better off. The proof of the First Welfare Theorem still goes through with the weaker version, but for simplicity we use strict inequalities.

for every consumer i. We know that consumer i maximizes utility subject to the budget constraint. Since the consumer chooses $\{c_1^i, c_2^i, \ldots, c_N^i\}$ even though $\{\bar{c}_1^i, \bar{c}_2^i, \ldots, \bar{c}_N^i\}$ yields higher utility, it has to be the case that $\{\bar{c}_1^i, \bar{c}_2^i, \ldots, \bar{c}_N^i\}$ violates the consumer's budget constraint:

$$\sum_{n=1}^{N} p_n \bar{c}_n^i > \sum_{n=1}^{N} p_n e_n^i. \tag{5.4}$$

Otherwise, the optimizing consumers would not have chosen the consumptions in the allocation $\{c_1^i, c_2^i, \ldots, c_N^i\}_{i=1}^I$ in the first place. Summing equation (5.4) over all consumers and rearranging yields:

$$\begin{aligned}
\sum_{i=1}^{I}\sum_{n=1}^{N} p_n \bar{c}_n^i &> \sum_{i=1}^{I}\sum_{n=1}^{N} p_n e_n^i, \\
\sum_{n=1}^{N}\sum_{i=1}^{I} p_n \bar{c}_n^i &> \sum_{n=1}^{N}\sum_{i=1}^{I} p_n e_n^i, \\
\sum_{n=1}^{N} p_n \sum_{i=1}^{I} \bar{c}_n^i &> \sum_{n=1}^{N} p_n \sum_{i=1}^{I} e_n^i, \text{ so:}
\end{aligned}$$

$$\sum_{n=1}^{N} p_n \left[\sum_{i=1}^{I} \bar{c}_n^i - \sum_{i=1}^{I} e_n^i\right] > 0.$$

We assumed that the allocation $\{\bar{c}_1^i, \bar{c}_2^i, \ldots, \bar{c}_N^i\}_{i=1}^I$ satisfied market-clearing. Therefore the terms inside the brackets are all zero. This implies $0 > 0$, which is a contradiction. Therefore, no such allocation $\{\bar{c}_1^i, \bar{c}_2^i, \ldots, \bar{c}_N^i\}_{i=1}^I$ can exist, and the original equilibrium allocation $\{c_1^i, c_2^i, \ldots, c_N^i\}_{i=1}^I$ is Pareto optimal.

Since any competitive equilibrium is Pareto optimal, there is no possibility of a redistribution of goods that makes everyone better off than before. Individual optimization together with the existence of markets imply that all gains from trade are exploited.

There is also a partial converse to the result that we just proved, the "Second Welfare Theorem". While the First Welfare Theorem says that every competitive equilibrium is Pareto efficient, the Second Welfare Theorem states that every Pareto optimum can be implemented as a competitive equilibrium, as long as wealth can be redistributed in advance. The Second Welfare Theorem rests on some extra assumptions and is harder to prove, so we omit it here. In economies with a single consumer there are no distributional issues, and the two theorems are equivalent.

Variable	Definition
N	Number of goods
p_n	Price of good n
I	Number of consumers
c_n^i	Consumption of good n by consumer i
$u_i(\cdot)$	Utility function of consumer i
e_n^i	Endowment with good n of consumer i
γ	Arbitrary proportionality factor

Table 5.1: Notation for Chapter 5

Exercises

Exercise 5.1 (Easy)
Show that Walras' law holds for the credit-market economy that we discussed in Chapter 3.2. That is, use the consumer's budget constraints and the market-clearing conditions for goods to derive the market-clearing condition for bonds in equation (3.9).

Exercise 5.2 (Hard)
Assume that the equilibrium price of one of the N goods is zero. What is the economic interpretation of this situation? Which of our assumptions ruled out that a price equals zero? Why? Does Walras' Law continue to hold? What about the First Welfare Theorem?

Chapter 6

The Labor Market

This chapter works out the details of two separate models. Section 6.1 contains a one-period model in which households are both demanders and suppliers of labor. Market clearing in the labor market determines the equilibrium wage rate. Section 6.2 further develops the two-period model from Chapter 3. In this case, the households are permitted to choose their labor supply in each period.

6.1 Equilibrium in the Labor Market

This economy consists of a large number of identical households. Each owns a farm on which it employs labor to make consumption goods, and each has labor that can be supplied to other farmers. For each unit of labor supplied to others, a household receives a wage w, which is paid in consumption goods. Households take this wage as given. In order to make the exposition clear, we prohibit a household from providing labor for its own farm. (This has no bearing on the results of the model.)

The first task of the representative household is to maximize the profit of its farm. The output of the farm is given by a production function $f(l_d)$, where l_d is the labor demanded (i.e., employed) by that farm. The only expense of the farm is its labor costs, so the profit of the farm is: $\pi = f(l_d) - wl_d$. The household that owns the farm chooses how much labor l_d to hire. The first-order condition with respect to l_d is:

$$\frac{\partial \pi}{\partial l_d} = f'(l_d^\star) - w = 0, \text{ so:}$$

$$w = f'(l_d^\star). \tag{6.1}$$

This implies that the household will continue to hire laborers until the marginal product of additional labor has fallen to the market wage. Equation (6.1) pins down the optimal labor input $l_d^\star$. Plugging this into the profit equation yields the maximized profit of the household: $\pi^\star = f(l_d^\star) - wl_d^\star$.

After the profit of the farm is maximized, the household must decide how much to work on the farms of others and how much to consume. Its preferences are given by $u(c, l_s)$, where c is the household's consumption, and l_s is the amount of labor that the household supplies to the farms of other households. The household gets income $\pi^\star$ from running its own farm and labor income from working on the farms of others. Accordingly, the household's budget is:

$$c = \pi^\star + wl_s,$$

so Lagrangean for the household's problem is:

$$\mathcal{L} = u(c, l_s) + \lambda[\pi^\star + wl_s - c].$$

The first-order condition with respect to c is:

$$u_1(c^\star, l_s^\star) + \lambda^\star[-1] = 0, \tag{FOC c}$$

and that with respect to l_s is:

$$u_2(c^\star, l_s^\star) + \lambda^\star[w] = 0. \tag{FOC l_s}$$

Solving each of these for λ and setting them equal yields:

$$-\frac{u_2(c^\star, l_s^\star)}{u_1(c^\star, l_s^\star)} = w, \tag{6.2}$$

so the household continues to supply labor until its marginal rate of substitution of labor for consumption falls to the wage the household receives.

Given particular functional forms for $u(\cdot)$ and $f(\cdot)$, we can solve for the optimal choices $l_d^\star$ and $l_s^\star$ and compute the equilibrium wage. For example, assume:

$$u(c, l) = \ln(c) + \ln(1 - l), \text{ and:}$$
$$f(l) = Al^\alpha.$$

Under these functional forms, equation (6.1) becomes:

$$w = A\alpha(l_d^\star)^{\alpha-1}, \text{ so:}$$
$$l_d^\star = \left(\frac{A\alpha}{w}\right)^{\frac{1}{1-\alpha}}. \tag{6.3}$$

This implies that the profit $\pi^\star$ of each household is:

$$\pi^\star = A\left(\frac{A\alpha}{w}\right)^{\frac{\alpha}{1-\alpha}} - w\left(\frac{A\alpha}{w}\right)^{\frac{1}{1-\alpha}}.$$

After some factoring and algebraic manipulation, this becomes:

$$\pi^\star = A(1-\alpha)\left(\frac{A\alpha}{w}\right)^{\frac{\alpha}{1-\alpha}}. \tag{6.4}$$

Under the given preferences, we have $u_1(c,l) = 1/c$ and $u_2(c,l) = -1/(1-l)$. Recall, the budget equation implies $c = \pi + wl_s$. Plugging these into equation (6.2) gives us:

$$\frac{\pi^\star + wl_s^\star}{1-l_s^\star} = w,$$

which reduces to:

$$l_s^\star = \frac{1}{2} - \frac{\pi^\star}{2w}.$$

Plugging in $\pi^\star$ from equation (6.4) yields:

$$l_s^\star = \frac{1}{2} - \left(\frac{1}{2w}\right) A(1-\alpha)\left(\frac{A\alpha}{w}\right)^{\frac{\alpha}{1-\alpha}},$$

which reduces to:

$$l_s^\star = \frac{1}{2} - \left(\frac{1-\alpha}{2\alpha}\right)\left(\frac{A\alpha}{w}\right)^{\frac{1}{1-\alpha}}.$$

Now we have determined the household's optimal supply of labor $l_s^\star$ as a function of the market wage w, and we have calculated the household's optimal choice of labor to hire $l_d^\star$ for a given wage. Since all household's are identical, equilibrium occurs where the household's supply equals the household's demand. Accordingly, we set $l_s^\star = l_d^\star$ and call the resulting wage $w^\star$:

$$\frac{1}{2} - \left(\frac{1-\alpha}{2\alpha}\right)\left(\frac{A\alpha}{w^\star}\right)^{\frac{1}{1-\alpha}} = \left(\frac{A\alpha}{w^\star}\right)^{\frac{1}{1-\alpha}}.$$

We gather like terms to get:

$$\frac{1}{2} = \left[1 + \left(\frac{1-\alpha}{2\alpha}\right)\right]\left(\frac{A\alpha}{w^\star}\right)^{\frac{1}{1-\alpha}}.$$

Further algebraic manipulation yields:

$$w^\star = A\alpha\left(\frac{1+\alpha}{\alpha}\right)^{1-\alpha}.$$

Finally, we plug this equilibrium wage back into our expressions for $l_s^\star$ and $l_d^\star$, which were in terms of w. For example, plugging the formula for $w^\star$ into equation (6.3) gives us:

$$l_d^\star = \left(\frac{A\alpha}{w^\star}\right)^{\frac{1}{1-\alpha}} = \left[\frac{A\alpha}{A\alpha\left(\frac{1+\alpha}{\alpha}\right)^{1-\alpha}}\right]^{\frac{1}{1-\alpha}} = \frac{\alpha}{1+\alpha}.$$

Of course, we get the same answer for $l_s^\star$, since supply must equal demand in equilibrium.

Given these answers for $l_s^\star$, $l_d^\star$, and $w^\star$, we can perform comparative statics to determine how the equilibrium values are influenced by changes in the underlying parameters. For example, suppose the economy experiences a positive shock to its productivity. This could be represented by an increase in the A parameter to the production function. We might be interested in how that affects the equilibrium wage:

$$\frac{\partial w^\star}{\partial A} = \alpha \left(\frac{\alpha+1}{\alpha}\right)^{1-\alpha} > 0,$$

so the equilibrium wage will increase. Just by inspecting the formula for $l_d^\star$ and $l_s^\star$, we know that labor supply and labor demand will be unchanged, since A does not appear. The intuition of this result is straightforward. With the new, higher productivity, households will be more inclined to hire labor, but this is exactly offset by the fact that the new wage is higher. On the other side, households are enticed to work more because of the higher wage, but at the same time they are wealthier, so they want to enjoy more leisure, which is a normal good. Under these preferences, the two effects cancel.

Variable	Definition
w	Wage in consumption goods per unit of labor
l_d	Labor demanded by owner of farm
$f(l_d)$	Output of farm
π	Profit of farm
c	Consumption of household
l_s	Labor supplied by household
$u(c, l_s)$	Utility of household
λ	Lagrange multiplier
$\mathcal{L}$	Lagrangean
A	Parameter of the production function
α	Parameter of the production function

Table 6.1: Notation for Section 6.1

6.2 Intertemporal Labor Choice

The model in this section is a pure extension of that developed in Section 3.2. In that model the representative household lived for two periods. Each period, the household got an endowment, e_1 and e_2. The household chose each period's consumption, c_1 and c_2, and the number of dollars of bonds b_1 to carry from period 1 to period 2.

The model presented here is almost identical. The only difference is that the household exerts labor effort in order to acquire goods instead of having them endowed exogenously. In particular, the household has some production function: $y_t = f(l_t)$. The household chooses each period's labor, l_1 and l_2. The income y_t takes the place of the endowment e_t in the model from Chapter 3.

The household's maximization problem is:

$$\max_{c_1,c_2,l_1,l_2,b_1} \{u(c_1,l_1) + \beta u(c_2,l_2)\}, \text{ subject to:}$$
$$Pf(l_1) = Pc_1 + b_1, \text{ and:}$$
$$Pf(l_2) + b_1(1+R) = Pc_2.$$

Refer to Chapter 3 for a discussion of: (i) the budget constraints, (ii) the meaning of the price level P and interest rate R, and (iii) how the bonds work. The Lagrangean is:

$$\mathcal{L} = u(c_1,l_1) + \beta u(c_2,l_2) + \lambda_1[Pf(l_1) - Pc_1 - b_1] + \lambda_2[Pf(l_2) + b_1(1+R) - Pc_2].$$

There are seven first-order conditions:

$$\text{(FOC } c_1\text{)} \qquad u_1(c_1^\star, l_1^\star) + \lambda_1^\star[-P] = 0,$$
$$\text{(FOC } c_2\text{)} \qquad \beta u_1(c_2^\star, l_2^\star) + \lambda_2^\star[-P] = 0,$$
$$\text{(FOC } l_1\text{)} \qquad u_2(c_1^\star, l_1^\star) + \lambda_1^\star[Pf'(l_1^\star)] = 0,$$
$$\text{(FOC } l_2\text{)} \qquad \beta u_2(c_2^\star, l_2^\star) + \lambda_2^\star[Pf'(l_2^\star)] = 0, \text{ and:}$$
$$\text{(FOC } b_1\text{)} \qquad \lambda_1^\star[-1] + \lambda_2^\star[(1+R)] = 0.$$

We leave off the FOCs with respect to λ_1 and λ_2 because we know that they reproduce the constraints. Solving equations (FOC c_1) and (FOC c_2) for the Lagrange multipliers and plugging into equation (FOC b_1) yields:

$$\frac{u_1(c_1^\star, l_1^\star)}{u_1(c_2^\star, l_2^\star)} = \beta(1+R). \tag{6.5}$$

This is the same Euler equation we saw in Chapter 3. Solving equations (FOC l_1) and (FOC l_2) for the Lagrange multipliers and plugging into equation (FOC b_1) yields:

$$\frac{u_2(c_1^\star, l_1^\star)}{u_2(c_2^\star, l_2^\star)} = \frac{\beta(1+R)f'(l_1^\star)}{f'(l_2^\star)}. \tag{6.6}$$

This is an Euler equation too, since it too relates marginal utilities in consecutive periods. This time, it relates the marginal utilities of labor.

We could analyze equations (6.5) and (6.6) in terms of the abstract functions, $u(\cdot)$ and $f(\cdot)$, but it is much simpler to assume particular functional forms and then carry out the analysis. Accordingly, assume:

$$u(c,l) = \ln(c) + \ln(1-l), \text{ and:}$$
$$f(l) = Al^\alpha.$$

Plugging the utility function into equation (6.5) yields:

$$\frac{c_2^\star}{c_1^\star} = \beta(1 + R),$$

just like in Chapter 3. All the analysis from that chapter carries forward. For example, this equation implies that a higher interest rate R implies that the household consumes more in period 2 relative to period 1. Equation (6.6) becomes:

$$\frac{(1 - l_2^\star)(l_2^\star)^{1-\alpha}}{(1 - l_1^\star)(l_1^\star)^{1-\alpha}} = \beta(1 + R). \tag{6.7}$$

Analysis of this equation is somewhat tricky. As a first step, let $D(l) = (1-l)l^{1-\alpha}$ be a helper function. Then equation (6.7) can be written as:

$$\frac{D(l_1^\star)}{D(l_2^\star)} = \beta(1 + R). \tag{6.8}$$

Now, let's consider how $D(l)$ changes when l changes:

$$\begin{aligned} D'(l) &= (1 - l)(\alpha - 1)l^{\alpha-2} + l^{\alpha-1}(-1) \\ &= l^{\alpha-2}(\alpha - \alpha l - 1 + l - l) \\ &= l^{\alpha-2}[\alpha(1 - l) - 1]. \end{aligned}$$

We know that $l^x > 0$ for all x, so $l^{\alpha-2} > 0$. Further, $\alpha(1 - l) < 1$, since l and α are both between zero and one. Putting these together, we find that $D'(l) < 0$, so increasing l causes $D(l)$ to decrease.

Now, think about what must happen to $l_1^\star$ and $l_2^\star$ in equation (6.8) if the interest rate R increases. That means that the right-hand side increases, so the left-hand side must increase in order to maintain the equality. There are two ways that the left-hand side can increase: either (i) $D(l_2^\star)$ increases, or (ii) $D(l_1^\star)$ decreases (or some combination of both). We already determined that $D(l)$ and l move in opposite directions. Hence, either $l_2^\star$ decreases or $l_1^\star$ increases (or some combination of both). Either way, $l_2^\star/l_1^\star$ decreases. The intuition of this result is as follows. A higher interest rate means the household has better investment opportunities in period 1. In order to take advantage of those, the household works relatively harder in period 1, so it earns more money to invest.

Exercises

Exercise 6.1 (Hard)
This economy contains 1,100 households. Of these, 400 own type-a farms, and the other 700 own type-b farms. We use superscripts to denote which type of farm. A household of type $j \in \{a, b\}$ demands (i.e., it hires) l_d^j units of labor, measured in hours. (The "d" is for

Variable	Definition
$U(\cdot)$	Overall utility
t	Time
c_t	Consumption at period t
l_t	Labor at period t
$u(\cdot)$	Period utility
β	Household's discount factor
y_t	Household's income in period t, in units of consumption
$f(l_t)$	Production function
P	Cost of a unit of consumption
R	Nominal interest rate
b_t	Number of dollars of bonds bought at period t
λ_t	Lagrange multiplier in period t
$\mathcal{L}$	Lagrangean
N	Number of households
$D(l)$	Helper function, to simplify notation

Table 6.2: Notation for Section 6.2

demand.) The type-j household supplies l_s^j of labor. (The "s" is for supply.) The household need not use its own labor on its own farm. It can hire other laborers and can supply its own labor for work on other farms. The wage per hour of work in this economy is w. This is expressed in consumption units, i.e., households can eat w. Every household takes the wage w as given. Preferences are:

$$u(c^j, l_s^j) = \ln(c^j) + \ln(24 - l_s^j),$$

where c^j is the household's consumption. Production on type-a farms is given by:

$$y = (l_d^a)^{0.5},$$

and that on type-b farms is:

$$y = 2(l_d^b)^{0.5}.$$

We are going to solve for the wage that clears the market. In order to do that, we need to determine demand and supply of labor as a function of the wage.

If an owner of a type-a farm hires l_d^a hours of labor at wage w per hour, the farm owner will make profit:

$$\pi^a = (l_d^a)^{0.5} - wl_d^a.$$

1. Use calculus to solve for a type-a farmer's profit-maximizing choice of labor l_d^a to hire as a function of the wage w. Call this amount of labor $l_d^{a\star}$. It will be a function of w. Calculate the profit of a type-a farmer as a function of w. Call this profit $\pi^{a\star}$.

2. If an owner of a type-b farm hires l_d^b hours of labor at wage w per hour, the farm owner will make profit:

$$\pi^b = 2(l_d^b)^{0.5} - wl_d^b.$$

   Repeat part 1 but for type-b farmers. Call a type-b farmer's profit-maximizing choice of labor $l_d^{b\star}$. Calculate the profit of a type-b farmer as a function of w. Call this profit $\pi^{b\star}$.

3. If a type-a farmer works l_s^a, then that farmer's income will be: $\pi^{a\star} + wl_s^a$. Accordingly, the budget constraint for type-a farmers is:

$$c^a = \pi^{a\star} + wl_s^a.$$

   A type-a household chooses its labor supply by maximizing its utility subject to its budget. Determine a type-a household's optimal choice of labor to supply for any given wage w. Call this amount of labor $l_s^{a\star}$.

4. Repeat part 3 but for type-b households. Call this amount of labor $l_s^{b\star}$.

5. Aggregate labor demand is just the sum of the demands of all the farm owners. Calculate aggregate demand by adding up the labor demands of the 400 type-a farmers and the 700 type-b farmers. This will be an expression for hours of labor l in terms of the market wage w. Call the result $l_d^\star$.

6. Aggregate labor supply is just the sum of the supplies of all the households. Calculate aggregate supply, and call it $l_s^\star$.

7. Use your results from parts 5 and 6 to solve for the equilibrium wage $w^\star$. (Set the two expressions equal and solve for w.)

Exercise 6.2 (Hard)
Consider an economy with many identical households. Each household owns a business that employs both capital (machinery) k and labor l_d to produce output y. (The "d" is for demand.) Production possibilities are represented by $y = Ak^{\frac{3}{10}}(l_d)^{\frac{7}{10}}$. The stock of capital that each household owns is fixed. It may employ labor at the prevailing wage w per unit of labor l_d. Each household takes the wage as given. The profit of each household from running its business is:

$$\pi = y - wl_d = Ak^{\frac{3}{10}}(l_d)^{\frac{7}{10}} - wl_d. \tag{6.9}$$

1. Determine the optimal amount of labor for each household to hire as a function of its capital endowment k and the prevailing wage w. Call this amount of labor $l_d^\star$.

2. Plug $l_d^\star$ back into equation (6.9) to get the maximized profit of the household. Call this profit $\pi^\star$.

3. Each household has preferences over its consumption c and labor supply l_s. These preferences are represented by the utility function: $u(c, l_s) = c^{\frac{1}{2}}(1 - l_s)^{\frac{1}{2}}$. Each household has an endowment of labor can be used in the household's own business or rented to others at the wage w. If the household supplies labor l_s, then it will earn labor income wl_s. Output, wages, and profit are all quoted in terms of real goods, so they can be consumed directly. Set up the household's problem for choosing its labor supply l_s. Write it in the following form:

$$\max_{\text{choices}} \{\text{objective}\} \quad \text{subject to: constraints}$$

4. Carry out the maximization from part 3 to derive the optimal labor supply $l_s^\star$.

5. Determine the equilibrium wage $w^\star$ in this economy.

6. How does the equilibrium wage $w^\star$ change with the amount of capital k owned by each household?

7. What does this model imply about the wage differences between the U.S. and Mexico? What about immigration between the two countries?

Chapter 8

Inflation

This chapter examines the causes and consequences of inflation. Sections 8.1 and 8.2 relate inflation to money supply and demand. Although the presentation differs somewhat from that in Barro's textbook, the results are similar. In Section 8.3 we extend Barro's analysis with a closer look at the real effects of inflation.

8.1 Money Supply and Demand

In most countries, the general level of prices tends to increase over time. This phenomenon is known as *inflation*. In this section we will link inflation to changes in the quantity of money in an economy.

The quantity of money is determined by money supply and demand. Before we can find out how supply and demand are determined, we have to make precise what exactly is meant by money. Money is defined as the *medium of exchange* in an economy. *Currency* (bank notes and coins) is a medium of exchange, but there are other commodities that fulfill this function as well. For example, deposits on checking accounts can be used as a medium of exchange, since a consumer can write a check in exchange for goods. There are other assets where it is not so clear whether they should be considered money or not. For example, savings deposits can be used as a medium of exchange by making transfers or withdrawals, but the main purpose of savings accounts is to serve as a store of value. In order to deal with these ambiguities, economists work with a number of different definitions of money. These definitions are often referred to as *monetary aggregates*. One of the most important monetary aggregates is called M1; this measure consists of the currency in circulation plus checking deposits at banks. Broader aggregates like M2 and M3 also contain savings and time deposits.[1] As a convention, in this chapter we will identify money with M1, although most of the analysis would also work if we had broader aggregates in mind.

Having defined money, let us turn to money supply. Since we use M1 as our definition of money, we have to find the determinants of the supply of currency and checking deposits. In most countries, the supply of currency is under control of the central bank. For example, in the United States the Federal Reserve is responsible for supplying currency. If the central bank decides to increase the supply of currency, all it needs to do is to print more bank notes and hand them out, most of the time to private banks. Conversely, the central bank can decrease the supply of currency by buying back its own money. The determination of the supply of checking deposits is a more difficult question. Even though the central bank does not directly control checking deposits at private banks, a number of monetary-policy instruments give the central bank indirect control over bank deposits. To explain exactly how this works is beyond the scope of the chapter. We will come back to this question in Chapter 17, which takes a closer look at central-bank policy and its relation to the banking industry. For the purposes of this chapter, we will simply assume that both currency and checking deposits are under direct control of the central bank. This approximation works well enough for a first analysis of inflation. From now on, we will use M_t to denote the overall quantity of money supplied by the central bank in year t. For convenience, we will measure M_t in dollars.

Let us now take a look at money demand. Money is demanded by households and firms. Households need money in order to purchase consumption goods. Firms need money to purchase inputs to production and to make change at cash registers. For a given year t, we will use Y_t to denote the total amount of purchases, measured in terms of consumption goods. For example, on Crusoe's island Y_t would be the number of coconuts consumed in year t. If we are thinking about a whole country, we can interpret Y_t as real GDP. Since Y_t is in terms of goods, we have to multiply it by the price level P_t to get the total amount of purchases in terms of dollars, P_tY_t. Actual money demand is lower than P_tY_t, because money can be used more than once in a year. The *velocity of money* is defined as the average number of times a piece of money turns over in a year. The more often money turns over, the less money is needed to carry out the planned purchases. Using V_t to denote velocity, actual money demand is given by P_tY_t/V_t. For example, if $V_t = 1$, then each unit of money will be used only once. This corresponds to a situation in which all purchases are carried out at the same time, so P_tY_t dollars will be needed. On the other hand, if each month only 1/12 of all purchases are made, only $P_tY_t/12$ dollars will be required, and V_t will be 12.

In equilibrium, money supply M_t and money demand P_tY_t/V_t have to be equal. If we set them equal and multiply by velocity V_t, we arrive at the *quantity equation*:

$$M_tV_t = P_tY_t.$$

The quantity equation relates the quantity of money M_t to the price level P_t. Still, as of now it does not provide an explanation for inflation, because we have not yet explained

[1]See Chapter 4 of Barro for precise definitions of these aggregates.

how velocity V_t and the the amount of purchases Y_t are determined.

8.2 The Quantity Theory

Our task is to add theoretical underpinnings to the quantity equation in order to better understand inflation. The best way to proceed would be to write down a model that explains how the decisions of optimizing agents determine velocity V_t and output Y_t. We will do that in the following section, but as a first step we will start with a simpler approach. We assume that velocity and output in each year are given constants that are determined independently of the money supply M_t and the price level P_t. Further, we assume that velocity does not change over time. Therefore we can drop the time subscript and use V to denote velocity. The central bank controls money supply M_t, so the price level P_t is the only free variable. Given these assumptions, the quantity equation implies that the central bank has perfect control over the price level. If the central bank changes money supply, the price level will change proportionally. We can see that by solving the quantity equation for P_t:

$$P_t = M_t V / Y_t. \tag{8.1}$$

Let us now see what this implies for inflation. The inflation rate π_t in a given year t is defined as the relative change in the price level from t to $t+1$, or:

$$\pi_t = \frac{P_{t+1} - P_t}{P_t}.$$

This can also be written as:

$$1 + \pi_t = \frac{P_{t+1}}{P_t}. \tag{8.2}$$

Taking the ratio of equation (8.1) for two consecutive years, we get:

$$\frac{P_{t+1}}{P_t} = \frac{M_{t+1} V Y_t}{M_t V Y_{t+1}}.$$

We know from equation (8.2) that P_{t+1}/P_t equals $1+\pi_t$, and the V terms cancel, so we have:

$$1 + \pi_t = \frac{M_{t+1}}{M_t} \frac{Y_t}{Y_{t+1}}. \tag{8.3}$$

We now take logs of both sides and use an approximation: $\ln(1+x) \approx x$ when x is not very large. Accordingly, equation (8.3) becomes:

$$\pi_t \approx [\ln M_{t+1} - \ln M_t] - [\ln Y_{t+1} - \ln Y_t].$$

This says that the inflation rate approximately equals the difference between the growth rate of money supply and the growth rate of output.[2] If output grows while the money

[2] See Chapter 1 for a general discussion of growth rates.

supply is constant, prices will have to fall so that money demand P_tY_t/V also stays constant. If money supply grows while output does not, prices will have to increase so that money demand increases in line with supply. Since the theory emphasizes the role of the quantity of money for the determination of inflation, it is known as the *quantity theory of money*.

Across countries and over time in a given country, we usually observe much higher variation in the growth rate of the money supply than in the rate of output growth. This indicates that variations in inflation are primarily attributable to variations in the rate of money growth. Empirical data gives strong support to this hypothesis. For example, Figure 7.1 in Barro shows that the money growth rate is almost perfectly proportional to the inflation rate in a sample of 80 countries.

While the quantity theory successfully explains the cause of inflation, it is not very helpful if we want to determine the consequences of inflation. In deriving the quantity theory, we assumed that money and prices were independent of all other variables in the economy. In the real world, high inflation is generally considered to be undesirable. If we want to understand why inflation might be bad, we have to determine the effects of inflation on real variables like output and consumption. This cannot be done within the quantity theory, since it assumed from the outset that such real effects did not exist. Instead, we need to go beyond the simplifying assumptions of the quantity theory.

To some degree we already did that in the discussion of money demand in Chapter 4, where we derived the optimal time T between a consumer's trips to the bank to get money. That time T between trips to get money was closely related to velocity V. In fact, $V = 2/T$.[3] In Chapter 4 we saw that the decision on T depended on the planned consumption expenditure and the nominal interest rate. Therefore the assumption of a constant velocity V that we made for the quantity theory was not correct. On the other hand, from an empirical point of view, the assumption of constant velocity seems to work relatively well as long as inflation rates are moderate.

The other assumption that we made for the quantity theory was that output Y_t was determined independently of monetary policy and inflation. We need to relax this assumption if we want to determine the real effects of inflation. In the next section, we will build a complete general equilibrium model that allows us to derive the impact of inflation on output and consumption.

[3]Velocity is given by $V = PY/M$. In Chapter 4, we derived that the average money holdings of a consumer were given by $\bar{m} = PcT/2$, where c was consumption. If we aggregate this over many consumers, the left-hand side becomes the aggregate money stock M, and individual consumption c sums to total output Y, so $M = PYT/2$. Plugging this into the formula for velocity yields $V = 2/T$.

8.3 A Cash-in-Advance Economy

In this section we derive the real effects of inflation. Unlike in the previous section, we will use a complete equilibrium model with optimizing consumers, because we want to understand how economic agents decide on consumption and output in the presence of inflation. The model builds on the general equilibrium framework developed in earlier chapters, but this model also contains a monetary sector.

This model is based on many identical consumers who live forever. In such a case, we say that consumers are *infinitely lived*. Since everyone is the same, it suffices to examine the choices of a single, representative consumer. The representative consumer has to decide on consumption c_t, labor supply l_t, savings s_{t+1}, and money holdings m_{t+1}. The utility function is:

$$\sum_{t=0}^{\infty} \beta^t [\ln(c_t) + \ln(1 - l_t)], \tag{8.4}$$

where β is a discount factor between zero and one. There is only one good in the economy, and the consumer can produce the good with the technology $y_t = l_t$, i.e., output equals labor input.

Monetary policy is conducted in a particularly simple way in this economy. There is no banking sector that intermediates between the central bank and consumers. Instead, the central bank hands out money directly to consumers. Monetary policy consists of printing money and giving it as a transfer τ_t to each consumer. When the central bank wants to contract the money supply, it taxes each consumer by making τ_t negative.

We will use R_t to denote the nominal interest rate on savings and P_t to denote the price of the consumption good in period t. The time-t budget constraint of the consumer is:

$$m_{t+1} + s_{t+1} = m_t + (1 + R_t)s_t + P_t l_t + \tau_t - P_t c_t. \tag{8.5}$$

On the left-hand side are the amounts of money and savings that the consumer carries into the next period. Therefore they are indexed by $t + 1$. On the right-hand side are all the receipts and payments during the period. The consumer enters the period with money m_t and savings plus interest $(1 + R_t)s_t$, both of which he carries over from the day before. During the day, the consumer also receives income from selling produced goods $P_t l_t$ and the transfer τ_t from the central bank. The only expenditures are purchases of the consumption good, $P_t c_t$. All funds that are left after the household purchases the consumption good are either invested in savings s_{t+1} or are carried forward as money m_{t+1}.

So far, there is no explanation for why the consumer would want to hold money. After all, savings earn interest, and money does not. In order to make money attractive, we assume that cash is required for buying the consumption good. The consumer cannot consume his own production and has to buy someone else's production in the market with cash. This introduces a new constraint faced by the consumer: expenditure on consumption goods has to be less than or equal to money holdings:

$$P_t c_t \leq m_t. \tag{8.6}$$

Since money that is to be used for buying consumption goods has to be put aside one period before it is spent, equation (8.6) is also called the *cash-in-advance constraint*, which explains the name of the model. From here on we will assume that equation (8.6) holds with equality. This will be the case as long as the nominal interest rate is positive, because in that situation it is more profitable to invest additional funds in savings instead of holding them as cash.

In this economy consumption equals output, so the cash-in-advance constraint aggregates up to be the quantity equation. This formulation implicitly assumes that velocity is one. A more sophisticated model would incorporate some version of the money-demand model of Chapter 4, allowing velocity to vary with inflation. However, such a model would be more complicated without adding much to our explanation of the real effects of monetary policy.

One way of understanding the cash-in-advance constraint is to think of the consumer as a family consisting of two members, a worker and a shopper. Each morning, the worker goes to his little factory, works, and sells the production to other consumers. Only late at night does the worker come home, so the income cannot be used for buying consumption goods that same day. The shopper also leaves each morning, taking the cash that was put aside the night before to do that day's shopping. Since the shopper does not see the worker during the day, only money that was put aside in advance can be used to make purchases.

The problem of the representative consumer is to maximize utility subject to the cash-in-advance constraint and the budget constraint:

$$\max_{\{c_t, l_t, s_{t+1}, m_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t [\ln(c_t) + \ln(1 - l_t)], \text{ subject to:}$$

$$P_t c_t = m_t, \text{ and:}$$

$$m_{t+1} + s_{t+1} = m_t + (1 + R_t)s_t + P_t l_t + \tau_t - P_t c_t.$$

In this model, the consumer's problem is much easier to analyze once we have the market-clearing conditions in place. Therefore we will complete the description of the economy first and derive the optimal decisions of the consumer later.

The next element of the model that needs to be specified is the monetary policy of the central bank. Instead of looking at aggregate money supply M_t, we will formulate monetary policy in terms of money per consumer m_t. This is merely a matter of convenience. We could recover the aggregate quantity of money by multiplying m_t by the number of consumers. However, since we are using a representative consumer, it is easier to formulate monetary policy on the level of individual consumers in the first place. We will assume a particularly simple policy: the central bank increases the money supply at a constant rate g. If the central bank wants to increase the money supply, it gives new cash to consumers.

Money supply in the next period is the sum of money supply in the current period and the transfer to the consumer. The money supply will grow at rate g if the transfer τ_t is given by $\tau_t = gm_t$, so we have:

$$m_{t+1} = m_t + \tau_t = (1 + g)m_t.$$

To close the model, we have to specify the three market-clearing conditions that must hold at each date t. The constraint for clearing the goods market states that consumption has to equal production:

$$c_t = l_t.$$

Clearing the credit market requires that total borrowing be equal to total savings. Since everyone is identical, there cannot be both borrowers and savers in the economy. In equilibrium savings have to be zero. Therefore the market-clearing constraint is:

$$s_t = 0.$$

In fact, we could omit savings from the model without changing the results. The only reason that we include savings is that this allows us to determine the nominal interest rate, which will play an important role in determining the real effects of monetary policy.

Finally, clearing the money market requires that the amount of cash demanded be the household equals the money supplied by the central bank. Since we use the same symbol m_t to denote money demand and supply, this market-clearing constraint is already incorporated in the formulation of the model.

An equilibrium for this economy is an allocation $\{c_t, l_t, s_t, m_t, \tau_t\}_{t=0}^{\infty}$ and a set of prices $\{P_t, R_t\}_{t=0}^{\infty}$ such that:

- Given the prices and transfers, $\{c_t, l_t, s_t, m_t\}_{t=0}^{\infty}$ is a solution to the household's problem; and
- All markets clear.

While this setup with infinitely-lived consumers might look complicated, having people live forever is actually a simplification that makes it easy to solve the model. The special feature of this framework is that the world looks the same in every period. The consumer always has infinitely many periods left, and the only thing that changes is the amount of money the consumer brings into the period (savings do not change since they are zero in equilibrium). The price level turns out to be proportional to the money stock, so the consumer always buys the same amount of the consumption good. In equilibrium, consumption c_t, labor l_t, and the nominal interest rate R_t are constant. Therefore we will drop time subscripts and denote interest by R and the optimal choices for consumption and labor by $c^\star$ and $l^\star$. Of course, we still need to show formally that $c^\star$, $l^\star$, and R are constant. This result will follow from the first-order conditions of the household's problem. We will plug in constants for consumption, labor, and interest, and we will be able to find prices such that the first-order conditions are indeed satisfied. For now, we just assume that $c^\star$ is constant.

As a first step in the analysis of the model, we examine the connection between monetary policy and inflation. This can be done in the same fashion as in the section on the quantity theory, without solving the consumer's problem explicitly.

The cash-in-advance constraint with constant consumption $c^\star$ is:

$$P_t c^\star = m_t. \tag{8.7}$$

The inflation rate π is defined by $1 + \pi = P_{t+1}/P_t$. Thus we can derive an equation for inflation by taking the ratio of the equation (8.7) for two consecutive periods:

$$1 + \pi = \frac{P_{t+1}}{P_t} = \frac{m_{t+1}}{m_t}.$$

Now we can use the fact that the money stock grows at a constant rate:

$$1 + \pi = \frac{m_{t+1}}{m_t} = \frac{m_t + \tau_t}{m_t} = \frac{(1+g)m_t}{m_t} = 1 + g.$$

Thus the inflation rate is equal to the growth rate of money supply. It is not surprising that we get this result. As in the quantity theory, we assume that velocity is constant. Since the cash-in-advance constraint is the quantity equation in this model, we had to come to the same conclusions as the quantity theory.

The main question that is left is how the level of consumption $c^\star$ (and hence equilibrium output) depends on inflation and monetary policy. To answer this question, we need to solve the household's problem.

We will use the Lagrangian method. The formulation of the Lagrangian differs from the one we used in the infinite-period model in Section 3.3, because here we multiply the Lagrange multipliers by the discount factor. This alternative formulation does not change results, and is mathematically more convenient. We use $\beta^t \mu_t$ for the multiplier on the time-t cash-in-advance constraint and $\beta^t \lambda_t$ as the multiplier on the time-t budget constraint. The Lagrangian for the household's problem is:

$$\begin{aligned}\mathcal{L} = \sum_{t=0}^{\infty} \beta^t [&\ln(c_t) + \ln(1 - l_t) + \mu_t(m_t - P_t c_t) \\ &+\lambda_t(m_t + (1 + R_t)s_t + P_t l_t + \tau_t - P_t c_t - m_{t+1} - s_{t+1})].\end{aligned}$$

The first-order conditions with respect to c_t, l_t, s_{t+1} and m_{t+1} are:

$$\text{(FOC } c_t\text{)} \qquad \beta^t \frac{1}{c_t^\star} - \beta^t(\mu_t + \lambda_t)P_t = 0;$$

$$\text{(FOC } l_t\text{)} \qquad -\beta^t \frac{1}{1 - l_t^\star} + \beta^t \lambda_t P_t = 0;$$

$$\text{(FOC } s_{t+1}\text{)} \qquad -\beta^t \lambda_t + \beta^{t+1}\lambda_{t+1}(1 + R_{t+1}) = 0; \text{ and:}$$

$$\text{(FOC } m_{t+1}\text{)} \qquad -\beta^t \lambda_t + \beta^{t+1}(\mu_{t+1} + \lambda_{t+1}) = 0.$$

We now guess that in equilibrium, consumption, labor, and interest are constants $c^\star$, $l^\star$, and R. If this guess were wrong, we would run into a contradiction later. (Take our word for it: this is not going to happen.) With consumption, labor, and interest being constants, the first-order conditions simplify to the following expressions:

$$\frac{1}{c^\star} = (\mu_t + \lambda_t)P_t, \tag{8.8}$$

$$\frac{1}{1 - l^\star} = \lambda_t P_t, \tag{8.9}$$

$$\lambda_t = \beta\lambda_{t+1}(1 + R), \text{and:} \tag{8.10}$$

$$\lambda_t = \beta(\mu_{t+1} + \lambda_{t+1}). \tag{8.11}$$

If we now solve equation (8.9) for λ_t and plug the result into equation (8.10), we get:

$$\frac{1}{(1 - l^\star)P_t} = \frac{1}{(1 - l^\star)P_{t+1}}\beta(1 + R), \text{ or:}$$

$$\frac{P_{t+1}}{P_t} = \beta(1 + R).$$

The left-hand side equals one plus the inflation rate. We determined already that the inflation rate is equal to the growth rate of money supply in this economy. Therefore we can express the nominal interest rate as:

$$1 + R = \frac{1 + \pi}{\beta} = \frac{1 + g}{\beta}. \tag{8.12}$$

This says that the nominal interest rate R moves in proportion to the growth rate g of money. Dividing the nominal interest rate by inflation yields the real interest rate r:[4]

$$1 + r = \frac{1 + R}{1 + \pi} = \frac{1}{\beta}.$$

This expression should look familiar. It is a version of the Euler equation (3.16) that we derived for in the infinite-period model of Chapter 3. In the model we are considering here, consumption is constant, so the marginal utilities drop out. To interpret this equation, keep

[4]See Barro, Chapter 4 for a discussion of real versus nominal interest rates.

in mind that there cannot be any borrowing in equilibrium because there is no one from whom to borrow. If β is low, then consumers are impatient. Therefore the interest rate has to be high to keep consumers from borrowing.

We still have to trace out the effect of inflation on consumption. By using equations (8.8) and (8.9), we can eliminate the multipliers from equation (8.11):

$$\frac{1}{(1-l^\star)P_t} = \beta \frac{1}{c^\star P_{t+1}}.$$

From the goods market-clearing constraint, we know that $c^\star = l^\star$. Therefore we get:

$$\frac{P_{t+1}}{P_t} = \beta \frac{1-c^\star}{c^\star}.$$

The left-hand side is equal to the inflation rate (which itself equals the money growth rate). We can use that fact to solve for $c^\star$:

$$1+g = \beta\frac{1-c^\star}{c^\star},$$
$$c^\star + gc^\star = \beta - \beta c^\star, \text{ so:}$$
$$c^\star = \frac{\beta}{1+g+\beta}. \tag{8.13}$$

This equation implies that consumption depends negatively on money growth, so consumption and inflation move in opposite directions. The intuition for this result is that inflation distorts the incentives to work. Income from labor cannot be used immediately for purchases of consumption, since consumption goods are bought with cash that has been put aside in advance. The labor income of today can be spent only tomorrow. When inflation is high, cash loses value over night. The higher inflation, the higher are prices tomorrow, and the fewer consumption goods can be bought for the same amount of labor. This implies that high rates of inflation decrease the incentive to work. Since consumption is equal to labor in equilibrium, consumption is low as well.

Given this relationship between consumption and inflation, which money growth rate should the central bank choose? In equilibrium, labor and consumption are equal. We can use this fact to find the optimal consumption, and then go backwards to compute the optimal money growth rate. The utility of consuming and working some constant $c = l$ forever is:[5]

$$\sum_{t=0}^{\infty} \beta^t[\ln(c) + \ln(1-c)] = \frac{1}{1-\beta}[\ln(c) + \ln(1-c)].$$

We will use $\hat{c}^\star$ to denote the optimal consumption. The first-order condition with respect to c is:

$$0 = \frac{1}{\hat{c}^\star} - \frac{1}{1-\hat{c}^\star}.$$

[5]Here we are using the formula for the sum of an infinite geometric series: $\sum_{n=0}^{\infty} a^n = 1/(1-a)$.

Solving for $\hat{c}^\star$ yields:

$$\hat{c}^\star = \frac{1}{2}. \tag{8.14}$$

Equation (8.13) gives us an expression for $\hat{c}^\star$ as a function of g. Combining this with equation (8.14) yields an equation involving the optimal rate of growth of the money stock $g^\star$:

$$\frac{1}{2} = \frac{\beta}{1 + g^\star + \beta}.$$

Solving this for $g^\star$ gives us:

$$g^\star = \beta - 1. \tag{8.15}$$

Since β is smaller than one, this equation tells us that $g^\star$ is negative: the optimal monetary policy exhibits shrinking money supply. Using equation (8.12) and our expression for $g^\star$, we can compute the optimal nominal interest rate:

$$1 + R = \frac{1 + g^\star}{\beta} = \frac{1 + (\beta - 1)}{\beta} = \frac{\beta}{\beta} = 1.$$

This implies $R = 0$, i.e., the nominal interest rate is zero. The intuition behind this result is as follows. The inefficiency in the model originates with the cash-in-advance constraint. The consumers are forced to hold an inferior asset, cash, for making purchases. If money were not needed for buying consumption goods and nominal interest rates were positive, everyone would save instead of holding cash. But if nominal interest rates were zero, cash and savings would earn the same return. Because prices fall in the equilibrium we calculated above, a consumer who holds money can buy more goods with this money in the future than he can buy now. This implies that the real interest rate on money is positive. Therefore incentives are not distorted if the nominal interest rate is zero. The recommendation of setting nominal interest rates to zero is known as the *Friedman rule*, after the Chicago economist Milton Friedman, who first came up with it. In Section 19.4, we will derive the Friedman rule once again within a different framework.

To summarize, the main outcomes of the cash-in-advance model are that: (1) the rate of money growth equals the inflation rate; (2) nominal interest rates move in proportion to inflation; and (3) output is negatively related to inflation. Empirical findings in the real world are consistent with these findings. The correlation of money growth and inflation was already addressed in the section on the quantity theory. Also, most of the variation in interest rates across countries can be explained by differences in inflation, which supports the second result. As to the third result, we observe that countries with very high inflation tend to do worse economically than countries with moderate inflation. However, within a set of countries with moderate inflation, the evidence is not conclusive.

There are a number of advanced issues concerning monetary policy and inflation that we will pick up later in this book. Chapter 18 is concerned with the coordination of monetary and fiscal policy, and in Chapter 19 we will return to the question of optimal monetary policy. While the prime emphasis of the cash-in-advance model is the inefficiency of holding cash instead of interest-bearing assets, Chapter 19 turns to the issue of expected versus unexpected inflation. You can think of the cash-in-advance model as describing the long-run consequences of expected inflation, while Chapter 19 considers the short-run consequences of a monetary policy that is not known in advance.

Variable	Definition
M_t	Aggregate quantity of money or cash
Y_t	Output
P_t	Price level
V	Velocity of money
π_t	Inflation rate
β	Discount factor of consumer
c_t	Consumption of consumer
l_t	Labor of consumer
$1 - l_t$	Leisure of consumer
m_t	Money or cash per consumer
s_t	Savings of consumer
τ_t	Central bank transfer of money to consumer
R_t	Nominal interest rate
r_t	Real interest rate
g	Growth rate of money supply

Table 8.1: Notation for Chapter 8

Exercises

Exercise 8.1 (Easy)
Consider an economy where velocity V equals 5, output grows at three percent a year, and money supply grows at five percent a year. What is the annual inflation rate?

Exercise 8.2 (Hard)
In the quantity theory, we assumed that velocity was constant. In reality, the velocity of money varies across countries. Would you expect countries with high inflation to have higher or lower velocity than low-inflation countries? Justify your answer. (Hint: You should draw both on Chapter 4 and Chapter 8 to answer this question.)

Chapter 9

Business Cycles

In this chapter we explore the causes of business cycles. Briefly, business cycles are the recurring fluctuations that occur in real GDP over time. For further descriptions of business cycles, refer to Barro's Chapter 9. Here, we concentrate on explaining business cycles. We begin with an overview of potential explanations. Then we work out a real business cycle model in detail.

While there are many different theories of business cycles, they share some properties. There is always a driving force behind economic fluctuations, some sort of shock or disturbance that is the original cause of the cycle. In addition, most theories build on a propagation mechanism that amplifies shocks. Unless the disturbances are already big enough by themselves to account for the fluctuations, there has to be some propagation mechanism that translates small, short-lived shocks into large, persistent economic fluctuations.

We will start our search for the cause of business cycles in Section 9.1 by listing a number of possible shocks and propagation mechanisms. Competing theories of the business cycle differ in which shocks and mechanisms they emphasize. In Section 9.2 we will concentrate on the real business cycle model, which is a straightforward extension of the market-clearing models that we developed in earlier chapters. Section 9.3 presents simulations for our real business cycle model and assesses the success of the model in matching real-world fluctuations.

9.1 Shocks and Propagation Mechanisms

Among the many shocks and disturbances that are present in an economy, only a few have received special attention in research on business cycles. Here are some of the more important candidates:

- **Technology shocks:** Real-world production functions change over time. New technologies like computers or robots alter the production process and raise overall productivity. Sometimes, production facilities break down or do not work as expected, so productivity falls. This technological change is not always smooth; it often comes in the form of shocks.

- **Weather shocks and natural disasters:** Many industries like agriculture or tourism are weather-dependent. Rainfall and sunshine influence the output of these sectors, so the weather is a a potential source of fluctuations. This is also true for disasters like earthquakes or landslides. El Niño is a shock of this kind that received a lot of attention lately. We can regard these kinds of shocks as one type of technology shock. Weather changes the production function for wheat, and an earthquake that wiped out, say, Silicon Valley, would change the production function for computers.

- **Monetary shocks:** We saw in Chapter 8 on inflation that there are real effects of monetary policy. Therefore random changes to money supply or interest rates are a potential source of fluctuations as well.

- **Political shocks:** The government influences the economy both directly through government enterprises and indirectly through regulation. Changes in tax laws, antitrust regulation, government expenditure and so on are a potential source of disruption in the economy.

- **Taste shocks:** Finally, it is also conceivable that shifts in preferences cause fluctuations. Fashion and fads change rapidly, and they may cause fluctuations in areas like the apparel, music, or movie industries.

While the shocks just mentioned are present to some degree in every economy, they are probably not large enough to serve as a direct explanation of business cycles. For example, in the United States real GDP fell by 2.8% between October 1981 and 1982. It is hard to imagine any shock that caused a direct output loss of almost 3% of GDP within only a year, and if there was one, we would probably be aware of it. It appears more likely that there are mechanisms present in the economy that amplify shocks and propagate them through time. Here are some candidates:

- **Intertemporal substitution:** Shocks that have a negative impact on productivity lower the marginal return to labor and other factors of production. If marginal products fall, consumer's might prefer to work less and consume leisure instead. Labor input would fall, which amplifies the negative impact on output. At the same time, since consumers prefer a smooth consumption profile they might prefer to lower savings for some time when a shock hits. On an aggregate level, this leads to lower investment and a lower capital stock in the future. Therefore a short-lived shock may have an impact in the future as well.

- **Sticky prices:** Market economies react to changes with price adjustments. For example, a negative productivity shock lowers the marginal product of labor, so that the real wage would have to move downward to adjust labor demand and supply. But if wages are inflexible for some reason, the adjustment cannot take place. The result is unemployment and an output loss that is larger than the direct effect of the shock. Similar effects arise if goods prices are sticky.

- **Frictions in financial sector:** Even small shocks can force the firms the are hit directly into bankruptcy. This will affect other firms and banks that lent money to the now bankrupt firms. Often additional firms have to declare bankruptcy, and sometimes even banks fail. Bank failures affect all creditors and debtors and therefore can have large economic consequences. Serious economic crises are often accompanied and amplified by series of bank failures. Examples are the great depression and the current Asian crisis.

Business cycle models can be broadly subdivided into two categories. Some theories regard cycles as a failure of the economic system. Because of frictions or imperfections of the market mechanism, the economy experiences depressions and fails to achieve the efficient level of output and employment. Models of this kind often rely on financial frictions, sticky prices, or other adjustment failures as the propagation mechanism. Both technology shocks and monetary shocks are considered to be important sources of fluctuations. The Keynesian model of output determination[1] falls into this category.

On the other hand, there is a class of models that regards business cycles as the optimal reaction of the economy to unavoidable shocks. Shocks are propagated through intertemporal substitution within an efficient market mechanism. Technology shocks are considered to be the main cause of economic fluctuations. Models of this kind are often referred to as *real business cycle* models.[2]

We can be fairly certain that there is some truth to both views of economic fluctuations. Major economic breakdowns like the great depression or the recent Asian crisis appear to be closely connected to disruptions in the financial sector. Bank failures and financial instability played an important role in both cases.

On the other hand, most business cycles are far less severe than the great depression or the Asian crisis. In the entire post-war history of the United States and the Western European countries there is not a single depression that caused an output loss similar to the one suffered between 1929 and 1933. The question is whether normal business cycles are caused by the same kind of frictions that caused the great depression. The Keynesian model with its emphasis on slow adjustments and sticky prices supports this view. Real business cycle theorists argue that breakdowns like the great depression are a phenomenon distinct from usual business cycles, and that usual cycles can be explained as the optimal reaction of an efficient market system to economic shocks.

[1] See Barro, Chapter 20.

[2] The term derives from the fact that shocks in real business cycle theory are real, as opposed to monetary, and that sluggish nominal adjustment does not play a role as a propagation mechanism.

In this chapter, we will primarily look for explanations for normal-scale business cycles, like those experienced in the United States since World War II. How can we determine whether such cycles are small-scale failures of the economic system rather than simply the markets' efficient reactions to shocks? A natural way to answer this question is to build a number of model economies that include alternative propagation mechanisms, expose the model economies to shocks, and see whether the outcomes look like real-world business cycles. This is exactly the road that has been taken by real business cycle theorists. They have taken standard equilibrium models as a point of departure and exposed them to productivity shocks. As it turns out, models of this kind are quite successful at explaining real-world business cycles. We will now take a closer look at such a real business cycle model.

9.2 A Real Business Cycle Model

Real business cycle models are straightforward extensions of equilibrium models of the kind that we use throughout this course. In most cases, the models feature infinitely lived consumers, and business cycles are generated by random disturbances to production possibilities. Unfortunately, solving that kind of model is difficult. Often no explicit solution is available, so numerical approximations have to be used. To keep the presentation tractable, in this chapter we will use a simpler framework in which people live for two periods only. The model does not fit the facts as well as a full-scale real business cycle model, but it serves its purpose as a simple illustration of the main ideas of real business cycle theory.

In the model world there is a sequence of overlapping generations. Each period a new generation of consumers is born, and each consumer lives for two periods. We will sometimes refer to the periods as years, and for simplicity we assume that exactly one consumer is born each year. People work in the first period when they are young. In the second period they are retired and live on savings. Throughout the model, superscripts refer to the year when a person was born, while subscripts refer to the actual year. For example, c_t^t is the period-t consumption of a consumer who was born in year t, so such a consumer would be young in period t. Similarly, c_{t+1}^t is the consumption of the same consumer in period $t+1$, when he is old. The consumers do not care about leisure. A consumer born in year t has the following utility function:

$$u(c_t^t, c_{t+1}^t) = \ln(c_t^t) + \ln(c_{t+1}^t).$$

We could introduce a discount factor, but for simplicity we assume that the consumers value both periods equally. Note that at each point of time there are exactly two people around: one who was just born and is young, and another who was born the year before and is now retired. In each period the young person supplies one unit of labor and receives wage income w_t. The labor supply is fixed, since consumers do not care about leisure. The wage income can be used as savings k_t and as consumption c_t^t. The budget constraint of a young worker is:

$$c_t^t + k_t = w_t,$$

i.e., consumption plus savings equals income from labor. In period $t+1$ the consumer born in t is old and retired. The old consumer lends his savings k_t to the firm. The firm uses the savings as capital and pays return r_{t+1} to the old consumer. A fraction δ of the capital wears out while being used for production and is not returned to the consumer. δ is a number between zero and one, and is referred to as the *depreciation rate*. The budget constraint for the retirement period is:

$$c_{t+1}^t = (1 - \delta + r_{t+1})k_t,$$

i.e., consumption equals the return from savings.

The household born in period t maximizes utility subject to the budget constraints, and takes prices as given:

$$\max_{c_t^t, c_{t+1}^t, k_t} \left\{ \ln(c_t^t) + \ln(c_{t+1}^t) \right\}, \text{ subject to:}$$
$$c_t^t + k_t = w_t, \text{ and:}$$
$$c_{t+1}^t = (1 - \delta + r_{t+1})k_t.$$

We can use the constraints to eliminate consumption and write this as:

$$\max_{k_t} \left\{ \ln(w_t - k_t) + \ln((1 - \delta + r_{t+1})k_t) \right\}.$$

This is similar to the problem of the consumer in the two-period credit market economy that we discussed in Section 3.2. From here on we will drop the practice of denoting optimal choices by superscripted stars, since the notation is already complicated as it is. The first-order condition with respect to k_t is:

$$0 = -\frac{1}{w_t - k_t} + \frac{1 - \delta + r_{t+1}}{(1 - \delta + r_{t+1})k_t}.$$

Solving this for k_t yields:

$$k_t = \frac{w_t}{2}. \tag{9.1}$$

Thus, regardless of the future return on capital, the young consumer will save half of his labor income. Again, this derives from the fact that wealth and substitution effects cancel under logarithmic preferences. This feature is is helpful in our setup. Since there will be productivity shocks in our economy and r_{t+1} depends on such shocks, the consumer might not know r_{t+1} in advance. Normally we would have to account for this uncertainty explicitly, which is relatively hard to do. In the case of logarithmic utility, the consumer does not care about r_{t+1} anyway, so we do not have to account for uncertainty.

Apart from the consumers, the economy contains a single competitive firm that produces output using capital k_{t-1} and labor l_t. Labor is supplied by the young consumer, while the supply of capital derives from the savings of the old consumer.The rental rate for capital is r_t, and the real wage is denoted w_t. The production function has constant returns to scale and is of the Cobb-Douglas form:

$$f(l_t, k_{t-1}) = A_t l_t^{\alpha} k_{t-1}^{1-\alpha}.$$

Here α is a constant between zero and one, while A_t is a productivity parameter. A_t is the source of shocks in this economy. We will assume that A_t is subject to random variations and trace out how the economy reacts to changes in A_t. The profit-maximization problem of the firm in year t is:

$$\max_{l_t, k_{t-1}} \left\{ A_t l_t^{\alpha} k_{t-1}^{1-\alpha} - w_t l_t - r_t k_{t-1} \right\}.$$

The first-order conditions with respect to l_t and k_{t-1} are:

(FOC l_t) $$A_t \alpha l_t^{\alpha-1} k_{t-1}^{1-\alpha} - w_t = 0; \text{ and:}$$

(FOC k_{t-1}) $$A_t (1-\alpha) l_t^{\alpha} k_{t-1}^{-\alpha} - r_t = 0.$$

Using the fact that the young worker supplies exactly one unit of labor, $l_t = 1$, we can use these first-order conditions to solve for the wage and return on capital as a function of capital k_{t-1}:

$$w_t = A_t \alpha k_{t-1}^{1-\alpha}; \text{ and:} \tag{9.2}$$

$$r_t = A_t (1-\alpha) k_{t-1}^{-\alpha}. \tag{9.3}$$

Since the production function has constant returns, the firm does not make any profits in equilibrium. We could verify that by plugging our results for w_t and r_t back into the firm's problem. Note that the wage is proportional to the productivity parameter A_t. Since A_t is the source of shocks, we can conclude that wages are procyclical: when A_t receives a positive shock, wages go up. Empirical evidence suggests that wages in the real world are procyclical as well.

To close the model, we have to specify the market-clearing constraints for goods, labor, and capital. At time t the constraint for clearing the goods market is:

$$c_t^t + c_t^{t-1} + k_t = A_t l_t^{\alpha} k_{t-1}^{1-\alpha} + (1-\delta) k_{t-1}.$$

On the left hand side are goods that are used: consumption c_t^t of the currently young consumer, consumption c_t^{t-1} of the retired consumer who was born in $t-1$, and savings k_t of the young consumer. On the right hand side are all goods that are available: current production and what is left of the capital stock after depreciation.

The constraint for clearing the labor market is $l_t = 1$, since young consumers always supply one unit of labor. To clear the capital market clearing we require that capital supplied by the old consumer be equal to the capital demanded by the firm. To save on notation, we use the same symbol k_{t-1} both for capital supplied and demanded. Therefore the market-clearing for the capital market is already incorporated into the model and does not need to be written down explicitly.

In summary, the economy is described by: the consumer's problem, the firm's problem, market-clearing conditions, and a random sequence of productivity parameters $\{A_t\}_{t=1}^{\infty}$. We assume that in the very first period there is already an old person around, who somehow fell from the sky and is endowed with some capital k_0.

Given a sequence of productivity parameters $\{A_t\}_{t=1}^{\infty}$, an equilibrium for this economy is an allocation $\{c_t^t, c_t^{t-1}, k_{t-1}, l_t\}_{t=1}^{\infty}$ and a set of prices $\{r_t, w_t\}_{t=1}^{\infty}$ such that:

- Given prices, the allocation $\{c_t^t, c_t^{t-1}, k_{t-1}, l_t\}_{t=1}^{\infty}$ gives the optimal choices by consumers and firms; and
- All markets clear.

We now have all pieces together that are needed to analyze business cycles in this economy. When we combine the optimal choice of savings of the young consumer (9.1) with the expression for the wage rate in equation (9.2), we get:

$$k_t = \frac{1}{2} A_t \alpha k_{t-1}^{1-\alpha}. \tag{9.4}$$

This equation shows how a shock is propagated through time in this economy. Shocks to A_t have a direct influence on k_t, the capital that is going to be used for production in the next period. This implies that a shock that hits today will lead to lower output in the future as well. The cause of this is that the young consumer divides his income equally between consumption and savings. By lowering savings in response to a shock, the consumer smoothes consumption. It is optimal for the consumer to distribute the effect of a shock among both periods of his life. Therefore a single shock can cause a cycle that extends over a number of periods.

Next, we want to look at how aggregate consumption and investment react to a shock. In the real world, aggregate investment is much more volatile than aggregate consumption (see Barro's Figure 1.10). We want to check whether this is also true in our model. First, we need to define what is meant by aggregate consumption and investment. We can rearrange the market-clearing constraint for the goods market to get:

$$c_t^t + c_t^{t-1} + k_t - (1-\delta)k_{t-1} = A_t l_t^{\alpha} k_{t-1}^{1-\alpha}.$$

On the right-hand side is output in year t, which we are going to call Y_t. Output is the sum of aggregate consumption and investment. Aggregate consumption C_t is the sum of the consumption of the old and the young person, while aggregate investment I_t is the difference between the capital stock in the next period and the undepreciated capital in this period[3]:

$$\underbrace{c_t^t + c_t^{t-1}}_{C_t} + \underbrace{k_t - (1-\delta)k_{t-1}}_{I_t} = \underbrace{A_t l_t^{\alpha} k_{t-1}^{1-\alpha}}_{Y_t}.$$

Consumption can be computed as the difference between output and investment. Using equation (9.4) for k_t yields:

$$\begin{aligned} C_t = Y_t - I_t &= A_t k_{t-1}^{1-\alpha} + (1-\delta)k_{t-1} - k_t \\ &= A_t k_{t-1}^{1-\alpha} + (1-\delta)k_{t-1} - \frac{1}{2}A_t \alpha k_{t-1}^{1-\alpha} \\ &= \left(1 - \frac{1}{2}\alpha\right) A_t k_{t-1}^{1-\alpha} + (1-\delta)k_{t-1}. \end{aligned} \tag{9.5}$$

Aggregate investment can be computed as output minus aggregate consumption. Using equation (9.5) for aggregate consumption yields:

$$\begin{aligned} I_t = Y_t - C_t &= A_t k_{t-1}^{1-\alpha} - \left(1 - \frac{1}{2}\alpha\right) A_t k_{t-1}^{1-\alpha} - (1-\delta)k_{t-1} \\ &= \frac{1}{2}A_t \alpha k_{t-1}^{1-\alpha} - (1-\delta)k_{t-1}. \end{aligned} \tag{9.6}$$

We are interested in how C_t and I_t react to changes in the technology parameter A_t. We will look at relative changes first. The *elasticity* of a variable x with respect to another variable y is defined the percentage change in x in response to a one percent increase in y. Mathematically, elasticities can be computed as $\frac{\partial x}{\partial y}\frac{y}{x}$. Using this formula, the elasticity of consumption with respect to A_t is:

$$\frac{\partial C_t}{\partial A_t}\frac{A_t}{C_t} = \frac{\left(1 - \frac{1}{2}\alpha\right) A_t k_{t-1}^{1-\alpha}}{\left(1 - \frac{1}{2}\alpha\right) A_t k_{t-1}^{1-\alpha} + (1-\delta)k_{t-1}} < 1,$$

and for investment we get:

$$\frac{\partial I_t}{\partial A_t}\frac{A_t}{I_t} = \frac{\frac{1}{2}A_t \alpha k_{t-1}^{1-\alpha}}{\frac{1}{2}A_t \alpha k_{t-1}^{1-\alpha} - (1-\delta)k_{t-1}} > 1.$$

It turns out that the relative change in investment is larger. A one-percent increase in A_t leads to an increase of more than one percent in investment and less than one percent in consumption. Investment is more volatile in response to technology shocks, just as real-world investment is. Of course, to compare the exact size of the effects we would have to specify the parameters, like α and δ, and to measure the other variables, like k_t.

[3]More precisely, I_t in the model is *gross investment*, which includes replacement of depreciated capital. The net difference between capital tomorrow and today $k_t - k_{t-1}$ is referred to as *net investment*.

If we look at absolute changes instead of relative changes, the results are less satisfactory. The absolute change is higher in consumption than in investment, while in the real world it is the other way around. This failure of the model derives from the fact that people are too short-lived. In real business cycle models, the smaller variations in consumption relative to investment result from consumers trying to smooth their consumption. In our model, the possibilities for smoothing are rather limited. The old person has no more time left and therefore cannot smooth at all, while the young person has only one more year to go. Therefore a comparatively large fraction of the shock shows up in consumption. In more-advanced real business cycle models with infinitely lived consumers, the absolute changes in consumption are much smaller than the absolute changes in investment.

9.3 Simulations

We can get an even better impression of the business cycle in our model by simulating the economy. This means that we specify all parameters, start at some initial capital stock, and generate a series of random shocks. We can use the solutions to the model to compute consumption, investment, output, and the capital stock in the economy for any number of periods. Then we can compare the results to real-world business cycles.

There are only two parameters to be specified in the model, α and δ. Our choices are $\alpha = .7$ and $\delta = .05$. The choice for α matches the labor share in the economy to real world data[4], while the value for δ is an estimate of the actual average depreciation rate in an industrialized economy. The initial capital stock k_1 was set to .22. The productivity parameter was generated by:

$$A_t = \bar{A} + \epsilon_t.$$

Here $\bar{A}$ is the average level of productivity, while the ϵ_t are random shocks. We set $\bar{A} = 1$. The ϵ_t where generated by a computer to be independent over time and uniformly distributed on the interval $[-.1,\ .1]$. Thus the shocks can change productivity by up to ten percent upward or downward.

Figure 9.1 shows the reactions to a single productivity shock of five percent. That is, in the first period A_t is equal to its average, $A_1 = 1$. In the second period the shock hits, $A_2 = 1.05$. From then on, A_t is back to one and stays there. We can see that even this single shock has an impact that can be felt for a long period of time. Figure 9.1 shows the absolute deviations of consumption, investment, and capital from their average values. It takes about eight periods until all variables are back to their average. In the second period, when the shock takes place, both consumption and investment are up. In period 3 the capital stock is higher because of the higher investment in period 2. At the same time, investment falls. Consumption is higher than average because the capital stock is

[4]The *labor share* in an economy is defined to be total wages as a fraction of output. See Chapter 11 to see why α is equal to the labor share.

higher, even though productivity is back to normal again. From then on, all variables slowly return to their average values. Note that from period 4 on no one is alive anymore who was present when the shock took place. The higher investment in the period of the shock has increased the capital stock, and the effects of that can be felt for a long time. Thus even a single shock has long-run effects, and investment goes through a full cycle in response to this shock.

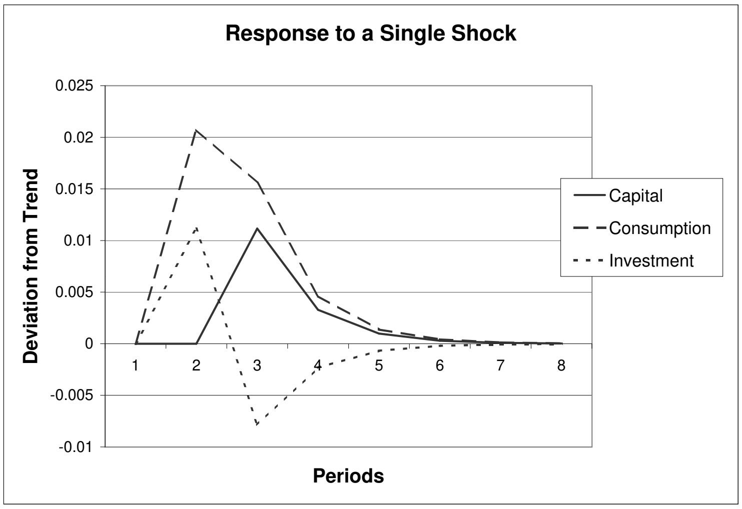

Figure 9.1: Response to a Five-Percent Productivity Shock

Figure 9.2 shows the same information as Figure 9.1, but variables are divided by their mean so that we can see the relative changes. Investment is by far the most volatile series. Compared to investment, the changes in capital and consumption are hardly visible.

By looking at a single shock, we were able to examine the propagation mechanism in isolation and to get an impression of the relative volatility of consumption and investment. But if we want to compare the model outcomes to real-world business cycles, we need to generate a whole series of shocks. Figure 9.3 shows such a simulation for our model economy. The combined effects of many shocks cause an outcome that looks similar to real-world business cycles. There are booms and depressions, the cycles vary in length within a certain interval, and investment is more volatile than consumption.

Our simple business cycle model is quite successful in emulating a number of business-cycle facts. Shape, length, and amplitude of business cycles are comparable to real-world data, investment is relatively more volatile than consumption, and the wage is procyclical. More-advanced real business cycle models are even better in matching the facts. By introducing variable labor supply we can generate procyclical employment. Using infinitely lived consumers would get the absolute changes in consumption and investment right.

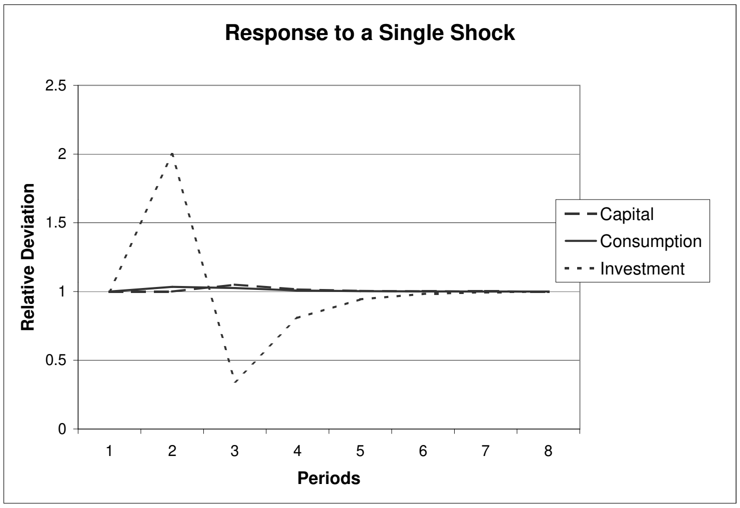

Figure 9.2: Relative Changes in Response to a Five-Percent Productivity Shock

State-of-the-art real business cycle models match most business cycle facts, and when fed with measured productivity shocks, they generate cycles that explain about 70% of the size of actual business cycles.

This success has led some researchers to the conclusion that business cycles are exactly what standard economic theory predicts. In the presence of shocks to production possibilities, optimal adjustments of households and firms within an efficient market system generate just the pattern of fluctuations that is observed in the real world. From this perspective, business cycles are no miracle at all. We would be surprised if there were no business cycles!

Even though technology shocks combined with efficient markets appear to provide a convincing explanation for business cycles, it cannot be ruled out that other shocks or propagation mechanisms also play a role. After all, real business cycle theory does not account for 100% of the amplitude of actual business cycles, so there have to be other factors as well. Other types of shocks can be analyzed within the real business cycle framework. There are also a number of models that emphasize other propagation mechanisms. The Keynesian model of output determination is the most prominent example[5] but models that combine monetary shocks with frictions in the financial sector have also received a lot of attention lately. However, so far none of these models matches the ability of the real business cycle model to mimic actual economic fluctuations.

[5]See Barro's Chapter 20 for a discussion of that model. For a critique of the real business cycle model from a Keynesian perspective, try the following article: Mankiw, N. Gregory. 1989. "Real Business Cycles: A New Keynesian Perspective". *Journal of Economic Perspectives* 3(3): 79-90.

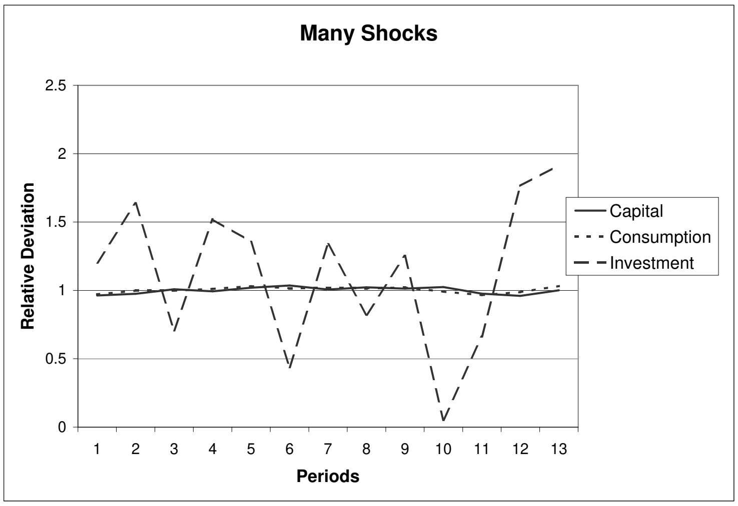

Figure 9.3: Capital, Consumption, and Investment with Many Shocks

Variable	Definition
c_t^t	Consumption of generation t when young
c_{t+1}^t	Consumption of generation t when old
$u(\cdot)$	Utility function
k_t	Capital saved in t and used in $t+1$
l_t	Labor
w_t	Wage
r_t	Rental rate of capital
$f(\cdot)$	Production parameter
A_t	Productivity parameter
α	Parameter in the production function
δ	Depreciation rate
C_t	Aggregate consumption
I_t	Aggregate investment
Y_t	Aggregate output
ϵ	Random shock

Table 9.1: Notation for Chapter 9

Exercises

The following exercises make up a project that can be done in groups or individually.

Exercise 9.1 (Moderate)
As the word "cycle" indicates, for a long time economists thought of business cycles as regular, recurrent events. The length and severity of business cycles was thought to be mostly constant. For example, the typical length of one full cycle (from boom through recession back to boom) was supposed to be between four and seven years. In this question you will examine the actual business cycles of a country of your choice and examine whether they seem to follow a regular pattern.

The first thing to do is to get the necessary data. Business cycles are roughly defined as deviations of real GDP from trend. Therefore you will need to acquire data on real GDP for some country. A good source is the Penn World Tables, a set of standardized measures of economic activity for most countries in the world. You can access the World Tables through a website at the University of Toronto. The address is:

`http://arcadia.chass.utoronto.ca/pwt/`

Once you are there, select "Alphabetical List of Topics", then "Real GDP per capita in constant dollars using chain index", then click on the country of your choice (not the United States), then use the "Submit Query" button to get the data. Load the data into a spreadsheet, and you are ready to go.

The first step is to compute the trend component of GDP. Good methods for computing the trend of a time series require a relatively high amount of complicated computations. Therefore we will offer you an ad hoc, quick-and-dirty method of computing the trend. Once we get to the business cycles, it turns out that this method works sufficiently well for our purposes. We will use GDP_t to denote real GDP at time t. The computation of the trend proceeds in steps:

- Compute the growth rate of GDP for each year. In terms of your spreadsheet, let us assume that column A is year and column B is real GDP. The first year is in row one. Now you can put the growth rates into column C. Put the growth rate from year 1 to 2 into cell C1, and so on.

- From now on, we are going to apply a method called *exponential smoothing* to get smooth versions of our data. Assume you want to get a smooth version of a times series x_t. Let us call the smooth version $\hat{x}_t$. Basically, the $\hat{x}_t$ are computed as a forecast based on past observations of x_t. The first $\hat{x}_t$ is set equal to the first x_t: $\hat{x}_1 = x_1$. From then on, the forecasts for the next period are computed as an average of the last forecast and the actual value: $\hat{x}_{t+1} = \beta\hat{x}_t + (1 - \beta)x_t$, where β is a number between zero and one. If you plug this formula recursively into itself, you will see that each $\hat{x}_t$ is a weighted average of past x_t.

Let us now put a smooth growth rate into column D. Since $\hat{x}_1 = x_1$, the first smooth value is equal to the original value: D1=C1. For the next value, we apply the smoothing formula. We recommend that you set β to .5: D2=.5*D1+.5*C1. In the same way, you can get the other smoothed growth rates. For future reference, We will call the smooth growth rates $\hat{g}_t$.

- In the next step, we are going to apply the same method to real GDP, but additionally we will use the smooth growth rates we just computed. This smooth real GDP is the trend we are looking for, and we will place it in column E. As before, in the first year the smooth version is identical to the original one: $\text{Trend}_1 = \text{GDP}_1$, thus E1=B1. From then on, we get the trend in the next period by averaging between the trend and the actual value (as before), but also applying the smooth growth rates we just computed. If we do not do that, our trend will always underestimate GDP. From year two on the formula is therefore:

$$\text{Trend}_{t+1} = (1 + \hat{g}_t)(0.5)\text{Trend}_t + (0.5)\text{GDP}_t.$$

  In terms of the spreadsheet, this translates into E2=(1+D1)*(.5*E1+.5*B1), and so on.

This completes the computation of the trend. Plot a graph of GDP and its trend. If the trend does not follow GDP closely, something is wrong. (Document your work, providing spreadsheet formulas, etc.)

Exercise 9.2 (Moderate)
Now we want to see the cyclical component of GDP. This is simply the difference between GDP and its trend. Because we are interested in relative changes, as opposed to absolute changes, it is better to use log-differences instead of absolute differences. Compute the cyclical component as ln(GDP) − ln(Trend). Plot the cyclical component. You will see the business cycles for which we have been looking. (Document your work, providing spreadsheet formulas, etc.)

Exercise 9.3 (Easy)
Now we will examine the cycles more closely. Define "peak" by a year when the cyclical component is higher than in the two preceding and following years. Define "cycle" as the time between two peaks. How many cycles do you observe? What is the average length of the cycle? How long do the shortest and the longest cycles last? Do the cycles look similar in terms of severity (amplitude), duration, and general shape? (Document your work, providing spreadsheet formulas, etc.)

Exercise 9.4 (Moderate)
Having seen a real cycle, the next step is to create one in a model world. It turns out that doing so is relatively hard in a model with infinitely lived agents. There we have to deal with uncertainty, which is fun to do, but it is not that easy as far as the math is concerned.

Therefore our model world will have people living for only one period. In fact, there is just one person each period, but this person has a child that is around in the next period, and so on. The person, let us call her Jill, cares about consumption c_t and the bequest of capital k_{t+1} she makes to her child, also named Jill. The utility function is:

$$\ln(c_t) + A\ln(k_{t+1}),$$

where $A > 0$ is a parameter. Jill uses the capital she got from her mother to produce consumption c_t and investment i_t, according to the resource constraint:

$$c_t + i_t = \sqrt{Bk_t} + \epsilon_t,$$

where $B > 0$ is a parameter, and ϵ_t a random shock to the production function. The shock takes different values in different periods. Jill knows ϵ_t once she is born, so for her it is just a constant. The capital that is left to Jill the daughter is determined by:

$$k_{t+1} = (1 - \delta)k_t + i_t,$$

where the parameter δ, the depreciation rate, is a number between zero and one. This just means that capital tomorrow is what is left over today after depreciation, plus investment.

Compute Jill's decision of consumption and investment as a function of the parameters k_t, and ϵ_t.

Exercise 9.5 (Moderate)
If we want to examine the behavior of this model relative to the real world, the next step would be to set the parameters in a way that matches certain features of the real world. Since that is a complicated task, we will give some values to you. B is a scale parameter and does not affect the qualitative behavior of the model. Therefore we set it to $B = .1$. δ is the depreciation rate, for which a realistic value is $\delta = .05$. A determines the relative size of c_t and k_t in equilibrium. A rough approximation is $A = 4$. Using these parameters, compare the reactions of c_t and i_t to changes in ϵ_t. (Use calculus.)

Exercise 9.6 (Moderate)
In the last step, you will simulate business cycles in the model economy. All you need to know is the capital k_1 at the beginning of time and the random shocks ϵ_t. As a starting capital, use $k_1 = 3.7$. You can generate the random shocks with the random number generator in your spreadsheet. In Excel, just type "=RAND()", and you will get a uniformly distributed random variable between zero and one. Generate 50 such random numbers, and use your formulas for c_t and i_t and the equation for capital in the next period, $k_{t+1} = (1 - \delta)k_t + i_t$, to simulate the economy. Plot consumption and investment (on a single graph). How does the volatility of the two series compare? Plot a graph of GDP, that is, consumption plus investment. How do the business cycles you see compare with the ones you found in the real world? You don't need to compute the length of each cycle, but try to make some concrete comparisons.

Exercise 9.7 (Easy)
Read the following article: Plosser, Charles. 1989. "Understanding Real Business Cycles". *Journal of Economic Perspectives* 3(3): 51-78. Plosser is one of the pioneers of real business cycle theory. What you have done in the previous exercises is very similar to what Plosser does in his article. His economy is a little more realistic, and he gets his shocks from the real world, instead of having the computer draw random numbers, but the basic idea is the same.

Describe the real business cycle research program in no more than two paragraphs. What question is the theory trying to answer? What is the approach to answering the question?

Exercise 9.8 (Moderate)
What does Plosser's model imply for government policy? Specifically, can the government influence the economy, and is government intervention called for?

Chapter 10

Unemployment

The study of unemployment is usually cast as the study of workers. Several theories seek to explain why the labor market might not clear at a particular wage. Among these are "search" models, in which unemployed people are in the process of looking for work. One such model is presented in Chapter 10 of the Barro textbook. More-sophisticated theories attempt to explain unemployment as the breakdown in a matching process between workers and jobs. Public discussions of unemployment often conflate the two.

In this chapter we will discuss some exciting new research on the statistical characteristics of jobs and employment in the United States. We will not attempt to provide theoretical explanations for the observed statistical patterns; rather, we will concentrate on the statistics themselves. The primary source for the material in this chapter is *Job Creation and Destruction*, by Davis, Haltiwanger, and Schuh.[1] Hereafter, we will refer to this book simply as "DHS". The book synthesizes research based on some important data sets regarding jobs and employment. This chapter can provide only a very broad outline of the book, and the interested reader is strongly encouraged to obtain his or her own copy. The book is short, accessible, and every page contains something worth knowing.

The authors use two previously untapped sets of data regarding manufacturing employment in the United States. They present evidence that the main statistical regularities of their data sets are also present in service industries and across countries. The data sets give them the number of *jobs* (defined as filled employment positions) at different *establishments* (roughly, factories) over time. Most importantly, they are able to track *gross* job flows over time, i.e., how many jobs are created and how many are destroyed at each establishment. Standard measures track only *net* job flows, i.e., the difference between the number created and the number destroyed. It turns out that net flows conceal an enormous amount. For

[1] The complete citation is: Davis, Steven J., John C. Haltiwanger, and Scott Schuh. *Job Creation and Destruction.* Cambridge, Mass.: MIT Press. 1996.

example, if we knew that the number of jobs in the U.S. grew 3% from 1998 to 1999, from say 100 million to 103 million, we would know the net change in jobs, but nothing about the gross changes in jobs. How many jobs were created? How many destroyed? Until DHS, there were simply no good answers to those questions.

The authors find that in a typical year 10% of jobs are created and that a roughly equal number are destroyed. The authors are also able to track job creation and destruction over the business cycle, and they find that job creation falls slightly during recessions, whereas job destruction grows strongly. Their data sets contain information about the nature of the establishments, so they are able to track job creation and destruction by employer and by factory characteristics. They convincingly explode one of the shibboleths of modern American political discourse: the myth of small-business job creation. It turns out that most jobs are created (and destroyed) by large, old plants and firms. This insight alone makes the book worth reading.

We begin with a primer on the notation used in DHS and then turn to a brief overview of the main conclusions of the book.

10.1 Job Creation and Destruction: Notation

Basic Notation

Variables in DHS can take up three subscripts. For example, the total number of filled employment positions at a plant is denoted X_{est}, where e denotes the *establishment* (that is, the plant), s denotes the *sector* (for example, the garment industry) and t denotes the *time period* (usually a specific year). If you find this notation confusing, ignore the differences among the first two subscripts e and s and just think of them as denoting the same thing: establishments. Capital letters will denote levels and lower-case letters will denote rates. The words "plant" and "establishment" mean the same thing. A *job* is defined as a filled employment position; no provision is made for considering unfilled positions.

Jobs are *created* when a plant increases the number of jobs from one period to the next, while jobs are *destroyed* when a plant decreases the number of jobs from one period to the next. *Gross job creation* is the sum of all new jobs at expanding and newly-born plants, while *gross job destruction* is the sum of all the destroyed jobs at shrinking and dying plants. Let X_{est} denote the number of jobs at establishment e in sector s at time t, and let S_t^+ be the set of establishments that are growing (i.e., hiring more workers) between periods $t-1$ and t. Then gross job creation is:

$$C_{st} = \sum_{e \in S_t^+} \Delta X_{est},$$

where Δ is the difference operator:

$$\Delta X_{est} \equiv X_{est} - X_{es,t-1}.$$

In words, C_{st} is the total number of new jobs at expanding and newly born plants in sector s between periods $t-1$ and t. Next we turn to job destruction. Let S_t^- be the set of establishments that are shrinking between periods $t-1$ and t. Then gross job destruction is:

$$D_{st} = \sum_{e \in S_t^-} |\Delta X_{est}|.$$

In words, D_{st} is the total number of all the jobs lost at shrinking and dying plants in sector s between periods $t-1$ and t. The absolute-value operator guarantees that D_{st} will be a positive number.

Next, we need a measure for the size of a plant. DHS use the average number of jobs between the current period and the last. For some establishment e in sector s at time t, DHS define its *size* Z_{est} as follows:

$$Z_{est} \equiv \frac{1}{2}\left(X_{est} + X_{es,t-1}\right).$$

Notice that the size in period t contains employment information for both periods t and $t-1$.

Suppose we discover that ten-thousand jobs were created in the mining sector in 1996. In our notation, we would write that as: $C_{m,1996} = 10{,}000$, where m is for "mining". This information would be more useful if compared with some measure of the number of jobs already present in the mining sector, which is what we call a *rate*. Then we could say, for example, that the gross rate of job creation in the mining sector was 10% in 1996.

The rate of employment growth at the plant level is defined as:

$$g_{est} = \frac{\Delta X_{est}}{Z_{est}}. \tag{10.1}$$

Let Z_{st} be the sum of all the plant sizes in sector s. Then the rate of job creation in a sector is defined as:

$$c_{st} = \frac{C_{st}}{Z_{st}}. \tag{10.2}$$

The rate of job destruction is defined similarly:

$$d_{st} = \frac{D_{st}}{Z_{st}}. \tag{10.3}$$

Notice that if some plant i dies between $t-1$ and t (so that $X_{i,t-1} > 0$ but $X_{it} = 0$), then the growth rate of the plant will be $g_{it} = -2$, while if plant i is born in t (so that $X_{i,t-1} = 0$ and $X_{it} > 0$), then the growth rate of the plant will be $g_{it} = 2$.

A Simple Example

Consider an economy with only one sector and three plants, P_1, P_2, and P_3. The following tables list the total employment figures for these three plants as well as: the gross levels of job creation and destruction at the plant level, average plant size, and gross rates of job creation and destruction.

Year	$X_{P_1,t}$	$X_{P_2,t}$	$X_{P_3,t}$	Total
1991	200	100	0	300
1992	0	300	200	500

Plant i	P_1	P_2	P_3	Total
$Z_{i,1992}$	100	200	100	400
$\Delta X_{i,1992}$	-200	200	200	200
$C_{i,1992}$	0	200	200	400
$D_{i,1992}$	200	0	0	200
$g_{i,1992}$	-2.0	1.0	2.0	

Notice that plant P_1 died between 1991 and 1992, so its growth rate was -2, the same as it would have been for any plant that died. Plant P_3 was born between 1991 and 1992, so its growth was 2, the same as it would have been for any plant that was born. This is because of the somewhat non-standard definition of plant size chosen by DHS. (See equation (10.1).) The economy went from 300 jobs in 1991 to 500 jobs in 1992, so it added 200 net jobs. However, two plants expanded, adding 200 jobs each, while one plant contracted, destroying 200 jobs. Thus gross job creation was 400 jobs, and gross job destruction was 200 jobs.

In 1992 the sizes $Z_{i,1992}$ of the three plants were 100, 200, and 100 jobs, respectively, so the aggregate plant size Z_{1992} was 400. Recall, S^+_{1992} is the set of plants that were growing between 1991 and 1992, so $S^+_{1992} = \{2, 3\}$. The rate of job creation c_{1992} for this economy was:

$$c_{1992} = \frac{C_{1992}}{Z_{1992}} = \frac{\sum_{i \in S^+_{1992}} \Delta X_{i,1992}}{Z_{1992}} = \frac{200 + 200}{400} = 1.$$

Now, recall that S^-_{1992} was the set of plants that were shrinking between 1991 and 1992. Using the same formulation, we can calculate the economy-wide rate of job destruction as:

$$d_{1992} = \frac{D_{1992}}{Z_{1992}} = \frac{\sum_{i \in S^-_{1992}} |\Delta X_{i,1992}|}{Z_{1992}} = \frac{|-200|}{400} = 0.5.$$

Job Reallocation, Net Job Creation, and Persistence

Let R_{st} be the sum of the number of jobs created and the number of jobs destroyed in sector s between periods $t-1$ and t. We call R_{st} the *level of job reallocation* in sector s at time t. Formally:

$$R_{st} = C_{st} + D_{st}.$$

Note that R_{st} is an upper bound for the number of workers who have to switch jobs to accommodate the redistribution of employment positions across plants.

We define the *employment status* of a citizen to be: "employed", "unemployed", or "not in the workforce". With that in mind, consider the previous example. For that one-sector economy, 600 jobs were reallocated. Imagine that all of the 400 newly created positions were filled with workers just entering the workforce and that none of the workers at the 200 destroyed jobs found employment. Then 600 workers changed employment status. Of course, if some of the workers at the 200 destroyed jobs had been hired to fill the newly created jobs, then the number of workers changing employment status would have been lower.

As before, we convert the level of job reallocation into a rate by dividing by our measure of plant size Z_{st}. Formally, the rate of job reallocation in sector s at time t is defined as:

$$r_{st} \equiv \frac{R_{st}}{Z_{st}} = \frac{C_{st} + D_{st}}{Z_{st}} = c_{st} + d_{st}.$$

Let NET_{st} be the difference between the gross levels of job creation and destruction in sector s at time t:

$$\mathrm{NET}_{st} = C_{st} - D_{st}.$$

This is the net level of job creation. Note that when job destruction is greater than job creation, NET_{st} will be negative. In the simple example above, $\mathrm{NET}_{1992} = 200$. Let net_{st} be the net rate of job creation in sector s at time t. Formally:

$$\mathrm{net}_{st} = c_{st} - d_{st}.$$

Now we are interested in creating a measure of the persistence of the changes in employment levels at establishments. We will first define a simple counting rule for determining how many of the new jobs created at a plant are still present after j periods, where j is an integer greater than or equal to one. Consider some plant i in year $t-1$ with $X_{i,t-1} = 100$ and $X_{it} = 110$. Thus $C_{it} = 10$, i.e., ten jobs were created (in gross) at plant i in year t.

Now consider the future year $t+j$. If employment $X_{i,t+j}$ at plant i in the year $t+j$ is 105, we say that five of the new jobs created at plant i in the year t have survived for j periods. If $X_{i,t+j} \leq 99$, we say that zero of the new jobs have survived. If $X_{i,t+j} \geq 110$, we say that all ten of the new jobs have survived.

Let $\delta_{it}(j)$ be the number of new jobs created at plant i in year t that have survived to year $t + j$, using the counting rule defined above. The *level of job persistence* $P_{it}^{c}(j)$ at plant i between periods t and $t + j$ is defined as the number of jobs created in year t that exist in all of the periods between t and $t + j$. Formally:

$$P_{it}^{c}(j) = \min\left\{\delta_{it}(1), \delta_{it}(2), \ldots, \delta_{it}(j)\right\}.$$

The rate of job persistence can be calculated by summing over all new and expanding establishments at time t and dividing by gross job creation at t. Using the fact that we have defined S_t^{+} to be the set of growing plants, we formally define the rate of job persistence as follows:

$$p_t^{c}(j) = \sum_{i \in S_t^{+}} \frac{P_{it}^{c}(j)}{C_{st}}.$$

Now we work through an example in order to fix ideas. The following chart gives employment levels for a firm between 1990 and 1995. X_t denotes the number of jobs at the plant in the year t, where all other subscripts have been dropped for convenience, and $\delta_{1991}(j)$ gives the number of jobs created in 1991 that still exist in the period $t + j$.

Year	Employment	Persistent jobs from 1991
1990	$X_{1990} = 100$	—
1991	$X_{1991} = 110$	—
1992	$X_{1992} = 109$	$\delta_{1991}(1) = 9$
1993	$X_{1993} = 108$	$\delta_{1991}(2) = 8$
1994	$X_{1994} = 107$	$\delta_{1991}(3) = 7$
1995	$X_{1995} = 108$	$\delta_{1991}(4) = 8$

We see that for this plant there were seven jobs that were created in 1991 and were also present in all periods from 1991 to 1995. Accordingly, $P_{1991}^{c}(4) = 7$. The subtle point is that one of the jobs of the 108 in 1995 was not one of those created in 1991.

Worker Reallocation and Excess Job Reallocation

We define the *level of worker reallocation* WR_t at time t to be the number of workers who change employment status or place of employment between periods $t - 1$ and t. There is no way to extract WR_t precisely from the data, since the data concentrate on jobs, not workers. However, we can provide upper and lower bounds on WR_t from the data on jobs.

Our job reallocation measure R_t may overstate the number of workers who change status or position. Consider a worker whose job is destroyed and then finds employment later within the sample period at a newly created job. This worker is counted twice in calculating

R_t. Only in the special case that all fired workers fail to find reemployment in the survey period would $\mathrm{WR}_t = R_t$. Hence R_t provides an upper bound on WR_t.

The minimum bound on WR_t is $\max\{C_t, D_t\}$. That is, if 75 jobs were created economy-wide and 50 were destroyed, we know that at least 75 workers had to change employment status or place of employment. Imagine that all 50 workers whose jobs were destroyed found employment at the newly created jobs. Then 25 other workers would have had to be drawn in to fill the remaining 25 new jobs. Hence the 50 fired workers changed place of employment and the 25 new workers changed employment status, moving from unemployed to employed.

10.2 Job Creation and Destruction: Facts

In this section we sketch briefly only the high points of the results in DHS. To answer the exercises at the back of this chapter, you will need to consult the text directly.

Over the sample period 1973-1988, the net manufacturing job creation rate ($c_t - d_t$) averaged -1.1%. This basic fact obscures the variation of job creation and destruction over the business cycle, by industry and by plant characteristic. In this section, we hit some of the high points.

The average annual rate of job destruction d_t in manufacturing was 10.3%, and the average rate of job creation c_t was slightly lower at 9.1%, so the average rate of job reallocation was 19.4%. The rate of job creation hit a peak of 13.3% in the recovery year of 1984, while the rate of job destruction peaked at 14.5% and 15.6% in the recession years of 1982-83. This points to the striking cyclical nature of job creation and destruction: in recessions job destruction spikes well above its mean, while job creation does not fall that much. Moreover, most of the job creation and destruction is concentrated in plants that open or close, rather than in plants that change size.

When DHS look at gross job flows across industries, they find that rates of job reallocation are uniformly high, i.e., all industries create and destroy lots of jobs. However, high-wage industries tend to have smaller gross job flows than low-wage industries. Finally, they find that the degree to which an industry faces competition from imports does not significantly affect job destruction. (For your information, imports make up less than 13% of the market for 80% of U.S. industries.)

Examining gross job flows by employer characteristics reveals that most jobs are not created by small business but rather by large, old firms. The pervasive myth of small-business job creation is fed by bureaucratic self-interest and by some elementary statistical errors. Understanding these errors is an instructive exercise in its own right and one of the most interesting parts of the DHS book.

Variable	Definition
X_{est}	Number of filled employment positions (jobs) at plant e, sector s, period t
S_t^+	Set of plants that grew between periods $t-1$ and t
S_t^+	Set of plants that shrunk between periods $t-1$ and t
C_{st}	Gross job creation
D_{st}	Gross job destruction
Z_{est}	Size of employment of plant e, sector s, period t
g_{est}	Plant-level rate of employment growth
c_{st}	Rate of job creation
d_{st}	Rate of job destruction
R_{st}	Level of job reallocation
r_{st}	Rate of job reallocation
NET_{st}	Net jobs created in sector s at time t
net_{st}	Net rate of creation of jobs in sector s at time t
$P_t^c(j)$	The number of jobs created in period t still present j periods later
$p_t^c(j)$	Rate of the persistence of job creation
WR_t	Level of worker reallocation

Table 10.1: Notation for Chapter 10

Finally, we touch on one last insight. Well-diversified plants tend to be more likely to survive recessions than single-output plants. This makes sense, since by producing a portfolio of products, a plant can spread the risk that a recession will completely stop demand for all of its output.

Exercises

Exercise 10.1 (Moderate)
Answer: True, False, or Uncertain, and explain.

1. "Did you know that America's 22 million small businesses are the principal source of new jobs?" (Source: Web page of the Small Business Administration.)

2. "In the next century, 20% of the population will suffice to keep the world economy

going.... A fifth of all job-seekers will be enough to produce all the commodities and to furnish the high-value services that world society will be able to afford" the remaining 80% will be kept pacified by a diet of "Tittytainment". (Source: Martin, Hans-Peter and Harald Schumann. *The Global Trap.* New York: St Martin's Press. 1996.)

Exercise 10.2 (Easy)
The plant-level rate of employment growth is defined as:

$$g_{est} = \frac{\Delta X_{est}}{Z_{est}},$$

where:

$$\Delta X_{est} = X_{est} - X_{es,t-1}.$$

That is, ΔX_{est} is the change in employment at plant e in sector s from $t-1$ to t. Show that $g_{est} = 2$ for all plants that are born between $t-1$ and t, and show that $g_{est} = -2$ for all plants that die between $t-1$ and t.

Exercise 10.3 (Easy)
Show the following:

$$c_{st} = \sum_{e \in S^+} \left(\frac{Z_{est}}{Z_{st}} \right) g_{est}, \text{ and:}$$

$$\text{net}_{st} = \sum_{e \in S} \left(\frac{Z_{est}}{Z_{st}} \right) g_{est}.$$

Here c_{st} is the average rate of job creation of all plants in sector s. What does the term Z_{est}/Z_{st} mean?

Exercise 10.4 (Moderate)
For the purposes of this exercise, assume that you have data on annual national job creation C_t and job destruction D_t for N years, so $t = 1 \ldots N$. Show that if annual national job reallocation R_t and net job creation NET_t have a negative covariance, then the variance of job destruction must be greater than the variance of job creation. Recall the definition of variance of a random variable X for which you have N observations, $\{x_i\}_{i=1}^{N}$:

$$\text{var}(X) = \frac{1}{N} \sum_{i=1}^{N} (x_i - \bar{x})^2,$$

where $\bar{x}$ is the mean of X. Similarly, recall the definition of the covariance of two variables X and Y. If there are N observations each, $\{x_i, y_i\}_{i=1}^{N}$, then:

$$\text{cov}(X, Y) = \frac{1}{N} \sum_{i=1}^{N} (x_i - \bar{x})(y_i - \bar{y}).$$

These definitions and the definitions of NET_t and R_t provide all the information necessary to answer this exercise.

Exercise 10.5 (Easy)
Consider the employment statistics in chart below. Compute each of the following five measures: (i) the economy-wide rate of job creation c_t; (ii) the economy-wide rate of job destruction d_t; (iii) the net rate of job creation net$_t$; (iv) the upper bound on the number of workers who had to change employment status as a result of the gross job changes; and (v) the lower bound on the number of workers who had to change employment status as a result of the gross job changes for each each of the years 1991, 1992, 1993, 1994, and 1995.

Year	$X_{1,t}$	$X_{2,t}$	$X_{3,t}$	c_t	d_t	net$_{st}$	UB	LB
1990	1000	0	500					
1991	800	100	800					
1992	1200	200	700					
1993	1000	400	600					
1994	800	800	500					
1995	400	1200	600					
1996	200	1400	600					
1997	0	2000	500					

Exercise 10.6 (Moderate)
For each of the following statements, determine if it is true, false or uncertain and why. If possible, back your assertions with specific statistical evidence from DHS.

1. Foreign competition is destroying American manufacturing jobs.
2. Robots and other capital improvements are replacing workers in factories.
3. Most job creation occurs at plants that grow about 10% and most job destruction occurs at plants that shrink about 10%.
4. Diversified plants are better able to withstand cyclical downturns.
5. Every year, high-wage manufacturing jobs are replaced by low-wage manufacturing jobs.

Chapter 11

Economic Growth

This chapter examines the determinants of economic growth. A startling fact about economic growth is the large variation in the growth experience of different countries in recent history. Some parts of the world, like the United States or Western Europe, experienced sustained economic growth over a period of more than 100 years, so by historical standards these countries are now enormously wealthy. This is not only true in absolute terms (i.e., GDP), but also if we measure wealth as income per capita (i.e., GDP per person). In contrast, there are countries where even today large parts of the population live close to the subsistence level, much the same as Europeans and Americans did some hundreds of years ago. Also, a group of countries that used to be relatively poor around the time of World War II managed to achieve even higher growth rates than the western industrialized countries, so their per capita incomes now approach those of western countries. Most of the members of this group are located in East Asia, like Japan, Singapore, Hong Kong, and so on.

It proves to be difficult to explain these different growth experiences within a single model. There are models that provide an explanation for the growth experience of the now industrialized countries, but most of these models fail to explain why much of the world is still poor. Models that seek to explain the difference between rich and poor countries are less successful at reproducing the growth facts for industrialized countries. We will therefore approach the topic of economic growth from a number of different angles. In Section 11.1 we present a number of facts about economic growth, facts that we will seek to explain with our growth models. Section 11.2 introduces the *Solow growth model*, a classic in the theory of economic growth. This model is quite successful at matching a number of facts about growth in industrialized countries. Section 11.3 introduces growth accounting, an empirical application of the Solow framework. This kind of accounting can be used to determine the sources of growth for a given country. In Section 11.4 we turn to the question why some countries are still poor today. A complete answer to this question is beyond the scope of this book; in fact, it is fair to say that a satisfactory answer has not been found yet.

Therefore we concentrate on only one important aspect of the growth experience of poor countries: the relationship between fertility, human capital, and growth.

11.1 Growth Facts

If we look at the group of industrialized countries only, we can identify a number of empirical regularities in the growth process. The British economist Nicholas Kaldor summarized these regularities in a number of stylized facts. Although he did that more than 50 years ago, the Kaldor facts still provide an accurate picture of growth in industrialized countries. Kaldor's first observation was that both output per worker and capital per worker grow over time. They also grow at similar rates, so the ratio of the aggregate capital stock to output or GDP does not change much over time. The return to capital, i.e., the interest that firms have to pay if they rent capital, is almost constant over time. Finally, the *labor share* and *capital share* are almost constant. The labor share is the fraction of output that goes to workers in the form of wages; it is computed as aggregate labor income divided by GDP. Similarly, the capital share is given by aggregate payments to capital divided by GDP. Notice that the Kaldor facts hold even if we consider long periods of time. For example, the capital-output ratio and the return to capital are not much different now from what they were 100 years ago, even though output is much higher now and the goods produced and the general technology have changed completely.

In addition to the Kaldor facts, another important fact about growth in the industrialized world is the convergence of per capita GDP of different countries and regions over time. For example, the relative difference in per capita GDP between the southern and northern states in the United States has diminished greatly since the Civil War. Similarly, countries like Germany and Japan that suffered greatly from World War II have grown fast since the war, so today per capita income in the United States, Japan, and Germany are similar again.

There are no empirical regularities comparable to the Kaldor facts that apply to both industrialized and developing countries. However, we can identify some factors that distinguish countries that went through industrialization and have a high income today from countries that remained relatively poor. An explanation of the role of such factors might be an important step toward understanding the large international differences in wealth. We are going to focus on the relationship between growth and fertility. Every now industrialized country has experienced a large drop in fertility rates, a process known as the *demographic transition*. All industrialized countries have low rates of population growth. Without immigration countries like Germany and Japan would actually shrink. Two centuries ago, fertility rates were much higher, as they are in most developing countries today. Today, almost all of the growth in world population takes place in developing countries. We will come back to these observations in the section on fertility and human capital, but first we present a model that accounts for the stylized facts about growth in developed countries.

11.2 The Solow Growth Model

A natural starting point for a theory of growth is the *aggregate production function*, which relates the total output of a country to the country's aggregate inputs of the factors of production. Consider the neoclassical production function:

$$Y_t = (A_t L_t)^\alpha K_{t-1}^{1-\alpha}. \tag{11.1}$$

We used a production function of this form already in the chapter on business cycles. Output depends on the aggregate labor input L_t, the aggregate capital input K_{t-1}, and a productivity parameter A_t. Of course, it is a simplification to consider only three determinants of output. We could include other factors like land or environmental quality, and our factors could be further subdivided, for example by distinguishing labor of different quality. It turns out, however, that a production function of the simple form in equation (11.1) is all we need to match the stylized facts of economic growth. The production function equation (11.1) exhibits constant returns to scale, which means that if we double both inputs, output also doubles. Our choice of a constant-returns-to-scale production function is not by accident: most results in this section hinge on this assumption.

Equation (11.1) indicates the potential sources of growth in output Y_t. Either the inputs L_t and K_{t-1} must grow, or productivity A_t must grow. If we want to explain economic growth, we need a theory that explains how the population (i.e., labor), the capital stock, and productivity change over time. The best approach would be to write down a model where the decisions of firms and households determine the changes in all these variables. The consumers would make decisions about savings and the number of children they want to have, which would explain growth in capital and population. Firms would engage in research and development, which would yield a theory of productivity growth. However, doing all those things at the same time results in a rather complicated model.

The model that we are going to present takes a simpler approach. Growth in productivity and population is assumed to be exogenous and constant. This allows us to concentrate on the accumulation of capital over time. Moreover, instead of modeling the savings decision explicitly, we assume that consumers invest a fixed fraction of output every period. Although these are quite radical simplifications, it turns out that the model is rather successful in explaining the stylized facts of economic growth in industrialized countries. It would be possible to write down a model with optimizing consumers that reaches the same conclusions. In fact, we wrote down that model already: the real business cycle model that we discussed in Chapter 9 used a neoclassical production function, and the optimal decision of the consumers was to invest a fixed fraction of their output in new capital. To keep the presentation simple, we will not go through individual optimization problems; instead, we will assume that it is optimal to save a fixed fraction of output. There are a number of names for the model. It is either referred to as the *Solow model* after its inventor Robert Solow, or as the *neoclassical growth model* after the neoclassical production function it uses, or as the *exogenous growth model* after the fact that there is no direct explanation for productivity growth.

The *law of motion* for a variable describes how the variable evolves over time. In the Solow model, the law of motion for capital is:

$$K_t = (1-\delta)K_{t-1} + I_t, \tag{11.2}$$

where I_t is investment and δ is the depreciation rate, which is between zero and one. We assume that investment is a fixed fraction $0 < s < 1$ of output:

$$I_t = sY_t = s(A_t L_t)^\alpha K_{t-1}^{1-\alpha}.$$

Productivity and labor grow at fixed rates μ and γ:

$$\begin{aligned} A_{t+1} &= (1+\mu)A_t, \text{ and:} \\ L_{t+1} &= (1+\gamma)L_t. \end{aligned}$$

We now have to find out how the economy develops, starting from any initial level of capital K_0, and then check whether the model is in line with the stylized facts of economic growth in industrialized countries.

We assume that there is a competitive firm operating the production technology. We can check one of the stylized facts, constant labor and capital share, just by solving the firm's problem. The profit maximization problem of the firm is:

$$\max_{L_t, K_{t-1}} \left\{ (A_t L_t)^\alpha K_{t-1}^{1-\alpha} - w_t L_t - r_t K_{t-1} \right\}.$$

The first-order conditions with respect to labor and capital yield formulas for wage and interest:

$$w_t = \alpha A_t^\alpha L_t^{\alpha-1} K_{t-1}^{1-\alpha}, \text{ and:} \tag{11.3}$$

$$r_t = (1-\alpha)(A_t L_t)^\alpha K_{t-1}^{-\alpha}. \tag{11.4}$$

We can use these to compute the labor and capital shares in the economy:

$$\begin{aligned} \frac{w_t L_t}{Y_t} &= \frac{\alpha A_t^\alpha L_t^{\alpha-1} K_{t-1}^{1-\alpha} L_t}{(A_t L_t)^\alpha K_{t-1}^{1-\alpha}} = \alpha, \text{ and:} \\ \frac{r_t K_{t-1}}{Y_t} &= \frac{(1-\alpha)(A_t L_t)^\alpha K_{t-1}^{-\alpha} K_{t-1}}{(A_t L_t)^\alpha K_{t-1}^{1-\alpha}} = 1-\alpha, \end{aligned}$$

so the labor share is α, and the capital share is $1-\alpha$. Thus both the labor and capital shares are indeed constant. This result is closely connected to the fact that the production function exhibits constant returns to scale. Actually, the fact that the labor and capital shares are about constant is one of the main arguments in favor of using production functions that exhibit constant returns to scale.

To continue, we have to take a closer look at the dynamics of capital accumulation in the model. It turns out that this is easiest to do if all variables are expressed in terms of units of effective labor $A_t L_t$. The product $A_t L_t$ is referred to as effective labor because increases in A_t make labor more productive. For example, $A_t = 2$ and $L_t = 1$ amounts to the same quantity of effective labor as $A_t = 1$ and $L_t = 2$. When put in terms of units of effective labor, all variables will be constant in the long run, which will simplify our analysis.

We will use lowercase letters for variables that are in terms of effective labor. That is, $y_t = Y_t/(A_t L_t)$, $k_{t-1} = K_{t-1}/(A_t L_t)$, and $i_t = I_t/(A_t L_t)$. Substituting $Y_t = y_t A_t L_t$ and so on into the production function, equation (11.1), yields:

$$y_t A_t L_t = (A_t L_t)^{\alpha}(k_{t-1} A_t L_t)^{1-\alpha}, \text{ or:}$$
$$y_t = k_{t-1}^{1-\alpha}. \tag{11.5}$$

From the law of motion for capital, equation (11.2), we get the law of motion in terms of effective labor:

$$k_t(1+\mu)A_t(1+\gamma)L_t = (1-\delta)k_{t-1}A_t L_t + i_t A_t L_t, \text{ or:}$$
$$k_t(1+\mu)(1+\gamma) = (1-\delta)k_{t-1} + i_t. \tag{11.6}$$

Finally, investment is determined by:

$$i_t = s y_t = s k_{t-1}^{1-\alpha}. \tag{11.7}$$

Plugging equation (11.7) into the law of motion in equation (11.6) yields:

$$k_t(1+\mu)(1+\gamma) = (1-\delta)k_{t-1} + s k_{t-1}^{1-\alpha}, \text{ or:}$$
$$k_t = \frac{(1-\delta)k_{t-1} + s k_{t-1}^{1-\alpha}}{(1+\mu)(1+\gamma)}. \tag{11.8}$$

This last equation determines the development of the capital stock over time. Dividing by k_{t-1} yields an expression for the growth rate of capital per unit of effective labor:

$$\frac{k_t}{k_{t-1}} = \frac{1-\delta + s k_{t-1}^{-\alpha}}{(1+\mu)(1+\gamma)}. \tag{11.9}$$

The expression k_t/k_{t-1} is called the *gross* growth rate of capital per unit of effective labor. The gross growth rate equals one plus the *net* growth rate. The growth rates in Chapter 1 were net growth rates.

Since the exponent on k_{t-1} in equation (11.9) is negative, the growth rate is inversely related to the capital stock. When a country has a lower level of capital per unit of effective labor, its capital and hence its output grow faster. Thus the model explains the convergence of GDP of countries and regions over time.

Since the growth rate of capital decreases in k_{t-1}, there is some level of k_{t-1} where capital per unit of effective labor stops growing. We say that the economy reaches a *steady state*. Once the economy arrives at this steady state, it stays there forever. Figure 11.1 is a graphical representation of the growth process in this economy. For simplicity, we assume for the moment that labor and productivity are constant, $\mu = \gamma = 0$. In that case, equation (11.8) simplifies to:

$$k_t = (1 - \delta)k_{t-1} + sk_{t-1}^{1-\alpha}, \text{ or:}$$
$$k_t - k_{t-1} = sk_{t-1}^{1-\alpha} - \delta k_{t-1}.$$

The change in capital per unit of effective labor is equal to the difference between investment and depreciation. Figure 11.1 shows the production function per unit of effective labor $y_t = k_{t-1}^{1-\alpha}$, investment $i_t = sk_{t-1}^{1-\alpha}$, and depreciation δk_{t-1}. Because the return to capital is diminishing, investment is a concave function of capital. For low values of capital, the difference between investment and depreciation is large, so the capital stock grows quickly. For larger values of capital, growth is smaller, and at the intersection of depreciation and investment the capital stock does not grow at all. The level of capital per unit of effective labor at which investment equals depreciation is the steady-state level of capital. In the long run, the economy approaches the steady-state level of capital per unit of effective labor, regardless of what the initial capital stock was. This is even true if the initial capital stock exceeds the steady-state level: capital per unit of effective labor will shrink, until the steady state is reached.

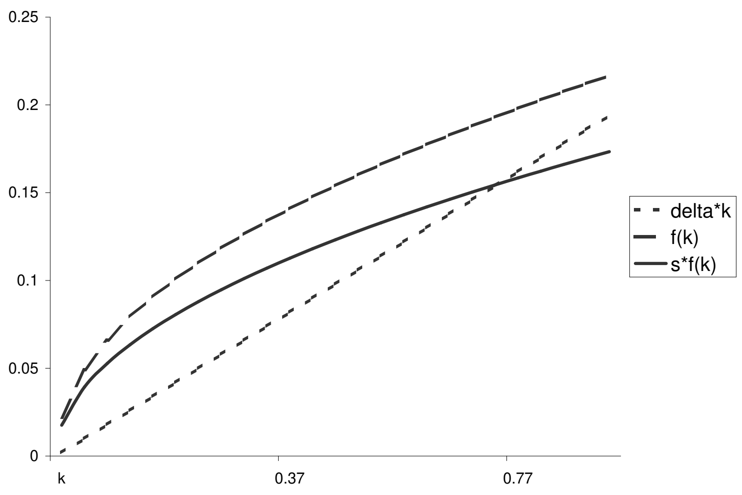

Figure 11.1: Output, Saving, and Depreciation in the Solow Model

At the steady state we have $k_t = k_{t-1}$. Using equation (11.8), we see that the steady-state level of capital per unit of effective labor $\bar{k}$ has to satisfy:

$$\bar{k}(1 + \mu)(1 + \gamma) = (1 - \delta)\bar{k} + s\bar{k}^{1-\alpha},$$

which yields:

$$\bar{k} = \left(\frac{s}{\delta + \mu + \gamma + \mu\gamma}\right)^{1/\alpha}. \tag{11.10}$$

We can use this equation to compute output, investment, and growth in the steady state. From equation (11.5), the steady-state level of output per effective labor unit is:

$$\bar{y} = \bar{k}^{1-\alpha} = \left(\frac{s}{\delta + \mu + \gamma + \mu\gamma}\right)^{\frac{1-\alpha}{\alpha}}.$$

The level of output depends positively on the saving rate. From equation (11.7), the steady-state investment per unit of effective labor is:

$$\bar{i} = s\left(\frac{s}{\delta + \mu + \gamma + \mu\gamma}\right)^{\frac{1-\alpha}{\alpha}}.$$

The steady-state growth rate of capital is $\mu + \gamma + \mu\gamma$:

$$\frac{K_t}{K_{t-1}} = \frac{\bar{k}(1+\mu)A_t(1+\gamma)L_t}{\bar{k}A_tL_t} = 1 + \mu + \gamma + \mu\gamma,$$

and the growth rate of output equals $\mu + \gamma + \mu\gamma$ as well. This implies that the long-run growth rate of an economy is independent of the saving rate. With a higher saving rate, the economy approaches a higher steady state, but the long-run growth rate is determined by growth in labor and productivity only.

There are still a number of stylized facts left to be checked. First, we will verify that the return to capital is constant. From equation (11.4), the return to capital is:

$$r_t = (1-\alpha)(A_tL_t)^{\alpha}K_{t-1}^{-\alpha} = (1-\alpha)\left(\frac{K_{t-1}}{A_tL_t}\right)^{-\alpha}.$$

In the steady state, capital per unit of effective labor is a constant $\bar{k}$. Therefore the return to capital in steady state is:

$$\bar{r}_t = (1-\alpha)\bar{k}^{-\alpha},$$

which is constant since $\bar{k}$ is constant. On the other hand, the wage is growing in the steady state, since the productivity of labor increases. The steady-state wage can be computed as:

$$w_t = \alpha A_t^{\alpha}L_t^{\alpha-1}K_{t-1}^{1-\alpha} = \alpha A_t\left(\frac{K_{t-1}}{A_tL_t}\right)^{1-\alpha} = \alpha A_t\bar{k}^{1-\alpha}, \text{ so:}$$

$$\frac{w_{t+1}}{w_t} = \frac{A_{t+1}}{A_t} = 1 + \mu,$$

which implies that the wage grows at the rate of technological progress.

The capital-output ratio in steady state is:

$$\frac{K_{t-1}}{Y_t} = \frac{\bar{k}A_tL_t}{\bar{y}A_tL_t} = \frac{\bar{k}}{\bar{y}},$$

which is a constant. This verifies the last stylized fact of economic growth on our list.

The Solow model succeeds in explaining all stylized facts of economic growth in industrialized countries. The key element of the model is the neoclassical constant-returns production function. Since returns to capital alone are decreasing, economies grow faster at lower levels of capital, until they approach the steady state, where units of effective labor and capital grow at the same rate. The model also explains why different saving rates in different industrialized countries do not translate into long-term differences in the growth rate. The saving rate affects the level of the steady state, but it does not affect the steady-state growth rate. The capital stock cannot grow faster than effective labor for a long time because of decreasing returns to capital.

Since the Solow model does well at matching the facts of economic growth, it forms the basis of many more-advanced models in macroeconomics. For example, our real business cycle model of Chapter 9 is a Solow model enriched by optimizing consumers and productivity shocks. On the other hand, the model works well only for countries that satisfy the assumptions of constant rates of population growth and technological progress. These assumptions are justified for industrialized countries, but they are not helpful for understanding the early stages of development of a country, which are usually accompanied by the demographic transition, so exogenous and constant population growth is not a useful assumption. We will look for possible explanations of fertility decisions below, but before that, we will introduce growth accounting, a method that allows us to decompose the growth rate of a country into growth in population, capital, and productivity.

11.3 Growth Accounting

In this section we will use the general framework of the Solow model to compute a decomposition of the rate of economic growth for a given country. Consider the neoclassical production function:

$$Y_t = (A_tL_t)^{\alpha}K_{t-1}^{1-\alpha}. \tag{11.11}$$

We will interpret Y_t as GDP, L_t as the number of workers, K_{t-1} as the aggregate capital stock, and A_t as a measure of overall productivity. We will be concerned with measuring the relative contributions of A_t, L_t, and K_{t-1} to growth in GDP. We assume that data for GDP, the labor force, and the aggregate capital stock are available. The first step is to compute the productivity parameter A_t. Solving the production function for A_t yields:

$$A_t = \frac{Y_t^{\frac{1}{\alpha}}}{L_tK_{t-1}^{\frac{1-\alpha}{\alpha}}}.$$

If α were known, we could compute the A_t right away. Luckily, we found out earlier that α is equal to the labor share. Therefore we can use the average labor share as an estimate of α and compute the A_t.

Now that A_t is available, the growth rates in A_t, L_t and K_{t-1} can be computed.[1] We can see how the growth rates in inputs and productivity affect the growth rate of GDP by taking the natural log of the production function:

$$\ln Y_t = \alpha \ln A_t + \alpha \ln L_t + (1 - \alpha) \ln K_{t-1}. \tag{11.12}$$

We are interested in growth between the years t and $t + k$, where k is some positive integer. Subtracting equation (11.12) at time t from the same equation at time $t + k$ yields:

$$\ln Y_{t+k} - \ln Y_t = \alpha(\ln A_{t+k} - \ln A_t) + \alpha(\ln L_{t+k} - \ln L_t) + (1 - \alpha)(\ln K_{t+k-1} - \ln K_{t-1}).$$

Thus the growth rate in output (the left-hand side) is α times the sum of growth in productivity and labor, plus $1 - \alpha$ times growth in capital. Using this, we can compute the relative contribution of the different factors. The fraction of output growth attributable to growth of the labor force is:

$$\frac{\alpha[\ln L_{t+k} - \ln L_t]}{\ln Y_{t+k} - \ln Y_t}.$$

The fraction due to growth in capital equals:

$$\frac{(1 - \alpha)[\ln K_{t+k-1} - \ln K_{t-1}]}{\ln Y_{t+k} - \ln Y_t}.$$

Finally, the remaining fraction is due to growth in productivity and can be computed as:

$$\frac{\alpha[\ln A_{t+k} - \ln A_t]}{\ln Y_{t+k} - \ln Y_t}.$$

It is hard to determine the exact cause of productivity growth. The way we compute it, it is merely a residual, the fraction of economic growth that cannot be explained by growth in labor and capital. Nevertheless, measuring productivity growth this way gives us a rough idea about the magnitude of technological progress in a country.

11.4 Fertility and Human Capital

In this section we will examine how people decide on the number of children they have. Growth and industrialization are closely connected to falling fertility rates. This was true for 19th century England, where industrialization once started, and it applies in the same way to the Asian countries that only recently began to grow at high rates and catch up with

[1]See Chapter 1 for a discussion of growth rates and how to compute them.

Western countries. Understanding these changes in fertility should help explain why some economies start to grow, while others remain poor.

The first economist to think in a systematic way about growth and fertility was Thomas Malthus. Back in 1798, he published his "Essay on Population", in which his basic thesis was that fertility was checked only by the food supply. As long as there was enough to eat, people would continue to produce children. Since this would lead to population growth rates in excess of the growth in the food supply, people would be pushed down to the subsistence level. According to Malthus's theory, sustained growth in per capita incomes was not possible; population growth would always catch up with increases in production and push per capita incomes down. Of course, today we know that Malthus was wrong, at least as far as the now industrialized countries are concerned. Still, his theory was an accurate description of population dynamics before the industrial revolution, and in many countries it seems to apply even today. Malthus lived in England just before the demographic transition took place. The very first stages of industrialization were accompanied by rapid population growth, and only with some lag did the fertility rates start to decline. We will take Malthus's theory as a point of departure in our quest for explanations for the demographic transition.

Stated in modern terms, Malthus thought that children were a normal good. When income went up, more children would be "consumed" by parents. We assume that parents have children for their enjoyment only, that is, we abstract from issues like child labor. As a simple example, consider a utility function over consumption c_t and number of children n_t of the form:

$$u(c_t, n_t) = \ln(c_t) + \ln(n_t).$$

We assume that the consumer supplies one unit of labor for real wage w_t and that the cost in terms of goods of raising a child is p. Therefore the budget constraint is:

$$c_t + pn_t = w_t.$$

By substituting for consumption, we can write the utility maximization problem as:

$$\max_{n_t} \left\{ \ln(w_t - pn_t) + \ln(n_t) \right\}.$$

The first-order condition with respect to n_t is:

$$-\frac{p}{w_t - pn_t} + \frac{1}{n_t} = 0, \text{ or:} \tag{FOC n_t}$$

$$n_t = \frac{w_t}{2p}. \tag{11.13}$$

Thus the higher the real wage, the more children are going to be produced.

If we assume that people live for one period, the number of children per adult n_t determines the growth rate of population L_t:

$$\frac{L_{t+1}}{L_t} = n_t.$$

To close the model, we have to specify how the wage is determined. Malthus's assumption was that the food supply could not be increased in proportion with population growth. In modern terms, he meant that there were decreasing returns to labor. As an example, assume that the aggregate production function is:

$$Y_t = A_t L_t^{\alpha},$$

with $0 < \alpha < 1$. Also assume that the real wage is equal to the marginal product of labor:

$$w_t = \alpha A_t L_t^{\alpha-1}. \tag{11.14}$$

We can combine equation (11.14) with the decision rule for the number of children in equation (11.13) to derive the law of motion for population:

$$\frac{L_{t+1}}{L_t} = \frac{\alpha A_t L_t^{\alpha-1}}{2p}, \text{ or:}$$

$$L_{t+1} = \frac{\alpha A_t L_t^{\alpha}}{2p}. \tag{11.15}$$

Notice that this last equation looks similar to the law of motion for capital in the Solow model. The growth rate of population decreases as population increases. At some point, the population stops growing and reaches a steady state $\bar{L}$. Using equation (11.15), the steady-state level of population can be computed as:

$$\bar{L} = \frac{\alpha A_t \bar{L}^{\alpha}}{2p}, \text{ or:}$$

$$\bar{L} = \left(\frac{\alpha A_t}{2p}\right)^{\frac{1}{1-\alpha}}.$$

In the steady state, we have $L_{t+1}/L_t = n_t = 1$. We can use this in equation (11.13) to compute the wage $\bar{w}$ in the steady state:

$$1 = \frac{\bar{w}}{2p}, \text{ or:}$$

$$\bar{w} = 2p.$$

Thus the wage in the steady state is independent of productivity A_t. An increase in A_t causes a rise in the population, but only until the wage is driven back down to its steady-state level. Even sustained growth in productivity will not raise per capita incomes. The population size will catch up with technological progress and put downward pressure on per capita incomes.

This Malthusian model successfully explains the relationship between population and output for almost all of history, and it still applies to large parts of the world today. Most developing countries have experienced large increases in overall output over the last 100 years. Unlike in Europe, however, this has resulted in large population increases rather than in increases in per capita incomes. Outside the European world, per capita incomes stayed virtually constant from 1700 to about 1950, just as the Malthusian model predicts.

Something must have changed in Europe in the nineteenth century that made it attractive to people to have less children, causing fertility rates to fall , so per capita incomes could start to grow. While these changes are by no means fully understood, we can identify a number of important factors. We will concentrate on two of them: the time-cost of raising children, and a quality-quantity tradeoff in decisions on children.

Human capital is a key element of the model that we are going to propose. So far, we considered all labor to be of equal quality. That might be a reasonable assumption for earlier times in history, but it certainly does not apply in our time, where special qualifications and skills are important. In the model, human capital consists of two components. First, there is innate human capital that is possessed by every worker, regardless of education. We will denote this component of human capital by H_0. This basic human capital reflects the fact that even a person with no special skill of any kind is able to carry out simple tasks that require manual labor only. In addition to this basic endowment, people can acquire extra human capital H_t through education by their parents. H_t reflects special skills that have to be taught to a worker. The total endowment with human capital of a worker is $H_0 + H_t$.

To come back to fertility decisions, we now assume that parents care both about the number n_t of their children and their "quality", or human capital $H_0 + H_{t+1}$. Preferences take the form:

$$u(c_t, n_t, H_{t+1}) = \ln(c_t) + \ln(n_t(H_0 + H_{t+1})).$$

The other new feature of this model is that parents must invest time, rather than goods, to raise children. In the Malthusian model, p units of the consumption good were needed to raise a child. We now assume that this cost in terms of goods is relatively small, so it can be omitted for simplicity. Instead, children require attention. For each child, a fraction h of the total time available has to be used to raise the child. In addition, the parents can decide to educate their children and spend fraction e_t of their time doing that. This implies that only a fraction $1 - hn_t - e_t$ is left for work. If w_t is the wage per unit of human capital when working all the time, the budget constraint is:

$$c_t = w_t(H_0 + H_t)(1 - hn_t - e_t). \tag{11.16}$$

The right-hand side says that income is the wage multiplied by human capital and the fraction of time worked. All this income is spent on consumption. We still have to specify the determination of the human capital of the children. We assume that the extra human capital of each child H_{t+1} depends on: the acquired human capital H_t of that child's parents, and the time e_t the parents spend teaching their children:

$$H_{t+1} = \gamma e_t H_t. \tag{11.17}$$

Here γ is a positive parameter. The interpretation of equation (11.17) is that parents who are skilled themselves are better at teaching their children. A person who does not have any skills is also unable to teach anything to his or her children.

We now want to determine how fertility is related to human capital in this model. If we plug the constraints in equations (11.16) and (11.17) into the utility function, the utility maximization problem becomes:

$$\max_{n_t, e_t} \left\{ \ln(w_t(H_0 + H_t)(1 - hn_t - e_t)) + \ln(n_t(H_0 + \gamma e_t H_t)) \right\}.$$

The first-order conditions with respect to n_t and e_t are:

$$\text{(FOC } n_t\text{)} \qquad -\frac{h}{1 - hn_t - e_t} + \frac{1}{n_t} = 0; \text{ and:}$$

$$\text{(FOC } e_t\text{)} \qquad -\frac{1}{1 - hn_t - e_t} + \frac{\gamma H_t}{H_0 + \gamma e_t H_t} = 0.$$

(FOC n_t) can be rewritten as:

$$hn_t = 1 - hn_t - e_t, \text{ or:}$$

$$e_t = 1 - 2hn_t. \tag{11.18}$$

Using equation (11.18) in (FOC e_t) allows us to compute the optimal fertility decision:

$$\begin{aligned} \gamma H_t(1 - hn_t - (1 - 2hn_t)) &= H_0 + (1 - 2hn_t)\gamma H_t; \text{ or:} \\ \gamma H_t hn_t &= H_0 + \gamma H_t - 2\gamma H_t hn_t; \text{ or:} \\ 3\gamma H_t hn_t &= H_0 + \gamma H_t; \text{ or:} \end{aligned}$$

$$n_t = \frac{1}{3h}\left[\frac{H_0}{\gamma H_t} + 1\right]. \tag{11.19}$$

According to equation (11.19), the key determinant of fertility is human capital H_t. If it is close to zero, the number of children is very high. If we added a cost of children in terms of goods to this model, for low values of H_t the outcomes would be identical to the Malthusian model. However, things change dramatically when H_t is high. Fertility falls, and if H_t continues to rise, the number of children reaches the steady state: $\bar{n} = 1/(3h)$. There are two reasons for this outcome. On the one hand, if human capital increases, the value of time also increases. It becomes more and more costly to spend a lot of time raising children, so parents decide to have less of them. The other reason is that people with high human capital are better at teaching children. That makes it more attractive for them to invest in the quality instead of the quantity of children.

The model sheds some light on the reasons why today fertility in industrialized countries is so much lower than that in developing countries. The theory also has applications within a given country. For example, in the United States teenagers are much more likely to become pregnant if they are school dropouts. The model suggests that this is not by accident. People with low education have a relatively low value of time, so spending time with children is less expensive for them.

The question that the model does not answer is how the transition from the one state to the other takes place. How did England manage to leave the Malthusian steady state?

In the model, only a sudden jump in H_t over some critical level could perform this task, which is not a very convincing explanation for the demographic transition. Still, the model is a significant improvement over theories that assume that population growth rates are exogenous and constant. More research on this and related questions will be needed before we can hope to find a complete explanation for the demographic transition and the wide disparity in wealth around the world.

Variable	Definition
Y_t	Aggregate output
y_t	Output per unit of effective labor
L_t	Aggregate labor input or population
K_t	Aggregate capital stock
k_t	Capital per unit of effective labor
A_t	Productivity parameter
I_t	Aggregate investment
i_t	Investment per unit of effective labor
w_t	Wage
r_t	Return on capital
δ	Depreciation rate
α	Parameter in the production function
$u(\cdot)$	Utility function
c_t	Consumption
n_t	Number of children
p	Cost of raising a child, in terms of goods
h	Cost of raising a child, in terms of time
e_t	Time spent on educating children
H_0	Innate human capital
H_t	Acquired human capital
γ	Parameter in the production function for human capital

Table 11.1: Notation for Chapter 11

Exercises

Exercise 11.1 (Easy)
Suppose the aggregate production technology is $Y = 3L^{.7}K^{.3}$ and that $L = 150$. Both the labor force and productivity are constant. Assume that the depreciation rate is 10% and that

20% of output is saved and invested each year. What is the steady-state level of output?

Exercise 11.2 (Moderate)
Assume that the Solow model accurately describes the growth experience of Kuwait. As a result of the Gulf war, much of the capital in Kuwait (oil extracting equipment, vehicles, structures etc.) was destroyed. Answer the following questions, and provide brief explanations.

- What will be the effect of this event on per capita income in Kuwait in the next five years?
- What will be the effect of this event on per capita income in Kuwait in the long run?
- What will be the effect of this event on the annual growth rate of per capita income in Kuwait in the next five years?
- What will be the effect of this event on the growth rate of per capita income in Kuwait in the long run?
- Will recovery in Kuwait occur faster if investment by foreigners is permitted, or if it is prohibited?
- Would Kuwaiti workers gain or lose by a prohibition of foreign investment? Would Kuwaiti capitalists gain or lose?

Exercise 11.3 (Moderate)
In this and the following two exercises, you will apply growth accounting to measure the determinants of growth in output per worker in a country of your choice. To start, you need to pick a country and retrieve data on real GDP per worker and capital per worker. You can get the time series you need from the Penn World Tables. See Exercise 9.1 for information about how to access this data set. You should use data for all years that are available.

In Section 11.3, we introduced growth accounting for output growth, while in this exercise we want to explain growth in output per worker. We therefore have to redo the analysis of Section 11.3 in terms of output per worker. The first step is to divide the production function in equation (11.11) by the number of workers L_t, which yields:

$$\frac{Y_t}{L_t} = \frac{(A_t L_t)^{\alpha} K_{t-1}^{1-\alpha}}{L_t} = \frac{A_t^{\alpha} L_t^{\alpha} K_{t-1}^{1-\alpha}}{L_t^{\alpha} L_t^{1-\alpha}} = A_t^{\alpha} \left(\frac{K_{t-1}}{L_t}\right)^{1-\alpha}. \tag{11.20}$$

Equation (11.20) relates output per worker Y_t/L_t to capital per worker K_{t-1}/L_t. If we use lower case letters to denote per-worker values ($y_t = Y_t/L_t$, $k_{t-1} = K_{t-1}/L_t$), we can write equation (11.20) as:

$$y_t = A_t^{\alpha} k_{t-1}^{1-\alpha}. \tag{11.21}$$

Use equation (11.21) to derive a formula for A_t and to derive a decomposition of growth in output per worker into growth in capital per worker and productivity growth. You can do that by following the same steps we took in Section 11.3.

Exercise 11.4 (Moderate)
Compute the productivity parameter A_t for each year in your sample. For your computations, assume that $1 - \alpha = .4$. This is approximately equal to the capital share in the United States, and we assume that all countries use the same production function. In fact, in most countries measures for $1 - \alpha$ are close to .4.

Exercise 11.5 (Moderate)
By using log-differences, compute the growth rate of GDP, productivity, and capital per worker for each year in your sample. Also compute the average growth rate for these three variables.

Exercise 11.6 (Moderate)
What percentage of average growth per worker is explained by growth in capital, and what percentage by productivity growth? For the period from 1965 to 1992, the average growth rate of output per worker was 2.7% in the United States, and productivity growth averaged 2.3%. How do these numbers compare to your country? Does the neoclassical growth model offer an explanation of the performance of your country relative to the United States? If not, how do you explain the differences?

Chapter 12

The Effect of Government Purchases

In this chapter we consider how governmental purchases of goods and services affect the economy. Governments tend to spend money on two things: wars and social services. Barro's Figure 12.2 shows that expenditures by the U.S. government have comprised a generally increasing fraction of GNP since 1928, but even today that fraction is nowhere near the peak it attained during WWII. This pattern is generally repeated across countries. The taste for social services seems to increase with national wealth, so the governments of richer countries tend to spend more, as a fraction of GDP, than the governments of poorer countries, especially during peacetime. Of course, there are exceptions to this pattern.

We will examine government spending in three ways:

1. We shall consider the effect of permanent changes in government spending in order to think about the secular peacetime increases in spending;
2. We shall consider temporary changes in government spending in order to think about the effect of sudden spikes like wars;
3. We shall begin an analysis of the effect of government social programs. Since government social programs (unemployment insurance, social security systems) are inextricably linked to tax systems, we will defer part of our analysis to the next chapter.

Since we have yet to fully discuss tax policy, for this chapter we will assume that the government levies a very special kind of tax: a *lump-sum tax.* That is, the government announces a spending plan and then simply removes that amount of money from the budget of the representative household. As we shall see in the next chapter, this kind of tax system does not distort the household's choices.

In the Barro textbook, the government budget constraint, in addition to lump sum taxes, also contains fiat currency. In this chapter we will assume that the government does not use the printing press to finance its purchases. In later chapters (especially Chapter 18) we will examine this effect in much greater detail.

12.1 Permanent Changes in Government Spending

Assume that the government announces a permanent level of government spending, G, to be levied each period. What is the role of these government expenditures? The government provides *productive* services, such as a court system for enforcing contracts and an interstate highway system for quickly and cheaply transporting goods. The government also provides *consumption* services such as public parks and entertainment spectacles such as trips to the moon and congressional hearings. We focus on the first role.

How should we model the productive services provided by the government? We shall analyze a model under two assumptions:

1. Government spending at some constant rate ϕ,
2. The effect of government spending G is augmented by the level of capital, K_t, so output Y increases by the amount $\phi G K_t$.

In the first case, \$100 of government spending increases output by 100ϕ regardless of the current level of capital, while in the second case, the same \$100 boosts output much more in nations with more capital.

The representative household lives forever and has preferences over consumption streams $\{C_t\}_{t=0}^{\infty}$ given by:

$$V(\{C_t\}_{t=0}^{\infty}) = \sum_{t=0}^{\infty} \beta^t U(C_t).$$

Where $U' > 0$ and $U'' < 0$. Here $0 < \beta < 1$ reflects impatience. In addition, to keep the algebra nice, we will say that:

$$\beta = \frac{1}{1+\rho}.$$

Here β is the discount factor and ρ the discount rate.

The household has access to a productive technology mapping capital K_t into *private* output Y_t^P of:

$$Y_t^P = K_t^{\alpha}.$$

Total output (and hence income) of the household will be the sum of private output and government-augmented output, Y_t^G. Government augmented output will take on one of two values:

$$Y_t^G = \phi G, \text{ or:} \tag{12.1}$$

$$Y_t^G = \phi G K_t. \tag{12.2}$$

Equation (12.1) corresponds to the case of government spending affecting total output the same amount no matter what the level of capital. Equation (12.2) corresponds to the case of government spending affecting total output more when the level of capital is high. We shall examine the effect of G on capital accumulation, aggregate output and consumption under both of these assumptions.

The household must split total income $Y_t = Y_t^P + Y_t^G$ into consumption C_t, investment I_t and *payments to the government* of G. Recall that we assumed the government would simply levy lump-sum taxes. Now we are using that assumption. The household's resource constraint is thus:

$$C_t + I_t + G \leq Y_t. \tag{12.3}$$

Finally, there is a law of motion for the capital stock K_t. Each period, a proportion δ of the capital stock vanishes due to physical depreciation, so only the remaining $(1-\delta)$ proportion survives into the next period. In addition, capital may be augmented by investment. Thus capital evolves according to:

$$K_{t+1} = (1 - \delta)K_t + I_t. \tag{12.4}$$

We assume that the representative household begins life with some initial stock of capital $K_0 > 0$.

We are interested in writing C_t as a function of next period's capital stock K_{t+1}. Combining equations (12.3) and (12.4) gives:

$$C_t = K_t^\alpha + (1 - \delta)K_t - K_{t+1} - G + \phi G, \text{ or:} \tag{BC1}$$

$$C_t = K_t^\alpha + (1 - \delta)K_t - K_{t+1} - G + \phi G K_t. \tag{BC2}$$

The differences between the two equations arises from which version of the government technology we use, equation (12.1) or (12.2).

Analysis with Equation (BC1)

Let us begin our analysis with the first version of the government spending technology, equation (12.1). Thus we are using as the relevant budget constraint equation (BC1). The household's problem becomes:

$$\max_{\{K_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t U\left[K_t^{\alpha} + (1-\delta)K_t - K_{t+1} - (1-\phi)G\right].$$

We take first-order conditions with respect to the choice of next period's capital K_{j+1} in some typical period j. Remember that K_{j+1} appears in two periods, j and $j+1$:

$$\beta^j U'(C_j)[-1] + \beta^{j+1} U'(C_{j+1})\left[\alpha K_{j+1}^{\alpha-1} + 1 - \delta\right] = 0.$$

For all $j = 0, 1, \ldots, \infty$. Here C_j is given by equation (BC1) above. Simplifying produces:

$$U'(C_j) = \beta U'(C_{j+1})[\alpha K_{j+1}^{\alpha-1} + 1 - \delta]. \tag{12.5}$$

For simplicity (and as in other chapters) we choose not to solve this for the transition path from the initial level of capital K_0 to the steady state level K_{ss}, and instead focus on characterizing the steady state. At a steady state, by definition the capital stock is constant:

$$K_t = K_{t+1} = K_{ss}.$$

As a result:

$$C_t = C_{t+1} = C_{ss}, \text{ and:}$$
$$I_t = I_{t+1} = I_{ss} = \delta K_{ss}.$$

Equation (12.5) at the steady-state becomes:

$$U'(C_{ss}) = \beta U'(C_{ss})[\alpha K_{ss}{}^{\alpha-1} + 1 - \delta].$$

Simplifying, and using the definition of β as $1/(1+\rho)$ produces:

$$1 + \rho = \alpha K_{ss}{}^{\alpha-1} + 1 - \delta.$$

We now solve for the steady-state capital level:

$$K_{ss} = \left(\frac{\alpha}{\rho+\delta}\right)^{\frac{1}{1-\alpha}}.$$

Notice immediately that, under this formulation of government spending the steady state capital level is independent of government spending. As we shall see in the next chapter, this is a direct consequence of the lump-sum tax technology. If the government had to use a distortionary tax, K_{ss} would be affected by G.

Given K_{ss}, it is easy to calculate the other variables that the household controls: steady-state private income, $Y_{ss}{}^P$, consumption C_{ss}, and investment, I_{ss}. From the technology, we know that $Y^P = K^{\alpha}$, so:

$$Y_{ss}{}^P = K_{ss}{}^{\alpha}.$$

Total output (GDP) is private output Y^P plus government output Y^G, or:

$$Y_{ss} = {K_{ss}}^{\alpha} + \phi G.$$

Consumption is, in this case, determined by the budget constraint equation (BC1). At the steady-state, then:

$$C_{ss} = {K_{ss}}^{\alpha} + (1 - \delta)K_{ss} - K_{ss} - (1 - \phi)G.$$

We can simplify this to produce:

$$C_{ss} = {K_{ss}}^{\alpha} - \delta K_{ss} - (1 - \phi)G.$$

At the steady-state, the household must be investing just enough in new capital to offset depreciation. Substituting into the law of motion for capital provides:

$$I_{ss} = \delta K_{ss}.$$

Now we are ready to determine the effect of government spending on total output, consumption and the capital level. When we think about changing G we are comparing two different steady states. Thus there may be short-term fluctuations immediately after the government announces its new spending plan, but we are concerned here with the long-run effects.

Notice immediately that:

$$\frac{dK_{ss}}{dG} = 0, \tag{12.6}$$

$$\frac{dY_{ss}}{dG} = \frac{d}{dG}(Y^P + Y^G) = \phi, \text{ and:} \tag{12.7}$$

$$\frac{dC_{ss}}{dG} = -(1 - \phi)G.. \tag{12.8}$$

That is, total output is increasing in G but consumption is decreasing in G if $\phi < 1$. Thus $\phi < 1$ is an example of *crowding out.* Think of it this way: the government spends \$1000 on a new factory, which produces 1000ϕ units of new output. The household pays the \$1000 in taxes required to construct the new factory, does not alter its capital level and enjoys the extra output of 1000ϕ as consumption. If $\phi < 1$ the household has lost consumption. Thus output has increased and consumption has decreased.

Why do we automatically assume that $\phi < 1$? This is equivalent to saying that the government is worse at building factories than the private sector. The government may be the only institution that can provide contract enforcement, police and national defense, but long history has shown that it cannot in general produce final goods as effectively as the private sector.

One final note before we turn our attention to the effect of production augmenting government spending. Government transfer payments, in which the government takes money from one agent and gives it to another, fit nicely into this category of expenditure. Transfer payments have absolutely no productive effects, and the government institutions required to administer the transfer payments systems will prevent the perfect transmission of money from one agent to another. Since we are working with a representative consumer, transfer payments appear as taxes which are partially refunded.

Analysis with Equation (BC2)

Now let us consider the effect of government spending whose benefits are proportional to capital stock. We will use precisely the same analysis as before. except that now consumption C_t as a function of capital K_t and K_{t+1} and government spending G will be given by equation (BC2) above.

The household's problem becomes:

$$\max_{\{K_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t U\left[K_t^{\alpha} + (1-\delta+\phi G)K_t - K_{t+1} - G\right].$$

We take first-order conditions with respect to the choice of next period's capital stock K_{j+1} in some typical period j. Remember the trick with these problems: K_{j+1} appears twice in the maximization problem, first negatively in period j and then positively in period $j+1$:

$$\beta^j U'(C_j)[-1] + \beta^{j+1} U'(C_{j+1})\left[\alpha K_{j+1}^{\alpha-1} + 1 - \delta + \phi G\right] = 0.$$

For all $j = 0, 1, \ldots, \infty$. C_j is given by equation (BC2). Simplifying produces:

$$U'(C_j) = \beta U'(C_{j+1})[\alpha K_{j+1}^{\alpha-1} + 1 - \delta + \phi G]. \tag{12.9}$$

Compare this with the previous simplified first-order condition, equation (12.5) above. Notice that in equation (12.5) the government spending term G does not appear. Here it does. This should alert us immediately that something new is about to happen. As before, we assume a steady state and characterize it. At the steady state:

$$U'(C_{\text{ss}}) = \beta U'(C_{\text{ss}})[\alpha K_{\text{ss}}{}^{\alpha-1} + 1 - \delta + \phi G].$$

Using our definition of β as $1/(1+\rho)$ this becomes:

$$1 + \rho = \alpha K_{\text{ss}}{}^{\alpha-1} + 1 - \delta + \phi G.$$

Hence the steady-state capital level is:

$$K_{\text{ss}} = \left(\frac{\alpha}{\rho + \delta - \phi G}\right)^{\frac{1}{1-\alpha}}.$$

Notice immediately that, under this formulation of government spending the steady-state capital level is increasing in government spending. If the government were forced to finance its spending with a distortionary tax this result might not go through.

Given the steady-state capital level, it is easy to calculate the steady-state levels of total output Y_{SS}, consumption C_{SS} and investment, I_{SS}. Since the steady-state capital level, K_{SS}, is now affected by G, both public output Y^G and private output Y^P are in turn affected by G. Given the production function, we see that:

$$Y_{SS} = {K_{SS}}^{\alpha} + \phi K_{SS} G.$$

From the budget constraint equation (BC2) above, we see that the steady state, consumption is:

$$C_{SS} = {K_{SS}}^{\alpha} - \delta K_{SS} - (1 - \phi K_{SS})G.$$

As before, the household must be investing just enough to overcome depreciation, to keep the capital level constant:

$$I_{SS} = \delta K_{SS}.$$

Now we can reconsider the effect of government spending on total output, consumption and the capital level. Some of these derivatives are going to be fairly involved, but if we break them down into their constituent pieces they become quite manageable.

Begin by defining:

$$X \equiv \frac{\alpha}{\rho + \delta - \phi G}.$$

Note that:

$$\frac{dX}{dG} = \frac{\phi}{\rho + \delta - \phi G} X.$$

The steady-state capital stock is:

$$K_{SS} = X^{\frac{1}{1-\alpha}},$$

so the derivative of the steady-state capital stock with respect to G is:

$$\frac{dK_{SS}}{dG} = \frac{1}{1-\alpha} X^{\frac{1}{1-\alpha}-1} \frac{dX}{dG}.$$

Plugging in dX/dG yields:

$$\begin{aligned}
\frac{dK_{SS}}{dG} &= \frac{1}{1-\alpha} X^{\frac{1}{1-\alpha}-1} \frac{\phi}{\rho + \delta - \phi G} X \\
&= \frac{1}{1-\alpha} \frac{\phi}{\rho + \delta - \phi G} X^{\frac{1}{1-\alpha}} \\
&= \frac{1}{1-\alpha} \frac{\phi}{\rho + \delta - \phi G} K_{SS}. \qquad (12.10)
\end{aligned}$$

Armed with this result we can tackle the other items of interest. First, consider the effect of increased spending on aggregate output:

$$\begin{aligned}
\frac{dY_{ss}}{dG} &= \frac{d}{dG}({Y_{ss}}^P + {Y_{ss}}^G) \\
&= \frac{d}{dG}\left({K_{ss}}^\alpha + \phi G K_{ss}\right) \\
&= \alpha {K_{ss}}^{\alpha-1}\frac{dK_{ss}}{dG} + \phi G \frac{dK_{ss}}{dG} \\
&= \alpha {K_{ss}}^{\alpha-1}\frac{1}{1-\alpha}\frac{\phi}{\rho+\delta-\phi G}K_{ss} + \phi G \frac{1}{1-\alpha}\frac{\phi}{\rho+\delta-\phi G}K_{ss} \\
&= \frac{1}{1-\alpha}\frac{\phi}{\rho+\delta-\phi G}\left[\alpha {K_{ss}}^\alpha + \phi G K_{ss}\right] \\
&= \frac{1}{1-\alpha}\frac{\phi}{\rho+\delta-\phi G}\left[\alpha {Y_{ss}}^P + {Y_{ss}}^G\right]. \qquad (12.11)
\end{aligned}$$

Compare the effect of government spending on aggregate output here with the effect of government spending on aggregate output when government spending simply augments output directly, equation (12.7) above. Notice that while previously every dollar of government spending translated into ϕ dollars of extra output no matter what the output level, now government spending is more productive in richer economies.

Finally, we turn our attention to consumption. Recall that before, for $\phi < 1$, consumption decreased as government spending increased, that is, consumption was crowded out. Now we shall see that, while consumption may be crowded out, it will not necessarily be crowded out. In fact, in rich economies, increases in government spending may increase consumption. Once again, this result will hinge to a certain extent on the assumption of a perfect tax technology. Begin by writing consumption as:

$$C_{ss} = {K_{ss}}^\alpha - \delta K_{ss} - G + \phi G K_{ss}, \text{ so:} \qquad (12.12)$$

$$\begin{aligned}
\frac{dC_{ss}}{dG} &= \frac{d}{dG}({K_{ss}}^\alpha + (\phi G - \delta)K_{ss} - G) \\
&= \alpha {K_{ss}}^{\alpha-1}\frac{dK_{ss}}{dG} + (\phi G - \delta)\frac{dK_{ss}}{dG} + \phi K_{ss} - 1 \\
&= \alpha {K_{ss}}^{\alpha-1}\frac{1}{1-\alpha}\frac{\phi}{\rho+\delta-\phi G}K_{ss} + (\phi G - \delta)\frac{1}{1-\alpha}\frac{\phi}{\rho+\delta-\phi G}K_{ss} + \phi K_{ss} - 1 \\
&= \frac{1}{1-\alpha}\frac{\phi}{\rho+\delta-\phi G}\left[\alpha {K_{ss}}^\alpha + (\phi - \delta)K_{ss}\right] + \phi K_{ss} - 1.
\end{aligned}$$

The first two terms are certainly positive. The question is, are they large enough to outweigh the -1? Even if $\phi < 1$, for large values of G this may indeed be the case.

Increasing Returns to Scale and Government Spending

Thus we have seen that the effect of government spending depends crucially on assumptions about how it is transformed into output. In the next chapter we will also see that it depends on how the government raises the revenue it spends.

Our second assumption about technology, embodied in equation (BC2), generated some exciting results about government spending. It seems that, if the world is indeed like the model, there is a potential for governments to provide us with a free lunch. Take a closer look at equation (12.2). If we assumed that the representative household controlled G directly (through representative government, for example) what level would it choose? Ignore the dynamics for a moment and consider the household's consumption C^a given that it has chosen some level of G and K:

$$C^a \equiv C(K,G) = K^\alpha + \phi GK - \delta K - G.$$

Now suppose the household doubles its inputs of K and G, so it is consuming some amount C^b:

$$C^b \equiv C(2K,2G) = 2^\alpha K^\alpha + 4\phi GK - 2\delta K - 2G.$$

For sufficiently large values of G and K it is easy to see that:

$$C^b > 2C^a.$$

In other words, by doubling G and K, the representative household could more than double net consumption. This is the standard free lunch of increasing returns to scale, in this case jointly in K and G. In the real world, are there increasing returns to scale jointly in government spending and capital? In certain areas this is almost certainly true. For example, by providing sewage and water-treatment services the government prevents epidemics and lowers the cost of clean water to consumers. This is a powerful direct benefit. This direct benefit is increasing in the population concentration (a small village probably would do fine with an outhouse, while 19th-century Chicago was periodically decimated by Cholera epidemics before the construction of the sanitary canal), and in turn encourages greater capital accumulation. No one business or household in 18th century Chicago would have found it worthwhile to build a sewage system, so it would have been difficult for private enterprise alone to have provided the improvements. Furthermore, since the Chicago sewage system depends in large measure on the Sanitary Canal, which had to be dug across previously-private land, it may have been impossible to build without the power of eminent domain.[1]

Unfortunately, there are few such clear-cut cases of increasing returns to scale combined with the requirement of government power. Why should a city government construct a stadium to lure sports teams? To build it, the government has to tax citizens who may experience no direct or indirect benefit.

[1] For more information on Chicago's sewer works, see Robin L. Einhorn, *Property Rules: Political Economy in Chicago, 1833-1872.*

Transitions in the Example Economies

We have so far ignored the problem of *transitions* in order to concentrate on steady-state behavior. But transition dynamics, describing the path that capital, consumption and the interest rate take as an economy transitions from low capital to the steady state capital level can be extremely interesting. In this subsection we will study transition dynamics by numerically simulating them on a computer.

Consider an example economy in which $G = 0.4$, $\phi = 0.1$, $\alpha = 0.25$, $\rho = 0.075$, $\delta = 0.1$ and $\beta = 1/(1 + \rho)$. Using the technology from equation (BC1), the steady-state capital level is $K_{ss} = 1.6089$, using the better technology from equation (BC2), the steady-state capital level is $K_{ss} = 2.2741$. Notice that, since $G = 0.4$, government spending as a fraction of output in these example economies is 0.3436 and 0.3033, respectively.

What happens if we endow the representative consumer with an initial capital stock $K_0 = 0.03$, which is far below the eventual steady-state level? We know generally that there will be growth to the steady-state, but little more.

The evolution of the capital stock under both assumptions about the government spending technology is plotted in Figure (12.1). The solid line gives the evolution with the high-return government spending technology (that is, equation (BC2)), while the dotted line gives the evolution with the low-return technology (that is, equation (BC1)). Notice that the economy based on equation (BC2) is initially poorer and slower-growing than the other economy. This is because, at low levels of capital, government spending is not very productive and is a serious drag on the economy. As capital accumulates and the complementarities with government spending kick in, growth accelerates and the economy based on equation (BC2) surpasses the economy based on equation (BC1).

In the same way, the time path of consumption is plotted in Figure (12.2). Finally, the real interest rate in these economies is plotted in Figure (12.3). For more about how to calculate the real interest rate in these models, please see the next section.

The Real Interest Rate

Now we turn our attention to the effect of permanent changes in government spending on the equilibrium real interest rate in this model. Recall that in infinite-horizon capital accumulation models, like the one we are studying here, it usual to assume there is a closed economy, so the representative household does not have access to a bond market. In this setting, the equilibrium interest rate becomes the interest rate at which the household, if offered the opportunity to use a bond market, would not do so. In other words, there is, as usual, no net borrowing or lending in a closed economy. We will refer to this condition as a market-clearing condition in the bond market, or simply market-clearing for short.

We shall see that, during the transition period while capital is still being accumulated, the

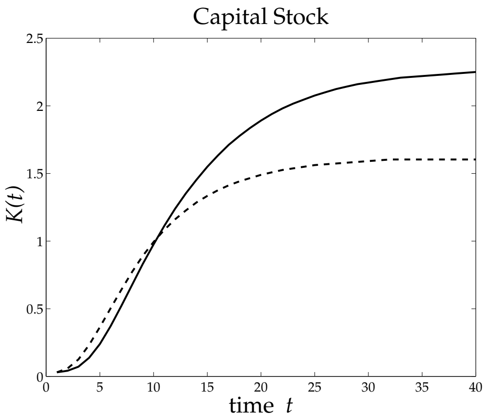

Figure 12.1: Evolution of capital stock. The solid line gives K_t assuming that government purchases affect output as in equation (BC2) and the dotted line assuming they affect output as in equation (BC1).

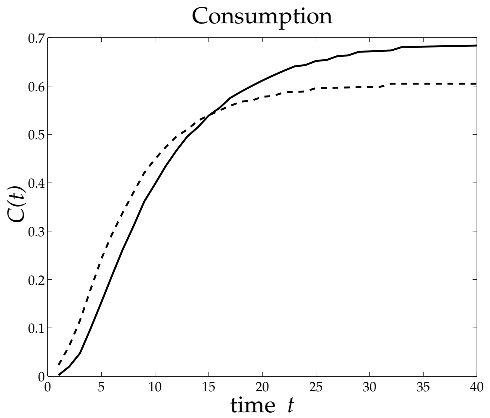

Figure 12.2: Time path of consumption. The solid line gives C_t assuming that government purchases affect output as in equation (BC2) and the dotted line assuming they affect output as in equation (BC1).

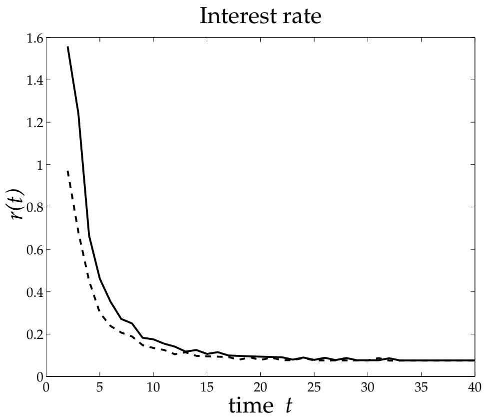

Figure 12.3: Time path of interest rates. The solid line gives r_t assuming that government purchases affect output as in equation (BC2) and the dotted line assuming they affect output as in equation (BC1).

interest rate is decreasing in the capital stock. At the steady state, however, when consumption is constant, the equilibrium interest rate will just be ρ, the discount rate. Because permanent changes in government spending lead eventually to a new steady-state, at which consumption is constant, permanent changes in the level of government spending will not affect the equilibrium interest rate at the steady state.

The easy way to see this is to notice that if the representative household has some endowment stream $\{e_t\}_{t=0}^{\infty}$, and the interest rate satisfies:

$$1 + r_t = \frac{1}{\beta}\frac{U'(C_t = e_t)}{U'(C_{t+1} = e_{t+1})},$$

then there will be no net borrowing or lending across periods. In our case the endowment stream $\{e_t\}_{t=0}^{\infty}$ is the result of a capital accumulation process which eventually reaches a steady state at which $e_t = e_{t+1} = e_{\text{ss}}$. Hence at a steady-state:

$$1 + r_{\text{ss}} = \frac{1}{\beta}\frac{U'(C_{\text{ss}})}{U'(C_{\text{ss}})} = \frac{1}{\beta} = 1 + \rho, \text{ so:}$$
$$r_{\text{ss}} = \rho.$$

No matter what the eventual steady-state level of capital, at the steady-state consumption becomes smooth, which forces the equilibrium interest rate to the discount rate. If $r_{\text{ss}} > \rho$ the household would wish to save on the bond market (consuming below endowment and thus violating market-clearing) and if $r_{\text{ss}} < \rho$ then then the household would wish to borrow on the bond market (consuming above endowment and again violating market-clearing).

12.2 Temporary Changes in Government Spending

Studying temporary changes in government spending requires studying the transition path of an economy from one steady-state to another and then back again. Imagine an economy of the type we studied in the previous section, in which the government is spending some low but constant amount G_0 each period. As time goes forward, the capital stock and consumption converge to their steady-state levels and the real interest rate converges to the discount rate. Suddenly the government must fight an expensive war. Government spending shoots up to some high level G_1 for a relatively short period of time. During the war, the capital stock will begin to transition to the steady-state implied by the new spending level G_1. Since wars tend to be short it may never get there. When the war is over, government spending drops to its accustomed pre-war level of G_0, and the capital stock slowly returns from wherever it was when the war ended to the old steady-state.

Analytically determining the trajectories of capital, consumption and the interest rate under temporary shifts in government spending is beyond the scope of this chapter. However, we can easily simulate them numerically, using precisely the same techniques we did to study the growth experience of economies.

All of the figures that follow make the following assumptions: That in periods 1-5 the economy is at its pre-war steady-state, that in periods 6-15 the economy is in a war, with increased government spending, and in periods 16-30 the economy is back at peace. During the war the economy begins its transition to a war steady-state, but the relatively short duration of the war prevents it from ever reaching that steady-state. After the war the economy transitions slowly back to its pre-war steady-state. We are also assuming that in the last period of peace before the war (period 5) the population learns of the impending war, and that in the last period of the war before peace begins again (period 15) the population learns of the coming peace.

The parameters used here are exactly those used in the section on transitions in the example economies (page 120) above. In addition, the peacetime spending level is $G_0 = 0$ and the wartime spending level is $G_1 = 0.4$.

The evolution of the capital stock under both assumptions about the government spending technology is plotted in Figure (12.4). The solid line gives the evolution with the high-return government spending technology (that is, equation (BC2)), while the dotted line gives the evolution with the low-return technology (that is, equation (BC1)).

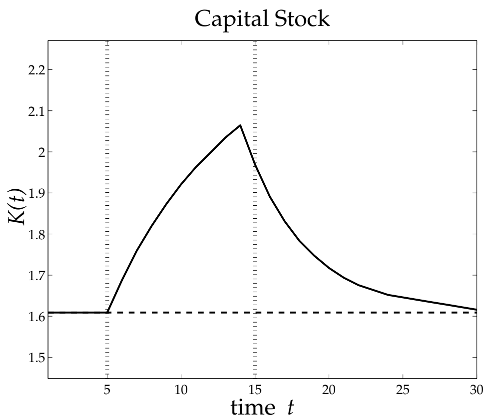

Figure 12.4: Time path of the capital stock before, during and after a war. The solid line gives K_t assuming that government purchases affect output as in equation (BC2) and the dotted line assuming they affect output as in equation (BC1).

In the same way, the time path of consumption is plotted in Figure (12.5). Finally, the *real interest rate* in these economies is plotted in Figure (12.6). It is surprising to note that sometimes the real interest rate is negative. From the section on the real interest rate (on page 120 above) we know that, given consumption decisions C_t and C_{t+1} that r_t must satisfy:

$$r_t = \frac{1}{\beta}\frac{U'(C_t)}{U'(C_{t+1})} - 1.$$

If C_{t+1} is quite small relative to C_t, then $U'(C_{t+1})$ will be large relative to $U'(C_t)$ and r_t might be negative. A negative real interest rate occurs in precisely those periods in which today's consumption must be high relative to tomorrow's, as in the last period of peacetime before the war, in order to prevent agents from carrying wealth forward into the next period. At the ends of wars, when today's consumption is low relative to tomorrow's (think March, 1945), real interest rates are quite high, to dissuade borrowing.

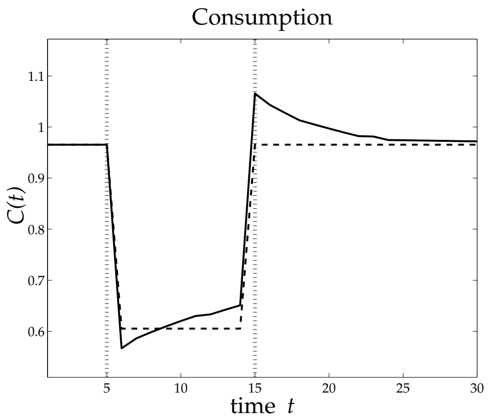

Figure 12.5: Time path of consumption before, during and after a war. The solid line gives C_t assuming that government purchases affect output as in equation (BC2) and the dotted line assuming they affect output as in equation (BC1).

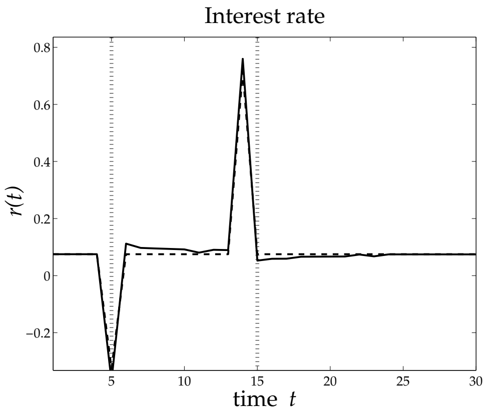

Figure 12.6: Time path of the interest rate before, during and after a war. Note the very low interest rates prevalent in the last period before the war and the generally higher interest rates during the war. The solid line gives r_t assuming that government purchases affect output as in equation (BC2) and the dotted line assuming they affect output as in equation (BC1).

In general, Barro presents evidence that, during wartime, interest rates tend to increase. That fits well with the experience of the second model presented here, the one in which government purchases affect output as in equation (BC2).

12.3 Social Security

The Social Security system is one of the largest components of U.S. government spending. There are some interesting theoretical issues associated with it that are worth examining. Social Security is an old-age pension system, in which young workers pay into a general fund with a *payroll tax* of about 7% of wages and old retirees receive payments from this same general fund. Thus although it maintains the illusion of being a national savings scheme (and many politicians and voters are convinced that it is exactly that) is in fact an *unfunded* or *pay-as-you-go* pension scheme. In an unfunded pension system, payments to retirees are paid for by taxes levied on the current young.

Other countries have adopted *funded* pension schemes, which are essentially forced savings systems. In a funded pension system, young workers are taxed, with the proceeds going to an individual account, invested in some securities (the precise type of investment mix, and whether these investments are under the control of the government or the worker vary from country to country). When workers become old and retire, they draw down their accumulated stock of savings.

Consider a world in which there are two types of agents: Young workers who earn an amount y in their working years, and old retirees who earn nothing. This is clearly a vast simplification over reality, since, in particular, the retirement date is exogenous. However, even this simple model will help us think clearly about pension schemes. A generation born in period t will have preferences over consumption while young C_0^t and old C_1^t of:

$$U(C_0^t, C_1^t) = 2\sqrt{C_0^t} + 2\beta\sqrt{C_1^t}.$$

Where $0 < \beta < 1$ reflects a preference for consumption while young.

Each period t there are N_t new young workers born, each of whom produces y with certainty in their youth. The youth population N_t evolves as:

$$N_{t+1} = (1 + n)N_t.$$

There is a bond market which pays a constant, riskless, real interest rate of $r \geq 0$, paid "overnight" on savings. Where does this bond market come from? We will not say here, leaving it simply outside of the scope of the model. If you are bothered by this, however, imagine that a certain portion of the population, instead of being workers, are entrepreneurs, who will accept funds from workers, use them as capital in a productive process of some kind, and then use the output from that production to repay the workers (now old) with interest. The interest rate gets set as the result of competition among entrepreneurs for funds.

Funded Pension Systems

Begin with an analysis of a funded system. The government levies a *tax rate* of τ on young workers' income y, taking τy. Since the young workers do not affect y, this is equivalent to a lump-sum tax. The government invests τy on behalf of the young workers, realizes the common real rate of return r on it, and returns it to the agents when they are retired. In addition, workers of generation t may save an amount $S_t \geq 0$ in the bond market on their own. Assume that τ is small relative to the savings needs of agents. This will prevent them from attempting to set $S_t < 0$, and will save us having to check corner conditions.

Given τ and S_t, we can calculate an agent's expected consumption path C_0^t, C_1^t:

$$C_0^t = (1 - \tau)y - S_t \tag{12.13}$$

$$C_1^t = (1 + r)(\tau y + S_t). \tag{12.14}$$

Because the government has taken τy from the agent while young, he is left only with $(1 - \tau)y$ to split between consumption while young and own-savings, S_t. When old, the agent gets the benefit of both public (government forced) savings τy and private (own) savings S_t. Consumption while old is merely the total volume of savings times the prevailing gross interest rate $1 + r$.

We are now ready to find S_t for this agent. The agent maximizes $U(C_0^t, C_1^t)$ where C_0^t as a function of S_t is given by equation (12.13) and C_1^t by equation (12.14). Thus the agent solves:

$$\max_{S_t} \left\{ 2\sqrt{(1 - \tau)y - S_t} + 2\beta\sqrt{(1 + r)(\tau y + S_t)} \right\}.$$

Assuming that the constraint $S_t \geq 0$ will not be binding, we take the derivative of this function with respect to S_t and set it to zero to find the optimal value of S_t. So:

$$\frac{-1}{\sqrt{(1 - \tau)y - S_t}} + \frac{\beta\sqrt{1 + r}}{\sqrt{\tau y + S_t}} = 0.$$

We cross-multiply to find:

$$\begin{aligned} \left(\beta\sqrt{(1 + r)}\right)\left(\sqrt{(1 - \tau)y - S_t}\right) &= \sqrt{\tau y + S_t}, \\ [\beta^2(1 + r)][(1 - \tau)y - S_t] &= \tau y + S_t, \\ [\beta^2(1 + r)](1 - \tau)y - \tau y &= [1 + \beta^2(1 + r)]S_t, \text{ and:} \\ \beta^2(1 + r)y - \tau y[1 + \beta^2(1 + r)] &= [1 + \beta^2(1 + r)]S_t. \end{aligned}$$

Dividing both sides by $1 + \beta^2(1 + r)$ produces:

$$S_t = \frac{\beta^2(1 + r)}{1 + \beta^2(1 + r)}y - \tau y. \tag{12.15}$$

Substituting back into equations (12.13) and (12.14) gives us optimal consumption choices in each period:

$$C_0^t = \frac{1}{1+\beta^2(1+r)}y,$$
$$C_1^t = \frac{\beta^2(1+r)^2}{1+\beta^2(1+r)}y.$$

Notice that the government-forced public savings policy τ does not affect the agent's choice of savings. If τ increases, the agent will merely decrease his choice of S_t.

If the government sets τ to exactly the agent's desired savings rate, that is:

$$\tau = \frac{\beta^2(1+r)}{1+\beta^2(1+r)},$$

then $S_t = 0$ and all saving is done by the government.

Unfunded Pension Systems

Now we turn our attention to unfunded pension systems (also known as *pay as you go* systems), in which the government taxes the current young workers to pay the current old retirees. The key insight will be that unfunded pension systems will dominate funded pension systems if the population is growing quickly enough.

In period t there are N_t young workers and N_{t-1} old retirees who were born in period $t-1$ and are now old. If the government taxes each young worker an amount τ it raises total revenue of:

$$G = \tau N_t y.$$

If it distributes this equally among the old, each old agent will get G/N_{t-1} or:

$$\frac{G}{N_{t-1}} \equiv g_{t-1} = \tau y \frac{N_t}{N_{t-1}}.$$

Recall that the population is growing at a rate n so that $N_t = (1+n)N_{t-1}$. Hence:

$$g_{t-1} = \tau y(1+n).$$

Notice that, since the population growth rate is constant at n, g_t does not vary with time, so we write merely g.

Consider again the agent's budget constraints as a function of τ and S_t, equations (12.13) and (12.14) above, only now using the unfunded pension system:

$$C_0^t = (1-\tau)y - S_t, \text{ and:}$$
$$C_1^t = (1+r)S_t + (1+n)\tau y.$$

We could solve this explicitly for S_t as a function of τ, y, n, r in much the same way that we did above (in fact, this is a good exercise to do on your own), but instead we are will simply provide intuition for the agent's choices.

If $n \neq r$ then the agent is no longer indifferent between public and private savings. If $n < r$, then public savings make the agent worse off. As τ increases more and more of the agent's wealth is being used in a relatively low-return activity. Agents would complain bitterly to their government about this (apparent) waste of their money.

On the other hand, if $n > r$, then the agent would prefer to save entirely by using the government pension system. Agents would demand that the system be increased until their private savings (in the relatively inefficient bond market) fell to zero.

Exercises

Exercise 12.1 (Easy)
For each of the following questions provide a brief answer.

1. (True, False or Uncertain) All things being equal, there is more total savings under a funded than under an unfunded pension system.
2. For the U.S., at the moment, is $n > r$?
3. Name three items in the Federal budget that account for more than 20% of all government expenditures (each).

Exercise 12.2 (Easy)
Assume that every dollar spent by the government augments total output by ϕ, where $0 < \phi < 1$. Assume that total private output is fixed at Y and that the government pays for its expenditures with lump-sum taxes. What is the absolute maximum amount of government spending, G? At this level, how much does the household consume and invest?

Exercise 12.3 (Moderate)
For this exercise assume that the representative household lives for only two periods and has preferences over consumption streams $\{C_0, C_1\}$ given by:

$$U(C_0) + \beta U(C_1),$$

where $\beta = 1/(1 + \rho)$ and $\rho > 0$. Here assume that $U' > 0, U'' < 0$. The household has a constant endowment stream $\{Y, Y\}$ which is not affected by government spending. Any government spending must be paid for by lump-sum taxes on the representative household. There is no capital stock. This is a closed economy. Answer the following questions:

1. Assume that the government spends the same amount G each period. What is the market-clearing interest rate, r_0?

2. Assume that the government spends different amounts in each period, $\{G_0, G_1\}$ and that $G_0 > G_1$. Now what is the market-clearing interest rate r_0^*?

3. Which is greater, r_0^* or r_0? Does this fit with your intuition about the effect of temporary government spending?

Exercise 12.4 (Moderate)
Consider again the model of Section 12.3 above. Calculate S_t explicitly when the return on public savings is n and the return to private savings is r. Assume $n \neq r$ and τ is small.

Exercise 12.5 (Moderate)
Grace lives for two periods. She has preferences over consumption streams c_0, c_1 of:

$$u(c_0, c_1) = \ln(c_0) + \beta \ln(c_1),$$

where $0 < \beta \leq 1$. Grace is endowed with one unit of time each period. In the first period, she can divide her time between working in a low-wage job at a wage of $w = 1$ or attending S hours of school. Grace earns nothing while in school, but she is augmenting her *human capital.* In the second period of life, Grace spends all of her time at her high-wage job, earning AK_1 where K_1 is her human capital and $A > 1$. Human capital is augmented by schooling by the simple formula $K_1 = S$, so given a choice for S, Grace earns $1 - S$ while young and AS while old. There is no bond market.

The government is interested in helping Grace go to school. It levies a lump-sum tax of G on Grace when she is young and uses it to augment her human capital so that $K_1 = S + \phi G$ where $\phi > 0$. Answer the following questions:

1. Assume $G = 0$. Find Grace's optimal schooling choice S and human capital K_1.

2. Assume $G > 0$. Find Grace's optimal schooling choice S and human capital K_1. Remember that K_1 is affected directly by G. Show that S is decreasing in G and that K_1 is decreasing in G is $\phi < 1$.

3. Now assume that the human capital augmentation is a a straight subsidy from the government, that is, the government has taxed someone else to pay for Grace's schooling, so she is not taxed at all while young. Now how do S and K_1 vary with G?

Variable	Definition
B_t	Household savings at the end of period t (if positive), or household debt (if negative).
C, C_t	Consumption by the household (in period t).
C^a, C^b	Specific consumption levels used in an example.
$\{e_t\}_{t=0}^{\infty}$	Sequence of household endowments over time.
G	Government spending (usually assumed to be constant).
I_t	Household's investment in the capital stock at time t.
K_t, K_0	Capital stock in period t (initial capital stock).
Y_t^P	Output from private productive processes.
Y_t^G	Output from government production which is refunded to the household.
Y_t	Total output in period t, the sum of private and government output.
β	Household discount factor, usually assumed to be $1/(1+\rho)$.
ρ	Household discount rate.
δ	Depreciation rate of capital.
$V(\{C_t\}_{t=0}^{\infty})$	Household's preferences over an entire stream of consumptions.
$U(C_t), u(C_t)$	Household utility in period t from consumption in that period of C_t.
C_0^t, C_1^t	Consumption of generation born in period t while young and old.
N_t	Population in period t.
n	Growth rate of population.
r	Real interest rate.
τ	Income tax rate.
y	Household income.
S_t	Household gross private savings.
G	The government's realized revenue from taxes on young.
g_t	Government's per-capita payments to the old, G/N_t, in period $t+1$.

Table 12.1: Notation for Chapter 12

Chapter 13

The Effect of Taxation

Taxes affect household behavior via income and substitution effects. The income effect is straightforward: as taxes go up, households are poorer and behave that way. For example, if leisure is a normal good, then higher taxes will induce consumers to consume less leisure. The substitution effect is trickier, but it can be much more interesting. Governments levy taxes on observable and verifiable actions undertaken by households. For example, governments often tax consumption of gasoline and profits from sales of capital assets, like houses. These taxes increase the costs to the households of undertaking the taxed actions, and the households respond by adjusting the actions they undertake. This can lead to outcomes that differ substantially from those intended by the government.

Since optimal tax policy is also a subject of study in microeconomics and public finance courses, we shall concentrate here on the effect of taxation on labor supply and capital accumulation. When modeling labor supply decisions we are going to have a representative agent deciding how to split her time between labor supply and leisure. Students might object on two grounds: First, that the labor supply is quite inelastic (since everyone, more or less, works, or tries to) and second, that everyone puts in the same number of hours per week, and the variation in leisure comes not so much in time as in expenditure (so that richer people take more elaborate vacations).

The representative household stands for the decisions of millions of underlying, very small, households. There is, to name only one example, mounting evidence that households change the timing of their retirement on the basis of tax policy. As taxes increase, more and more households choose to retire. At the level of the representative household, this appears as decreasing labor supply. As for the observation that everyone puts in either 40 hours a week or zero, this misses some crucial points. The fact is that jobs differ significantly in their characteristics. Consider the jobs available to Ph.D. economists: they range from Wall Street financial wizard, big-time university research professor, to small-time college instructor. The fact is that a Wall Street financial wizard earns, on her first day on the job, two or three times as much as a small-time college instructor. Of course, college teachers have a much more relaxed lifestyle than financiers (their salary, for example, is computed assuming that they only work nine months out of the year). The tax system can easily distort a freshly-minted Ph.D.'s choices: Since she consumes only the after-tax portion of her income, the Wall Street job may only be worth 50% more, after taxes, than the college instructor's job. The point is not that every new economics Ph.D. would plump for the college instructor's job, but that, as the tax on high-earners increased, an increasing fraction would. Again, we can model this with a representative household choosing how much leisure to consume.

We begin with a general overview of tax theory, discuss taxation of labor, then taxation of capital and finally consider attempts to use the tax system to remedy income (or wealth) inequality.

13.1 General Analysis of Taxation

In this section we will cast the problem of taxation in a very general framework. We will use this general framework to make some definitions and get some initial results.

Notation

Assume that the household take some observed action a in A (this discussion generalizes to the case when a is a vector of choices). For example, a could be hours worked, number of windows in one's house, or the number of luxury yachts the household owns (or, if a is a vector, all three). The set A is the set of allowed values for a, for example 0 to 80 hours per week, $\{0, 1, 2, \ldots, 500\}$ windows per house or 0 to ten luxury yachts (where we are assuming that no house may have more than 500 windows and no household can use more than 10 luxury yachts).

The government announces a *tax policy* $\mathcal{H}(a; \psi)$, where $\mathcal{H}(\psi) : A \to \mathbb{R}$. That is, a tax policy is a function mapping observed household choices into a *tax bill* which the household has to pay (if positive), or takes as a subsidy to consumption (if negative). The term ψ (which may be a vector) is a set of parameters to the tax policy (for example, deductions). The household is assumed to know the function $\mathcal{H}(a; \psi)$ and ψ before it takes action a.

An example of a tax policy $\mathcal{H}$ is the *flat income tax.* In a flat income tax, households pay a fixed fraction of their income a in taxes, so $\psi = \tau$, where τ is the flat tax rate. A more complex version of the flat income tax allows for *exemptions* or *deductions,* which are simply a portion of income exempt from taxation. If the exempt income is E, then the parameters to the tax system are $\psi = \{E, \tau\}$ and $\mathcal{H}(a;\psi)$ is:

$$\mathcal{H}(a;\psi) = \begin{cases} 0, & a \leq E \\ \tau(a-E), & a \geq E. \end{cases}$$

Definitions

We can use our notation to make some useful definitions. The *marginal tax rate* is the tax paid on the next increment of a. So if one's house had 10 windows already and one were considering installing an 11th window, the marginal tax rate would be the increase in one's tax bill arising from that 11th window. More formally, the marginal tax rate at a is:

$$\frac{\partial \mathcal{H}(a;\psi)}{\partial a}.$$

Here we are assuming that a is a scalar and smooth enough so that $\mathcal{H}(a;\psi)$ is at least once continuously differentiable. Expanding the definition to cases in which $\mathcal{H}(a;\psi)$ is not smooth in a (in certain regions) is straightforward, but for simplicity, we ignore that possibility for now.

The *average tax rate* at a is defined as:

$$\frac{\mathcal{H}(a;\psi)}{a}.$$

Note that a flat tax with $E = 0$ has a constant marginal tax rate of τ, which is just equal to the average tax rate.

If we take a to be income, then we say that a tax system is *progressive* if it exhibits an increasing marginal tax rate, that is if $\mathcal{H}'(a;\psi) > 0$. In the same way, a tax system is said to be *regressive* if $\mathcal{H}'(a;\psi) < 0$.

Household Behavior

Let us now turn our attention to the household. The household has some technology for producing income $\mathcal{Y}$[1] that may be a function of the action a, so $\mathcal{Y}(a)$. If a is hours worked, then $\mathcal{Y}$ is increasing in a, if a is hours of leisure, then $\mathcal{Y}$ is decreasing in a and if a is house-windows then $\mathcal{Y}$ is not affected by a. The household will have preferences directly over action a and *income net of taxation* $\mathcal{Y}(a) - \mathcal{H}(a;\psi)$. Thus preferences are:

$$U\left[a, \mathcal{Y}(a) - \mathcal{H}(a;\psi)\right].$$

There is an obvious maximization problem here, and one that will drive all of the analysis in this chapter. As the household considers various choices of a (windows, hours, yachts),

[1] We use the notation $\mathcal{Y}$ here to mean income to emphasize that income is now a function of choices a.

it takes into consideration both the direct effect of a on utility and the indirect effect of a, through the tax bill term $\mathcal{Y}(a) - \mathcal{H}(a;\psi)$. Define:

$$V(\psi) \equiv \max_{a \in A} U\left[a, \mathcal{Y}(a) - \mathcal{H}(a;\psi)\right].$$

For each value of ψ, let $a_{\max}(\psi)$ be the choice of a which solves this maximization problem. That is:

$$V(\psi) = U\left[a_{\max}(\psi), \mathcal{Y}(a_{\max}(\psi)\right] - \mathcal{H}\left[a_{\max}(\psi);\psi)\right].$$

Assume for the moment that U, $\mathcal{Y}$ and $\mathcal{H}$ satisfy regularity conditions so that for every possible $|psi$ there is only one possible value for $a_{\max}$.

The government must take the household's response $a_{\max}(\psi)$ as given. Given some tax system $\mathcal{H}$, how much revenue does the government raise? Clearly, just $\mathcal{H}\left[a_{\max}(\psi);\psi\right]$. Assume that the government is aware of the household's best response, $a_{\max}(\psi)$, to the government's choice of tax parameter ψ. Let $\mathcal{T}(\psi)$ be the revenue the government raises from a choice of tax policy parameters ψ:

$$\mathcal{T}(\psi) = \mathcal{H}\left[a_{\max}(\psi);\psi\right]. \tag{13.1}$$

Notice that the government's revenue is just the household's tax bill.

The functions $\mathcal{H}(a;\psi)$ and $\mathcal{T}(\psi)$ are closely related, but you should not be confused by them. $\mathcal{H}(a;\psi)$ is the tax system or tax policy: it is the legal structure which determines what a household's tax bill is, given that household's behavior. Households choose a value for a, but the tax policy must give the tax bill for all possible choices of a, including those that a household might never choose. Think of $\mathcal{H}$ as legislation passed by Congress. The related function $\mathcal{T}(\psi)$ gives the government's actual revenues under the tax policy $\mathcal{H}(a;\psi)$ when households react optimally to the tax policy. Households choose the action a which makes them happiest. The mapping from tax policy parameters ψ to household choices is called $a_{\max}(\psi)$. Thus the government's actual revenue given a choice of parameter ψ, $\mathcal{T}(\psi)$, and the legislation passed by Congress, $\mathcal{H}(a;\psi)$, are related by equation (13.1) above.

The Laffer Curve

How does the function $\mathcal{T}(\psi)$ behave? We shall spend quite a bit of time this chapter considering various possible forms for $\mathcal{T}(\psi)$. One concept to which we shall return several times is the *Laffer curve.* Assume that, if a is fixed, that $\mathcal{H}(a;\psi)$ is increasing in ψ (for example, ψ could be the tax rate on house windows). Further, assume that if ψ is fixed, that $\mathcal{H}(a;\psi)$ is increasing in a. Our analysis would go through unchanged if we assumed just the opposite, since these assumptions are simply naming conventions.

Given these assumptions, is $\mathcal{T}$ necessarily increasing in ψ? Consider the total derivative of $\mathcal{T}$ with respect to ψ. That is, compute the change in revenue of an increase in ψ, taking in to account the change in the household's optimal behavior:

$$\frac{d\mathcal{T}(\psi)}{d\psi} = \frac{\partial\mathcal{H}\left[a_{\max}(\psi);\psi\right]}{\partial a}\frac{\partial a_{\max}(\psi)}{\partial\psi} + \frac{\partial\mathcal{H}\left[a_{\max}(\psi);\psi\right]}{\partial\psi}. \tag{13.2}$$

The second term is positive by assumption. The first term is positive if $a_{\max}$ is increasing in ψ. If $a_{\max}$ is decreasing in ψ, and if the effect is large enough, then the government revenue function may actually be decreasing in ψ despite the assumptions on the tax system $\mathcal{H}$. If this happens, we say that there is a Laffer curve in the tax system.

A note on terms: the phrase "Laffer curve" has become associated with a bitter political debate. We are using it here as a convenient shorthand for the cumbersome phrase, "A tax system which exhibits decreasing revenue in a parameter which increases government revenue holding household behavior constant because the household adjusts its behavior in response". Do tax systems exhibit Laffer curves? Absolutely. For example, a Victorian-era policy which levied taxes on the number of windows (over some minimum number designed to exempt the middle class) in a house, over a span of years, resulted in grand houses with very few windows. As a result, the *hoi polloi* began building more modest homes also without windows and windowlessness became something of a fashion. Increases in the window tax led, in the long term, to decreases in the revenue collected on the window tax. The presence of a Laffer curve in the U.S. tax system is an empirical question outside the scope of this chapter.

Finally, the presence of a Laffer curve in a tax system does not automatically mean that a tax cut produces revenue growth. The parameter set ψ must be in the downward-sloping region of the government revenue curve for that to be the case. Thus the U.S. tax system could indeed exhibit a Laffer curve, but only at very high average tax rates, in which case tax cuts (given the current low level of taxation) would lead to decreases in revenue.

Lump-sum Taxes

Now consider the results if the government introduced a tax system with the special characteristic that the tax bill did not depend on the household's decisions. That is,

$$\frac{\partial\mathcal{H}(a;\psi)}{\partial a} = 0,$$

for all choices of ψ. Notice that the household's optimal decisions may still change with ψ, but that the government's revenue will not vary as $a_{\max}$ varies. Let us determine what happens to the derivative of the government revenue function $\mathcal{T}$ from equation (13.2) above:

$$\begin{aligned}\frac{d\mathcal{T}(\psi)}{d\psi} &= \frac{\partial\mathcal{H}\left[a_{\max}(\psi);\psi\right]}{\partial a}\frac{\partial a_{\max}(\psi)}{\partial\psi} + \frac{\partial\mathcal{H}\left[a_{\max}(\psi);\psi\right]}{\partial\psi} \\ &= (0)\left[\frac{\partial a_{\max}(\psi)}{\partial\psi}\right] + \frac{\partial\mathcal{H}\left[a_{\max}(\psi);\psi\right]}{\partial\psi} \\ &= \frac{\partial\mathcal{H}\left[a_{\max}(\psi);\psi\right]}{\partial\psi}.\end{aligned}$$

This is always greater than zero by assumption. Hence there is never a Laffer curve when the tax system has the property that $\partial\mathcal{H}/\partial a = 0$, that is, with lump-sum taxes.

Taxes which do not vary with household characteristics are known as *poll taxes* or *lump-sum taxes.* Poll taxes are taxes that are levied uniformly on each person or "head" (hence the name). Note that there is no requirement that lump sum taxes be uniform, merely that household actions cannot affect the tax bill. A tax lottery would do just as well. In modern history there have been relatively few examples of poll taxes. The most recent use of poll taxes was in England, where they were used from 1990-1993 to finance local governments. Each council (roughly equivalent to a county) divided its expenditure by the number of adult residents and delivered tax bills for that amount. Your correspondent was, at the time, an impoverished graduate student living in the Rotherhithe section of London, and was presented with a bill for £350 (roughly $650 at the time). This policy was deeply unpopular and led to the "Battle of Trafalgar Square"—the worst English riot of the 20th century. It is worth noting that this tax did not completely meet the requirements of a lump sum tax since it did vary by local council, and, in theory, households could affect the amount of tax they owed by moving to less profligate councils, voting Conservative or rioting. These choices, though, were more or less impossible to implement in the short-term, and most households paid.

Lump-sum taxes, although something of a historical curiosity, are very important in economic analysis. As we shall see in the next section, labor supply responds very differently to lump-sum taxes than to income taxes.

The Deadweight Loss of Taxation

Lump sum taxes limit the amount of *deadweight loss* associated with taxation. Consider the effect of an increase in taxes which causes an increase in government revenue: revenue increases slightly and household income net of taxes decreases by slightly more than the revenue increase. This difference is one form of deadweight loss, since it is revenue lost to both the household and the government.

It is difficult to characterize the deadweight loss of taxation with the general notation we have established here (we will be much more precise in the next section). However, we will be able to establish that the deadweight loss is increasing in the change of household behavior. That is, the more sensitive $a_{\max}(\psi)$ is to ψ, the larger the deadweight loss.

Consider a tax policy $\mathcal{H}(a;\psi)$ and two different parameter sets for the tax policy, ψ_0 and ψ_1. Assume that, for fixed a, $\mathcal{H}(a;\psi_0) < \mathcal{H}(a;\psi_1)$. The household's utility at each of the tax parameters is:

$$V(\psi_0) = U\left(a_{\max}(\psi_0), \mathcal{Y}\left[a_{\max}(\psi_0)\right] - \mathcal{H}\left[a_{\max}(\psi_0), \psi_0\right]\right), \text{ and:}$$
$$V(\psi_1) = U\left(a_{\max}(\psi_1), \mathcal{Y}\left[a_{\max}(\psi_1)\right] - \mathcal{H}\left[a_{\max}(\psi_1), \psi_1\right]\right).$$

The claim is that the change in household net income exceeds the change in government revenue, or:

$$\begin{aligned}(\mathcal{Y}\left[a_{\max}(\psi_0)\right] - \mathcal{H}\left[a_{\max}(\psi_0), \psi_0\right]) - (\mathcal{Y}\left[a_{\max}(\psi_1)\right] - \mathcal{H}\left[a_{\max}(\psi_1), \psi_1\right]) \\ > \mathcal{T}(\psi_1) - \mathcal{T}(\psi_0).\end{aligned} \tag{13.3}$$

Recall that $\mathcal{T}(\psi) = \mathcal{H}\left[a_{\max}(\psi);\psi\right]$. Equation (13.3) is true only if:

$$\mathcal{Y}[a_{\max}(\psi_0)] > \mathcal{Y}[a_{\max}(\psi_1)].$$

That is, the more household gross (that is, pre-tax) income falls in response to the tax, the greater the deadweight loss. But since household gross income is completely under the household's control through choice of a, this is tantamount to saying that the more a changes, the greater the deadweight loss. This is a very general result in the analysis of taxation: the more the household can escape taxation by altering its behavior, the greater the deadweight loss of taxation.

If we further assume that there are no pure income effects in the choice of a, then lump-sum taxes will not affect the household's choice of a and there will be no deadweight loss to taxation (a formal proof of this point is beyond the scope of this chapter). The assumption of no income effects is relatively strong, but, as we shall see later, even without it lump-sum taxes affect household behavior very differently than income taxes.

13.2 Taxation of Labor

In this section we shall assume that households choose only their effort level or labor supply L. We will assume that they have access to a technology for transforming labor into the consumption good of wL. Think of w as a wage rate. Although we will not clear a labor market in this chapter, so w is not an endogenous price, we can imagine that all households have a backyard productive technology of this form.

Households will enjoy consumption and dislike effort, but will be unable to consume without expending effort. They will balance these desires to arrive at a labor supply decision. Government taxation will distort this choice and affect labor supply.

A Simple Example

As a first step, consider a household with a utility function over consumption C and effort L of the form:

$$U(C, L) = 2\sqrt{C} - L.$$

The household's income takes the form:

$$\mathcal{Y}(L) = wL.$$

Assume that there is a simple flat tax, so the tax policy is:

$$\mathcal{H}(L; \tau) = \tau\mathcal{Y}(L).$$

Hence the household's budget constraint becomes:

$$C = \mathcal{Y}(L) - \mathcal{H}(L; \tau) = wL - \tau wL = (1 - \tau)wL.$$

Substituting this budget constraint into the household's utility function produces:

$$V(\tau) = \max_{L} \left\{ 2\sqrt{(1-\tau)wL} - L \right\}.$$

This function is just the household's utility given a tax rate τ. We can solve the maximization problem to find $V(\tau)$ directly. Take the derivative with respect to the single choice variable, labor supply L, and set it to zero to find:

$$\sqrt{\frac{(1-\tau)w}{L}} - 1 = 0.$$

Solving for L produces:

$$L(\tau) = (1 - \tau)w.$$

We can substitute the labor supply decision $L(\tau)$ back into the government's tax policy to find the government's revenue function:

$$\mathcal{T}(\tau) = \mathcal{H}\left[L(\tau); \tau\right] = \tau w L(\tau) = w^2\tau(1 - \tau).$$

Does this system exhibit a Laffer curve? Indeed it does. Clearly, $\mathcal{T}(\tau)$ in this case is simply a parabola with a maximum at $\tau = 0.5$. (See Figure (13.1)).

The effect of the income tax was to drive a wedge between the productivity of the household (constant at w) and the payment the household received from its productive activity. The household realized an *effective wage rate* of $(1 - \tau)w$. As the flat tax rate τ moved to unity, the effective wage rate of the household falls to zero and so does its labor supply. Compare this tax structure with one in which the household realizes the full benefit of its effort, after paying its fixed obligation. Thus we turn our attention next to a lump-sum tax.

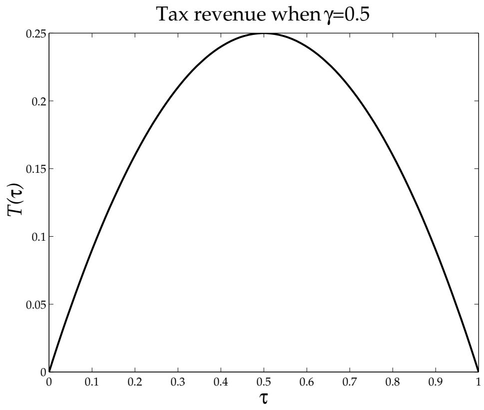

Figure 13.1: A government revenue function that exhibits a Laffer curve.

A Lump-sum Tax

Now let us introduce a lump-sum tax of amount τ_L.[2] No matter what income the household accumulates, it will be forced to pay the amount τ_L. On the other hand, after paying τ_L, the household consumes all of its income. Previously, with the income tax, the household faced an effective wage rate of $(1-\tau)w$, which decreased as τ increased. Now the household's effective wage will be w (after the critical income of τ_L is reached). Does this mean that effort will be unaffected by τ_L? Recall from the previous section that this will only happen if there are no wealth effects. Examining the utility function reveals that it is not homogeneous of degree 1 in wealth, hence we can expect labor supply to vary to with τ_L. In particular, since leisure is a normal good, we will expect that labor supply will be increasing in τ_L. The household's budget constraint, with this tax policy, becomes:

$$C = wL - \tau_L,$$

so the household's maximization problem is:

$$V(\tau_L) = \max_L \left\{ 2\sqrt{wL - \tau_L} - L \right\}.$$

The first-order condition for optimality is:

$$\frac{w}{\sqrt{wL - \tau_L}} - 1 = 0.$$

Solving for L produces:

$$L(\tau_L) = w + \frac{\tau_L}{w}.$$

[2]The notation τ_L is meant to imply lump-sum tax: there is a surfeit of notation involving τ and L in this chapter. Please refer to the table at the end if you become confused.

We see that labor supply is in fact increasing in the lump-sum tax amount τ_L. The household increases its labor supply by just enough to pay its poll tax obligation. What is the government revenue function? It is, in this case, simply:

$$\mathcal{T}(\tau_L) = \tau_L.$$

So there is no Laffer curve with a lump-sum tax (of course).

General Labor Supply and Taxation

With the assumption of a square-root utility function, we were able to derive very interesting closed-form solutions for labor supply and the government revenue function. Our results, though, were hampered by being tied to one particular functional form. Now we introduce a more general form of preferences (although maintaining the assumption of linear disutility of effort). We shall see that a Laffer curve is not at all a predestined outcome of income taxes. In fact, when agents are very risk-averse, and when zero consumption is catastrophic, we shall see that the Laffer curve vanishes from the income tax system.

Consider agents with preferences over consumption C and labor supply L of the form:

$$U(C, L) = \begin{cases} \frac{C^\gamma}{\gamma} - L, & \gamma \neq 0, \quad \gamma \leq 1 \\ \ln(C) - L, & \gamma = 0. \end{cases} \tag{13.4}$$

Notice the immediate difference when $0 < \gamma \leq 1$ and when $\gamma \leq 0$. In the former case, a consumption of zero produces merely zero utility, bad, but bearable; while in the latter case, zero consumption produces a utility of negative infinity, which is unbearable. Agents will do anything in their power to avoid any possibility of zero consumption when $\gamma \leq 0$. Recall that in our previous example (when $\gamma = 0.5$) labor supply dropped to zero as the income tax rate increased to unity. Something very different is going to happen here.

Given a distortionary income tax rate of τ, the household's budget constraint becomes:

$$C = (1 - \tau)wL,$$

as usual. The household's choice problem then becomes:

$$V(\tau) = \max_L \left\{ \frac{[(1-\tau)wL]^\gamma}{\gamma} - L \right\}.$$

The first-order necessary condition for maximization is:

$$[(1-\tau)w]^\gamma L^{\gamma-1} - 1 = 0.$$

This in turn implies that:

$$L^{1-\gamma} = [(1-\tau)w]^\gamma, \text{ so:}$$
$$L = [(1-\tau)w]^{\frac{\gamma}{1-\gamma}}.$$

Notice that if $\gamma < 0$, then L is decreasing in w.

The government revenue function $\mathcal{T}(\tau)$ is:

$$\mathcal{T}(\tau) = \tau w L(\tau) = \tau(1-\tau)^{\frac{\gamma}{1-\gamma}} w^{\frac{1}{1-\gamma}}.$$

The question becomes, when does this tax system exhibit a Laffer curve? This is tantamount to asking when, if ever, the government revenue function is decreasing in τ. We begin by taking the derivative of $\mathcal{T}$ with respect to τ:

$$\begin{aligned}
\mathcal{T}'(\tau) &= w^{\frac{1}{1-\gamma}} \left[(1-\tau)^{\frac{\gamma}{1-\gamma}} - \frac{\gamma}{1-\gamma}\tau(1-\tau)^{\frac{\gamma}{1-\gamma}-1}\right] \\
&= w^{\frac{1}{1-\gamma}}(1-\tau)^{\frac{\gamma}{1-\gamma}-1}\left[1-\tau-\frac{\gamma}{1-\gamma}\tau\right] \\
&= w^{\frac{1}{1-\gamma}}(1-\tau)^{\frac{\gamma}{1-\gamma}-1}\left[1-\frac{1}{1-\gamma}\tau\right].
\end{aligned}$$

Notice immediately that the derivative $\mathcal{T}'(\tau)$ has the same sign as the term:

$$\left[1-\frac{1}{1-\gamma}\tau\right],$$

since the term outside of the brackets is positive by the assumptions that $w > 0$ and $\gamma < 1$. Thus $\mathcal{T}'(\tau)$ will be negative only if:

$$\begin{aligned}
1-\frac{1}{1-\gamma}\tau &\leq 0, \text{ or:} \\
\tau &\geq 1-\gamma.
\end{aligned}$$

The tax rate τ must satisfy $0 \leq \tau \leq 1$. Thus we notice two things immediately: (1) If $\gamma \leq 0$ there is no Laffer curve, and (2) If $0 < \gamma < 1$, then there is a Laffer curve, and the peak occurs at $\tau = 1-\gamma$.

What is the real-world significance of this sharp break in behavior at $\gamma = 0$? Agents with $\gamma \leq 0$ are very risk-averse and are absolutely unwilling to countenance zero consumption (the real world equivalent would be something like bankruptcy). In addition, their labor supply is decreasing in the wage rate w. In contrast, agents with $\gamma > 0$ are less risk-averse (although by no means risk neutral), are perfectly willing to countenance bankruptcy and have labor supply curves which are increasing in the wage rate w. In a world with many households, each of whom has a different value of γ, and a government which imposes a common tax rate τ, we would expect greater distortions among those households that are less risk-averse and harder-working.

Finally, the reader may find it an instructive exercise to repeat this analysis with a lump-sum tax. Households will all respond to a lump-sum tax by increasing their labor effort by precisely the same amount, τ_L/w, no matter what their value of γ.

13.3 Taxation of Capital

Now we turn our attention to the problem of a household which owns some capital stock and a technology for transforming capital into output. We shall see that, if households are allowed to deduct investment from their taxes (that is, if investment is tax-exempt), then there will not be a Laffer curve in income taxes. If on the other hand investment is not tax-deductible, or only a portion is, then there will be a Laffer curve in income taxation.

The household lives forever and has preferences over consumption streams $\{C_t\}_{t=0}^{\infty}$ of:

$$\sum_{t=0}^{\infty} \beta^t u(C_t),$$

where $\beta = 1/(1+\rho)$. Here $\rho > 0$ is the discount rate.

The household begins life with an initial stock of capital $K_0 > 0$. In addition, the income of the household each period, Y_t, is:

$$Y_t = K_t^{\alpha},$$

where K_t is the household's capital stock in period t and α is a production parameter satisfying $0 < \alpha < 1$. The capital stock evolves according to the law of motion:

$$K_{t+1} = (1-\delta)K_t + I_t,$$

where I_t is investment in physical capital (a choice of the household) and δ is the depreciation rate of capital. The economy is closed, so there is no bond market.

Assume that the government's tax policy is a flat tax on income from capital. We will consider two forms: one in which investment is exempt and one in which it is not. Thus without the exemption, in period t, the legislative tax code requires households to pay:

$$\mathcal{H}_t(K_t;\tau) = \tau K_t^{\alpha}.$$

If investment is exempt, then the legislative tax code requires, in period t, that households pay:

$$\mathcal{H}_t(K_t;\tau) = \tau(K_t^{\alpha} - I_t).$$

The household chooses K_t (so it plays the role of a) in response to changes in the tax code. We will study the steady-state capital level as a function of taxes $K_{\mathrm{ss}}(\tau)$. Thus the steady-state revenue raised each period by the government is:

$$\mathcal{T}_{ss}(\tau) = \mathcal{H}_{ss}\left[K_{\mathrm{ss}}(\tau);\tau\right].$$

This will vary depending on whether investment is deductible or not.

The household's budget constraint is:

$$C_t + I_t + (\text{tax bill})_t = Y_t.$$

Begin by assuming that investment is non-deductible. The tax bill then becomes the tax rate τ times income Y_t, or:

$$(\text{tax bill})_t = \tau Y_t.$$

Hence the household's budget constraint becomes:

$$C_t = (1 - \tau)Y_t - I_t. \tag{13.5}$$

Now assume that investment is tax deductible. The government levies a tax at rate τ on every dollar earned above investment. This also sometimes called *paying for investment with pre-tax dollars.* That is:

$$(\text{tax bill})_t = \tau(Y_t - I_t).$$

The household's budget constraint now becomes:

$$C_t = (1 - \tau)(Y_t - I_t). \tag{13.6}$$

We shall see that, because the tax system in equation (13.5) raises the implicit price of investment, the steady-state level of capital will be distorted away from its first-best level. Thus as the tax rate increases, investment and the steady-state capital level fall, so there is a Laffer curve lurking in the tax system. In contrast, the tax system in equation (13.6) leaves the implicit price of investment in terms of output unaffected by the tax rate, hence we shall see that the steady-state capital level will be unaffected by the tax rate. As a result, the Laffer curve is banished from the system, and government revenues become linear in the tax rate τ.

Analysis When Investment is Not Exempt

We want to collapse the budget constraint in equation (13.5) and the law of motion for capital into one equation, giving consumption C_t as a function of K_t and K_{t+1}, where in period t the household takes K_t as given and chooses K_{t+1}. Thus write consumption C_t as:

$$C_t = (1 - \tau)K_t^{\alpha} + (1 - \delta)K_t - K_{t+1}. \tag{13.7}$$

Here we have substituted in K_t^{α} for income Y_t from the agent's technology.

The household's problem thus becomes:

$$\max_{\{K_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(C_t),$$

where C_t is given in equation (13.7). Take the derivative with respect to the capital choice K_{j+1} for some arbitrary period j (where we avoid taking derivatives with respect to capital in period T because t is the time index in the summation). Remember the trick to these problems: K_{j+1} will appear in period j and period $j + 1$. Hence optimality requires:

$$\beta^j u'(C_j)[-1] + \beta^{j+1} u'(C_{j+1})[\alpha(1-\tau)K_{j+1}^{\alpha-1} + 1 - \delta] = 0.$$

Divide by the common factor β^j. Now assume that a steady-state has been reached. At a steady-state $K_t = K_{t+1} = K_{ss}$ and $C_t = C_{t+1} = C_{ss}$. Hence:

$$u'(C_{ss}) = \beta u'(C_{ss})[\alpha(1-\tau){K_{ss}}^{\alpha-1} + 1 - \delta].$$

Recall that $\beta^{-1} = 1 + \rho$. Divide both sides by $\beta u'(C_{ss})$ to find:

$$1 + \rho = \alpha(1-\tau){K_{ss}}^{\alpha-1} + 1 - \delta.$$

Hence:

$$K_{ss} = (1-\tau)^{\frac{1}{1-\alpha}} \left(\frac{\alpha}{\rho+\delta}\right)^{\frac{1}{1-\alpha}}.$$

Notice immediately that the steady-state capital level is decreasing in the tax rate τ. Gross income each period at the steady-state is:

$$Y_{ss} = (1-\tau)^{\frac{\alpha}{1-\alpha}} \left(\frac{\alpha}{\rho+\delta}\right)^{\frac{\alpha}{1-\alpha}}.$$

Let $\mathcal{T}_t(\tau)$ be the tax revenue of the government each period when the tax rate is τ. At the steady-state:

$$\mathcal{T}_t(\tau) = \tau Y_{ss} = \tau(1-\tau)^{\frac{\alpha}{1-\alpha}} \left(\frac{\alpha}{\rho+\delta}\right)^{\frac{\alpha}{1-\alpha}}. \tag{13.8}$$

Since $0 < \alpha < 1$, $\mathcal{T}_t$ is a parabola with a peak at $\tau = 1 - \alpha$. Thus this tax system exhibits a Laffer curve.

Analysis When Investment is Exempt

As before, we begin by collapsing the budget constraint (now equation (13.6)) and the law of motion for capital into one equation. Thus write:

$$C_t = (1-\tau)(K_t^\alpha + (1-\delta)K_t - K_{t+1}).$$

Notice the difference that exempting investment makes. The entire right-hand-side is now multiplied by $1 - \tau$, so the household cannot escape taxation by altering its mix of investment and consumption.

Once again, we choose sequences of capital to maximize:

$$\max_{\{K_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(C_t),$$

where C_t is given above. The derivative with respect to K_{j+1} is now:

$$\beta^j u'(C_j)[-(1-\tau)] + \beta^{j+1} u'(C_{j+1})[(1-\tau)(\alpha K_{j+1}^{\alpha-1} + 1 - \delta)] = 0.$$

Divide the equation by the common factor $(1-\tau)\beta^j$. Notice that the tax rate has vanished. Now assume a steady-state. Hence:

$$u'(C_{\text{ss}}) = \beta u'(C_{\text{ss}})[\alpha K_{\text{ss}}{}^{\alpha-1} + 1 - \delta].$$

Solving for K_{ss} produces:

$$K_{\text{ss}} = \left(\frac{\alpha}{\rho+\delta}\right)^{\frac{1}{1-\alpha}}.$$

Notice that the steady-steady state capital level is unaffected by the tax rate τ. Gross income at the steady state is Y_{ss} and is given by:

$$Y_{\text{ss}} = K_{\text{ss}}{}^{\alpha}.$$

Again, this is not a function of τ.

The government's period-by-period revenue function $\mathcal{T}_t(\tau)$ is now simply:

$$\mathcal{T}_t(\tau) = \tau(Y_{\text{ss}} - I_{\text{ss}}),$$

where I_{ss} is the steady-state investment level (which is tax-exempt). We can find I_{ss} by solving the law of motion for capital:

$$K_{t+1} = (1-\delta)K_t + I_t$$

for the steady-state level of capital:

$$K_{\text{ss}} = (1-\delta)K_{\text{ss}} + I_{\text{ss}}, \text{ so:}$$
$$I_{\text{ss}} = \delta K_{\text{ss}}.$$

Hence $\mathcal{T}_t$ becomes:

$$\mathcal{T}_t(\tau) = \tau(Y_{\text{ss}} - I_{\text{ss}}) = \tau\left(K_{\text{ss}}{}^{\alpha} - \delta K_{\text{ss}}\right) = \tau\left[\left(\frac{\alpha}{\rho+\delta}\right)^{\frac{\alpha}{1+\alpha}} - \delta\left(\frac{\alpha}{\rho+\delta}\right)^{\frac{1}{1+\alpha}}\right].$$

Thus government revenue each period is just a linear function of the tax rate τ, and there is no Laffer curve in this tax system.

13.4 Redistribution and Taxation

Now we turn our attention, as promised, to fiscal policies aimed at redistribution. We shall write down a model with two agents. One agent will be low-productivity and the other agent will be high-productivity. Without government intervention, there will be income inequality in this model. Why is the government interested in redistributing income? For now, let us simply take it as given that the fiscal authority will attempt to remedy income equality by taxes and transfers. We might expect the government to address the underlying causes of the agents' productivity gap, but since they are likely the result of, in a best case, different schooling histories, they are not likely to be remedied over the short-term.

Agents of type i, $i = a, b$ have a common utility function over consumption C^i and labor effort L^i of the (familiar) form:

$$\frac{C^{i\gamma}}{\gamma} - L^i,$$

where $\gamma < 1$. There is a technology transforming labor effort into the consumption good of the form:

$$Y^i = w^i L^i,$$

where $i = a, b$. Assume that $w^a > w^b$ so agents of type a are more productive than type b agents.

The government will tax type a agents at a rate τ in order to make a lump-sum transfer to type b agents of v. Hence agents face budget constraints of the form:

$$\begin{aligned} C^a &= (1-\tau)w^a L^a, \text{ and:} \\ C^b &= w^b L^b + v. \end{aligned}$$

Agents of type a face precisely the same problem that we solved in Section 13.2. Agents of type b face a "negative lump sum tax" of v. There is thus nothing unfamiliar about this problem.

The government has a budget constraint which requires it to balance transfer payments v with tax revenue $\mathcal{T}^a$ from agents of type a. Assume that there are an equal number of type a and type b agents, so the government budget constraint becomes:

$$v = \mathcal{T}^a = \tau w^a L^a(\tau).$$

From our analysis in Section 13.2 above, we know that $L^a(\tau)$ is:

$$L^a(\tau) = [(1-\tau)w^a]^{\frac{\gamma}{1-\gamma}}.$$

Thus:

$$\mathcal{T}^a(\tau) = \tau(1-\tau)^{\frac{\gamma}{1-\gamma}} w^{a\,\frac{1}{1-\gamma}}.$$

Agents of type b get a lump-sum subsidy of $v = \mathcal{T}^a$. Solving for their optimal labor supply gives:

$$L^b = w^{b\,\frac{1}{1-\gamma}} - \frac{v}{w^b}.$$

Thus agents of type b certainly decrease their effort as v increases.

We know that if $\gamma > 0$ that agents of type a will also decrease their effort as τ increases. Hence if $\gamma > 0$, redistribution will certainly lower both agent's labor supply and total national output.

If $\gamma < 0$, an increase in the tax on type a agents will increase their labor supply. This effect will, we shall see, never be large enough to overcome the decrease in type b labor effort. As a result, increases in redistribution will again lower national income. To see this, begin by calculating the effect on L^a and L^b of an increase in τ:

$$\begin{aligned}
\frac{\partial L^a}{\partial \tau} &= w^{a\,\frac{\gamma}{1-\gamma}} \left[\frac{\gamma}{1-\gamma}(1-\tau)^{\frac{\gamma}{1-\gamma}-1}(-1)\right] \\
&= -(1-\tau)^{\frac{\gamma}{1-\gamma}-1} w^{a\,\frac{\gamma}{1-\gamma}} \left[\frac{\gamma}{1-\gamma}\right] \\
&= -L^a \frac{\gamma}{1-\gamma}\frac{1}{1-\tau}. \qquad (13.9) \\
\frac{\partial L^a}{\partial \tau} &= -\frac{w^a}{w^b}\left[L^a - \tau\frac{\gamma}{1+\gamma}\frac{L^a}{1-\tau}\right] \\
&= -L^a \frac{w^a}{w^b}\left[1 - \frac{\gamma}{1-\gamma}\frac{\tau}{1-\tau}\right]. \qquad (13.10)
\end{aligned}$$

Armed with these derivatives, we can consider the effect on total national output (GDP) of an increase in τ:

$$\frac{dY}{d\tau} = \frac{dY^a}{d\tau} + \frac{dY^b}{d\tau} = \frac{d}{d\tau}w^a L^a + \frac{d}{d\tau}w^b L^b = -\frac{\gamma}{1-\gamma}\frac{1}{1-\tau}w^a L^a - w^a L^a\left[1 - \frac{\gamma}{1-\gamma}\frac{\tau}{1-\tau}\right].$$

We divide by $w^a L^a$, to find that $dY/d\tau < 0$ if and only if:

$$\begin{aligned}
-\frac{\gamma}{1-\gamma}\frac{1}{1-\tau} - \left[1 - \frac{\gamma}{1-\gamma}\frac{\tau}{1-\tau}\right] &< 0, \text{ or:} \\
-\frac{\gamma}{1-\gamma}\frac{1}{1-\tau} + \frac{\gamma}{1-\gamma}\frac{\tau}{1-\tau} - 1 &< 0, \text{ or:} \\
-\frac{\gamma}{1-\gamma}\frac{1}{1-\tau}(1-\tau) &< 1, \text{ or:} \\
-\frac{\gamma}{1-\gamma} &< 1.
\end{aligned}$$

Since we are assuming here that $\gamma < 0$, this is always true. In this model of labor supply, redistribution financed by distortionary taxes leads to a decrease in total national income.

Variable	Definition
$\mathcal{H}(a;\psi)$	Tax policy or system: legal mapping from actions of household a and parameters ψ to tax liability of household.
$\mathcal{T}(\psi)$	Realized revenue of the government under tax policy $\mathcal{H}$ with parameters ψ. The household is assumed to be using its best response, $a_{\max}(\psi)$.
a, A	Action of household a must lie in the set of possible actions A.
ψ	Vector of parameters in tax system or policy $\mathcal{H}$.
$a_{\max}(\psi)$	Household's optimal choice of action a under tax policy $\mathcal{H}$ with parameters ψ.
$\mathcal{Y}(a)$	Household's gross (pre-tax) income as a function of household action choice a.
$U\left[a, \mathcal{Y}(a) - \mathcal{H}(a;\psi)\right]$	Household's utility over action a and net (post-tax) income $\mathcal{Y} - \mathcal{H}$.
$V(\psi)$	Household's indirect utility with parameters ψ: $U\left(a_{\max}(\psi), \mathcal{Y}(a_{\max}(\psi) - \mathcal{H}\left[a_{\max}(\psi);\psi\right)\right]\right)$.
τ, E	Parameters of a flat tax system: the flat tax rate and the exemption.

Table 13.1: General Tax Notation for Chapter 13.

Exercises

Exercise 13.1 (Hard)
Consider an economy with infinitely many agents, each of whom is very very small. An agent i has preferences over consumption c^i and labor effort ℓ^i of:

$$u^i(c^i, \ell^i) = c^i - \gamma^i \ell^i.$$

The preference parameter γ^i is distributed uniformly on the interval $[0, 1]$. So the fraction of agents with preference parameters γ less than some number x is just x, for $0 \leq x \leq 1$. For example, exactly half of the population have values of γ less than or equal to 0.5.

Agents may only choose whether or not to work, not how many hours to work. If an agent chooses to work, she supplies exactly one unit of labor effort to the common backyard technology transforming labor effort into output as $y^i = \ell^i$. If an agent chooses not work, her labor effort is zero, she produces nothing and consumes nothing. All agents have the same backyard technology.

The government levies a flat income tax at a rate $0 \leq \tau \leq 1$. The tax rate τ is common to all

Variable	Definition
$u(C_t)$	One-period utility function when consumption is C_t
C_t, c_t	Aggregate consumption, household consumption.
I_t	Household investment in capital at time t.
K_t	Capital stock at time t.
L	Representative household's labor supply.
w	Wage rate on labor.
τ_L	Lump sum tax amount.
α, δ	Production parameters: the marginal product of capital and the depreciation rate.
γ	Preference parameter.
w^a, w^b	Wage rates of agents of type a and type b.
C^a, C^b	Consumption of agents of type a and type b.
L^a, L^b	Labor supply of agents of type a and type b.
v	Lump-sum transfer from type a agents to type b agents.

Table 13.2: Other Notation for Chapter 13

agents (that is, all agents face the same tax rate). Answer the following questions:

1. Given the tax rate τ, how much does an agent consume if she works? If she does not work?
2. For agent i, with preference parameter γ^i, calculate the utility of working (so that labor supply is $\ell^i = 1$) and of not working (so that labor supply is $\ell^i = 0$). What determines whether or not an agent works?
3. Given τ, find the largest value of γ such that agents prefer to work, or are at least indifferent between working and not working. Call this critical value $\gamma^*(\tau)$.
4. Given τ, multiply revenue per worker by the fraction of agents willing to work. Call this $\mathcal{T}(\tau)$, the government revenue function. Draw $\mathcal{T}(\tau)$ as a function of τ.
5. Is there a Laffer curve in the tax system?

Exercise 13.2 (Easy)
Suppose I ran the following regression:

$$G_t = b_0 + b_1\tau_t + u_t,$$

for $t = 1948, 1949, \ldots, 1997$, and where G_t are real government receipts and τ_t are some measure of the marginal tax rates faced by a typical American for the indicated years. Suppose that my estimated coefficient $\hat{b}_1$ is negative and statistically significant. Would you conclude from this that there is a Laffer curve in the U.S. economy, and that we are on its downward-sloping portion? Why or why not?

Exercise 13.3 (Moderate)
The representative household lives for one period and has preferences over consumption C and labor supply L as follows:

$$U(C, L) = 2\sqrt{C} - L.$$

The government levies a flat tax at rate τ and a lump-sum tax of S. Money spent on the lump-sum tax is exempt from the flat tax (that is, the lump-sum tax is paid with *pre-tax* dollars). The household can transform labor effort into the consumption good at a rate of one-to-one (the wage rate is unity). Answer the following questions:

1. What are the parameters of this tax system? What is the action chosen by the household?

2. Write down the tax system function $\mathcal{H}(a; \psi)$ in this case (replace a with the household's action and ψ with the parameters of the tax system).

3. If the household works some amount L, write down its tax bill.

4. Write down the household's consumption as a function of L, τ and S.

5. Solve the household's problem. What is $L(\tau, S)$?

6. Find the government's revenue function $\mathcal{T}(\psi)$.

Exercise 13.4 (Easy)
Assume that a household lives for one period and has preferences over consumption c and labor supply ℓ as follows:

$$4\sqrt{c} - \ell.$$

The household earns a constant wage of $w = 1$ for each unit worked. There is a flat tax of rate τ. Answer the following questions:

1. Given that the household works an amount ℓ, find the household's tax bill, $\mathcal{H}(\ell; \tau)$ and consumption, $c(\ell, \tau)$.

2. Find the household's optimal choice of labor effort, $\ell(\tau)$.

3. Find the government revenue function $\mathcal{T}(\tau)$.

4. The government wishes to raise an amount 3/4 in tax revenue from this household. Which tax rate(s) can the government use. Assume that the government is benevolent (that is, nice). Which tax rate does the government use?

Exercise 13.5 (Hard)
Tammy lives for two periods. In the first period of life she works ℓ hours at a wage of w per hour. In the second period of life she is retired and supplies no labor. There is a perfect bond market on which Tammy saves an amount b in the first period of life, which earns a net interest rate of $r = 0$ (for a gross interest rate $1 + r = 1$), so Tammy's savings pay off $(1 + r)b = b$ in the second period of life. Tammy has preferences over consumption sequences $\{c_1, c_2\}$ and effort ℓ as follows:

$$u(\{c_1, c_2\}, \ell) = \sqrt{c_1} + \sqrt{c_2} - \ell.$$

The government taxes Tammy at a rate τ_1 on income earned in the first period of life, $w\ell$. Tammy owes no tax in the second period of life.

1. Write down Tammy's budget constraint in each period of her life, substitute out the savings term b and show that Tammy's present-value budget constraint is:

$$c_1 + \frac{1}{1+r}c_2 = c_1 + c_2 \leq (1 - \tau_1)w\ell.$$

2. Calculate Tammy's optimal choices of consumption, work, and savings: c_1, c_2, ℓ, and b, as a function of the tax rate τ_1.

3. How much revenue does the government raise as a function of the tax rate τ_1, $\mathcal{T}(\tau_1)$. Find τ_1^*, the tax rate at which government revenue is maximized. What is the maximum amount of revenue the government can raise (that is, what is $\mathcal{T}_1(\tau_1^*)$)?

4. Is there a Laffer curve?

Exercise 13.6 (Hard)
Use the same preferences and technology as in Exercise (13.5) above, except that here we call the government tax rate τ_2 instead of τ_1. The government will allow Tammy to *exempt* her savings b from taxes (as in a 401(k) plan), so she owes tax at a rate τ_2 in the first period only on the portion of her income that she does not save, $w\ell - b$. Tammy still owes no tax in the second period of life. If Tammy saves an amount b she consumes $(1 - \tau_2)(w\ell - b)$ in period 1 and b in period 2.

1. Write down Tammy's present-value budget constraint. How is it different from the one you calculated in part (1) of Exercise (13.5) above?

2. Solve Tammy's optimization problem. What are her optimal choices of consumption, work, and savings c_1, c_2, ℓ, and b as a function of the tax rate τ_2?

3. Now how much revenue does the government raise as a function of the tax rate τ_2 (call the revenue function $\mathcal{T}_2(\tau_2)$)? What is the maximum amount of revenue the government can raise?

4. Is there a Laffer curve?

5. How do your answers differ from those in Exercise (13.5)? Why?

Chapter 14

The Optimal Path of Government Debt

Up to this point we have assumed that the government must pay for all its spending each period. In reality, governments issue debt so as to spread their costs across several periods, just like households do. The path of governmental debt over time very often corresponds to events of major historical import, such as wars. For example, England has paid for her wars by issuing debt, resulting in debt peaks during the Seven Years War, the Napoleonic Wars, and especially World War I. Indeed, some economists argue that the sophistication of England's capital markets contributed to her eventual successes in the wars of the 17th and 18th centuries.

The U.S. and several European countries have run persistent peacetime deficits since about 1979. Quite a bit has been written in the popular press about the dire consequences of the ever-mounting debt. In this chapter we will not consider that a large debt may be inherently bad, instead we will treat the debt as a tool for use by a benevolent government. In the previous two chapters we saw that government spending may crowd out consumption and investment, and that government taxes may decrease labor supply and capital accumulation, but in this chapter we will have nothing bad to say about debt. In Chapter 18, however, we argue that under certain circumstances, large and persistent government deficits may be inflationary. In this chapter we will continue to ignore the price level and the ability of the government to raise revenue by printing money, so we will not be able to consider inflation directly.

We will begin by considering the government budget deficit and defining some terms. The reader should be thoroughly familiar with the terms defined there, as well as the historical paths of the debt, deficits, debt to GDP ratios and so on.

Next we will consider a very simple theory of the debt, that of *Barro-Ricardo Equivalence.*

Barro-Ricardo Equivalence is named for Robert Barro, at Harvard, and David Ricardo, the 19th-century economist. In this theory, the timing of government taxes and spending (and hence the path of the debt) do not matter. Only the present discounted value of these objects is important. We shall see that Barro-Ricardo Equivalence requires some strong assumptions. As we relax those assumptions, the timing of taxes begins to matter.

Requiring the government to use distortionary taxes is one way of breaking Barro-Ricardo Equivalence. In the final two sections of this chapter, we construct a fairly sophisticated theory of government debt based on precisely this assumption. This is known as the *Ramsey Optimal Tax problem* (or simply the Ramsey problem, for short). In the Ramsey problem the government has access only to a distortionary tax (in this case, an excise tax), and must raise a specific amount of revenue in the least-distortionary manner. In this example, the government will have to finance a war, modeled as a spike in planned government expenditures, with a period-by-period excise tax. By finding the optimal path of tax revenues, we can find the optimal path of government deficits and surpluses. This will provide us with a theory of government debt and deficits.

One of the features of the Ramsey model will be that both the government and the household will have access to a perfect loan market at a constant interest rate. This interest rate will not vary with the amount actually borrowed or lent, nor will it vary across time for other reasons. This is sometimes known as the "small open economy" equilibrium, but truthfully we are simply abstracting from the question of equilibrium entirely. No markets will clear in this example.

14.1 The Government Budget Constraint

Let $\mathcal{T}_t$ be the real revenue raised by the government in period t, let G_t be real government spending in period t (including all transfer payments) and let B_t^g be the real outstanding stock of government debt at the end of period t. That is, $B_t^g > 0$ means that the government is a net borrower in period t, while $B_t^g < 0$ means that the government is a net lender in period t. There is a real interest rate of r_t that the government must pay on its debt.

Assuming that the government does not alter the money supply, the government's budget constraint becomes:

(14.1) $$G_t + r_{t-1}B_{t-1}^g = \mathcal{T}_t + (B_t^g - B_{t-1}^g).$$

The left hand side gives expenditures of the government in period t. Notice that the government not only has to pay for its direct expenditures in period t, G_t, it must also *service the debt* by paying the interest charges $r_{t-1}B_{t-1}^g$. Of course, if the government is a net lender, then B_t^g is negative and it is collecting revenue from its holdings of other agents' debt.

The right hand side of the government budget constraint gives revenues in period t. The government raises revenue directly from the household sector by collecting taxes $\mathcal{T}_t$. In addition, it can raise revenue by *issuing net new debt* in the amount $B_t^g - B_{t-1}^g$.

Government debt is a *stock* while government deficits are a *flow*. Think of the debt as water in a bathtub: tax revenue is the water flowing out of drainhole and spending is water running in from the tap. In addition, if left to itself, the water grows (reflecting the interest rate). Each period, the level of water in the tub goes up or down (depending on $G_t, \mathcal{T}_t$ and r_t) by the amount $B_t^g - B_{t-1}^g$.

Call the *core deficit* the difference between real government purchases G_t and real government tax revenue $\mathcal{T}_t$. In the same way, define the *reported deficit* (or simply the deficit) to be the difference between all government spending, $G_t + r_{t-1}B_{t-1}^g$ and revenues from taxes $\mathcal{T}_t$. Thus:

$$\begin{aligned}
(\text{core deficit})_t &= G_t - \mathcal{T}_t, \text{ and:}\\
(\text{reported deficit})_t &= G_t - \mathcal{T}_t + r_{t-1}B_{t-1}^g.
\end{aligned}$$

The reported deficit is what is reported in the media each year as the government wrangles over the deficit. The U.S. has been running a core surplus since about 1990.

We can convert the period-by-period budget constraint in equation (14.1) into a single, infinite-horizon, budget constraint. For the rest of this chapter we will assume that the real interest rate is constant, so that $r_t = r$ all periods t. Assume further (again, purely for simplicity) that the government does not start with a stock of debt or with any net wealth, so $B_{-1}^g = 0$. Thus for convenience rewrite equation (14.1) as:

$$G_t + (1+r)B_{t-1}^g = \mathcal{T}_t + B_t^g.$$

The government's period-by-period budget constraints, starting with period zero, will therefore evolve as:

$$\begin{aligned}
(t=0) \qquad & G_0 + (1+r)\cdot 0 = \mathcal{T}_0 + B_0^g, \text{ so:}\\
& B_0^g = G_0 - \mathcal{T}_0.\\
(t=1) \qquad & G_1 + (1+r)B_0^g = \mathcal{T}_1 + B_1^g, \text{ so:}\\
& B_0^g = \frac{1}{1+r}(\mathcal{T}_1 - G_1) + \frac{1}{1+r}B_1^g.\\
(t=2) \qquad & G_2 + (1+r)B_1^g = \mathcal{T}_2 + B_2^g, \text{ so:}\\
& B_1^g = \frac{1}{1+r}(\mathcal{T}_2 - G_2) + \frac{1}{1+r}B_2^g.\\
(t=3) \qquad & G_3 + (1+r)B_2^g = \mathcal{T}_3 + B_3^g, \text{ so:}\\
& B_2^g = \frac{1}{1+r}(\mathcal{T}_3 - G_3) + \frac{1}{1+r}B_3^g.
\end{aligned}$$

Now recursively substitute backwards for B_t^g in each equation. That is, for the $t = 2$ budget constraint, substitute out the B_2^g term from the $t = 3$ budget constraint to form:

$$G_2 + (1+r)B_1^g = \mathcal{T}_2 + \frac{1}{1+r}(\mathcal{T}_3 - G_3) + \frac{1}{1+r}B_3^g, \text{ so:}$$
$$B_1^g = \frac{1}{1+r}(\mathcal{T}_2 - G_2) + \left(\frac{1}{1+r}\right)^2 (\mathcal{T}_3 - G_3) + \left(\frac{1}{1+r}\right)^2 B_3^g.$$

Eventually, this boils down to:

$$B_0^g = \frac{1}{1+r}(\mathcal{T}_1 - G_1) + \left(\frac{1}{1+r}\right)^2 (\mathcal{T}_2 - G_2) + \left(\frac{1}{1+r}\right)^3 (\mathcal{T}_3 - G_3) + \left(\frac{1}{1+r}\right)^3 B_3^g.$$

Since we also know that $B_0^g = G_0 - \mathcal{T}_0$ we can rewrite this as:

$$G_0 - \mathcal{T}_0 = \frac{1}{1+r}(\mathcal{T}_1 - G_1) + \left(\frac{1}{1+r}\right)^2 (\mathcal{T}_2 - G_2) + \left(\frac{1}{1+r}\right)^3 (\mathcal{T}_3 - G_3) + \left(\frac{1}{1+r}\right)^3 B_3^g.$$

Collect all of the G_t terms on the left hand side and all of the $\mathcal{T}_t$ terms on the right hand side to produce:

$$G_0 + \frac{1}{1+r}G_1 + \left(\frac{1}{1+r}\right)^2 G_2 + \left(\frac{1}{1+r}\right)^3 G_3 =$$
$$\mathcal{T}_0 + \frac{1}{1+r}\mathcal{T}_1 + \left(\frac{1}{1+r}\right)^2 \mathcal{T}_2 + \left(\frac{1}{1+r}\right)^3 \mathcal{T}_3 + \left(\frac{1}{1+r}\right)^3 B_3^g.$$

In the same way, we can start solving backwards from any period $j \geq 0$ to write the government's budget constraint as:

$$\sum_{t=0}^{j} \left(\frac{1}{1+r}\right)^t G_t = \sum_{t=0}^{j} \left(\frac{1}{1+r}\right)^t \mathcal{T}_t + \left(\frac{1}{1+r}\right)^j B_j^g.$$

If we assume that:

$$\lim_{j\to\infty} \left(\frac{1}{1+r}\right)^j B_j^g = 0, \tag{14.2}$$

then we can continue to recursively substitute indefinitely (that is, we can let j grow arbitrarily large), to produce the single budget constraint:

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t G_t = \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t \mathcal{T}_t. \tag{14.3}$$

Notice that the left hand side is the present discounted value of government expenditures while the right hand side is the present discounted value of government tax revenues. The government debt terms, B_t^g, have disappeared, since, at the limit *all government borrowing* *must be repaid.* The condition in equation (14.2), sometimes known as a *transversality condition,* prohibits the government from always borrowing to pay its debt. At some point in the future, all government expenditures must be backed by government tax revenues.

In future sections, we will mainly work with constraints of the form in equation (14.3) to find optimal sequences of tax revenue $\mathcal{T}_t$ and then infer what the sequence of government debt must be.

14.2 Barro-Ricardo Equivalence

Barro-Ricardo Equivalence is the statement that the timing of government taxes do not matter, since households internalize the government budget constraint and save to pay the expected future taxes. This is an old idea, first formulated by David Ricardo in the 19th century, that has returned to prominence with the 1974 paper "Are Government Bonds Net Wealth?" by Robert Barro.[1] In that paper, Barro argued that debt-financed tax cuts could not affect output, since households would use the increased net income to save for the coming increased taxes. This argument was of particular interest during the early 1980s when debt-financed tax cuts were a centerpiece of the government's economic strategy. In this section we will examine the proposition in a simple two-period model and then again in an infinite horizon model.

Assumptions for Barro-Ricardo Equivalence

Since the time path of government debt is determined entirely by the difference between spending and taxes, Barro-Ricardo equivalence says that the optimal path of government debt is indeterminate: only the present discounted values of spending and taxes matter. Barro-Ricardo equivalence rests on three key assumptions, and we will have to break at least one of them to get a determinate theory of optimal government debt. Barro-Ricardo equivalence holds if:

1. There is a perfect capital market, on which the government and households can borrow and lend as much as desired without affecting the (constant) real interest rate.
2. Households either live forever or are altruistic towards their offspring.
3. The government can use lump-sum taxes.

Since Barro-Ricardo Equivalence requires the government and households to completely smooth out transitory spikes in spending or taxes, it is obvious why a perfect capital market is important. If households were not altruistic towards their offspring, and did not

[1] Barro, Robert J. (1974) "Are Government Bonds Net Wealth?" *Journal of Political Economy* 82(6) pp1095-1117.

live forever, then they would consume from a debt-financed tax cut without saving and bequeathing enough to their offspring to repay the debt. Finally, if the government cannot use lump-sum taxes, then large taxes cause large distortions, encouraging the government to use low taxes to spread the deadweight loss out over several periods. In the next section, we force the government to use distortionary taxes, which breaks Barro-Ricardo Equivalence.

A Two-Period Example

Consider a government which must make real expenditures of $\{G_0, G_1\}$. It levies lump-sum taxes each period of $\{\mathcal{T}_0, \mathcal{T}_1\}$. The household has a fixed endowment stream of $\{Y_0, Y_1\}$. Both the government and the household have access to a perfect bond market, and can borrow and lend any amount at the constant real interest rate r. The government's initial stock of debt, $B^g_{-1} = 0$, and the government must repay all that it borrows by the end of period $t = 1$.

The household has preferences over consumption streams $\{C_0, C_1\}$ given by:

$$U(C_0, C_1) = u(C_0) + \beta u(C_1),$$

where $0 < \beta < 1$. We assume that $u' > 0, u'' < 0$. The government's two-period (flow) budget constraints are:

$$(t = 0) \qquad G_0 = \mathcal{T}_0 + B^g_0, \text{ and:}$$

$$(t = 1) \qquad G_1 + (1 + r)B^g_0 = \mathcal{T}_1.$$

These can be collapsed (by substituting out the debt term B^g_0) into a single budget constraint, expressed in terms of present discounted value:

$$G_0 + \frac{1}{1+r}G_1 = \mathcal{T}_0 + \frac{1}{1+r}\mathcal{T}_1. \tag{14.4}$$

This is the form of the government's budget constraint with which we will work. The household's two-period (flow) budget constraints are:

$$(t = 0) \qquad C_0 + \mathcal{T}_0 + B_0 = Y_0, \text{ and:}$$

$$(t = 1) \qquad C_1 + \mathcal{T}_1 = Y_1 + (1 + r)B_0.$$

Here we are using Barro's notation that, for private individuals, B_t denotes the stock of *savings* at the end of period t. If $B_t > 0$ then the household is a net lender. Collapsing the two one-period budget constraints into a single present-value budget constraint produces:

$$C_0 + \frac{1}{1+r}C_1 = (Y_0 - \mathcal{T}_0) + \frac{1}{1+r}(Y_1 - \mathcal{T}_1). \tag{14.5}$$

Notice that the government's lump sum taxes, $\mathcal{T}_t$, which form revenue for the government, are a cost to the household.

We will now use equation (14.4) to rewrite equation (14.5) without the tax terms. Notice that equation (14.5) may be written as:

$$C_0 + \frac{1}{1+r}C_1 = Y_0 + \frac{1}{1+r}Y_1 - \left(\mathcal{T}_0 + \frac{1}{1+r}\mathcal{T}_1\right).$$

But from the government's present-value budget constraint equation (14.4) we know that:

$$\mathcal{T}_0 + \frac{1}{1+r}\mathcal{T}_1 = G_0 + \frac{1}{1+r}G_1.$$

Thus we can rewrite the household's present-value budget constraint as:

$$C_0 + \frac{1}{1+r}C_1 = Y_0 + \frac{1}{1+r}Y_1 - \left(G_0 + \frac{1}{1+r}G_1\right).$$

Notice that the household's budget constraint no longer contains tax terms $\mathcal{T}_t$. Instead, the household has internalized the government's present-value budget constraint, and uses the perfect bond market to work around any fluctuations in net income caused by sudden increases or decreases in taxes.

An Infinite-Horizon Example

The infinite horizon version is a very simple extension of the previous model. Now governments will have a known, fixed, sequence of real expenditures $\{G_t\}_{t=0}^{\infty}$ that they will have to finance with some sequence of lump-sum taxes $\{\mathcal{T}_t\}_{t=0}^{\infty}$. The household has some known endowment sequence $\{Y_t\}_{t=0}^{\infty}$. Both the household and the government can borrow and lend freely on a perfect bond market at the constant interest rate r.

The household lives forever and has preferences over sequences of consumption $\{C_t\}_{t=0}^{\infty}$ of:

$$U(\{C_t\}_{t=0}^{\infty}) = \sum_{t=0}^{\infty} \beta^t u(C_t),$$

where $0 < \beta < 1$. Here we again assume that $u' > 0, u'' < 0$. To make the notation in this section simpler, define:

$$\mathbf{G} = \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t G_t \qquad \mathbf{T} = \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t \mathcal{T}_t$$

$$\mathbf{Y} = \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t Y_t \qquad \mathbf{C} = \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t C_t.$$

That is, $\mathbf{G}$ is the present discounted value of government spending, $\mathbf{T}$ is the present discounted value of government revenue, $\mathbf{Y}$ is the present discounted value of the household's endowment stream and $\mathbf{C}$ the present discounted value of the household's consumption stream.

The government's present-value budget constraint, as in equation (14.3), may now be written:

$$\mathbf{G} = \mathbf{T}.$$

The household's present-value budget constraint, in the same way, may be written:

$$\mathbf{C} = \mathbf{Y} - \mathbf{T}.$$

But since the government budget constraint requires $\mathbf{T} = \mathbf{G}$, the household's budget constraint becomes:

$$\mathbf{C} = \mathbf{Y} - \mathbf{G}.$$

Once again, the timing of taxes ceases to matter. The household only cares about the present discounted value of government spending.

As a final step, we shall solve the household's problem. For simplicity, assume that $1 + r = \beta^{-1}$. The household's Lagrangian is:

$$\mathcal{L}(\{C_t\}_{t=0}^{\infty}, \lambda) = \sum_{t=0}^{\infty} \beta^t u(C_t) + \lambda(\mathbf{Y} - \mathbf{G} - \mathbf{C}).$$

To find the optimal choices of consumption given the constraint, we maximize the Lagrangian with respect to consumption. The first-order necessary conditions for maximization are formed by taking the derivative with respect to consumption in some typical period j, and from the constraint. Recall that:

$$\frac{\partial \mathbf{C}}{\partial C_j} = \left(\frac{1}{1+r}\right)^j .$$

So the first-order conditions are:

$$\beta^j u'(C_j) - \lambda \left(\frac{1}{1+r}\right)^j = 0, \text{ for all } j = 0, \dots, \infty, \text{ and:}$$
$$\mathbf{C} = \mathbf{Y} - \mathbf{G}.$$

With the assumption that $1 + r = \beta^{-1}$, we find that:

$$u'(C_j) = \lambda, \text{ for all } j = 0, 1, \dots, \infty.$$

But λ is constant, so $u'(C_j)$ must also be constant. We conclude that consumption is also constant, $C_j = C^\star$ in all periods j. If consumption is constant at $C^\star$, we can substitute back into the budget constraint to find $C^\star$, using the fact that $C_t = C^\star$ for all t:

$$\begin{aligned}
\mathbf{Y} - \mathbf{G} = \mathbf{C} &= \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t C_t \\
&= \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t C^\star \\
&= C^\star \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t \\
&= C^\star \frac{1}{1 - \frac{1}{1+r}} \\
&= C^\star \frac{1+r}{r}, \text{ so:} \\
C^\star &= \frac{r}{1+r}(\mathbf{Y} - \mathbf{G}).
\end{aligned}$$

Thus household consumption is constant over time, increasing in $\mathbf{Y}$ and decreasing in $\mathbf{G}$. Household consumption is utterly unaffected by the timing of the taxes $\mathcal{T}_t$ used to finance the government's spending.

14.3 Preliminaries for the Ramsey Problem

Before we lay out the Ramsey model, we are going to need to define some terms. In particular, readers may be unfamiliar with *excise taxes*, which are used extensively in this chapter. Also, we will define in general terms the structure of Ramsey problems. Finally, we will define *indirect utility*, an important concept with which the reader may be unfamiliar.

Excise Taxes

An *excise tax* is a constant tariff levied on each unit of a good consumed. An example would be a \$1/gallon gasoline tax, or a \$0.25/pack cigarette tax. These are not sales taxes. Sales taxes are levied as a percentage of the total value of the goods purchased. Excise taxes are unaffected by the price of the taxed good. If there were a vector of n goods $\{x_i\}_{i=1}^n$, with an associated vector of prices $\{p_i\}_{i=1}^n$, and a consumer had m total dollars to spend on these goods, her budget constraint would be:

$$\sum_{i=1}^{n} p_i x_i \leq m.$$

Now the government levies an excise tax of τ_i on each good $i = 1 \dots n$. The consumer's budget constraint becomes:

$$\sum_{i=1}^{n} (p_i + \tau_i)\, x_i \leq m,$$

where the price paid by the consumer is now $p_i + \tau_i$. What would a sales tax look like?[2] Think of excise taxes like this: for each good x_i the consumer buys, she pays p_i to the firm, and τ_i to the government.

Under an excise tax system the government's revenue $\mathcal{H}(\cdot,\cdot)$ from the tax system, without taking into consideration the household's reactions (see Chapter 13), is:

$$\mathcal{H}(x_1, \dots, x_n; \tau_1 \dots, \tau_n) = \sum_{i=1}^{n} \tau_i x_i.$$

Households will adjust their choices of consumption x_i, $i = 1 \dots n$ in response to the taxes (this plays the role of $a_{\max}$ from Chapter 13). Thus, taking into consideration the household's best response, the government raises:

$$\mathcal{T}(\tau_1 \dots \tau_n) = \mathcal{H}(x_1^\star, \dots, x_n^\star; \tau_1, \dots, \tau_n) = \sum_{i=1}^{n} \tau_i x_i^\star(p_1, \dots, p_n; m; \tau_1, \dots, \tau_n).$$

Here $x_i^\star(\cdot;\cdot;\cdot)$ is the household's *Marshallian demand* for good i. As you recall from intermediate microeconomic theory, the Marshallian demand by an agent for a product gives the quantity of the product the agent would buy given prices and her income.

Structure of the Ramsey Problem

The government announces a sequence of excise tax rates $\{\tau_t\}_{t=0}^{\infty}$, which households take as given in making their decisions about consumption, borrowing and saving. This is actually quite a strong assumption, when you stop to think about it. The government has committed to a sequence of actions, when deviation might help it raise more revenue. What mechanism does a sovereign government have to enforce its commitment? Policies change, heads of state topple and constitutions are rewritten every year. Quite a bit of extremely interesting research centers on how governments ought to behave when they cannot credibly commit to a policy and all the agents in the model know it. See Chapter 19

[2] Okay, I'll tell you. Let's say the government levies a sales tax of t_i on each good i. Then the agent's budget constraint becomes:

$$\sum_{i=1}^{n} (1 + t_i)\, p_i x_i \leq m,$$

where now the consumer owes $t_i p_i x_i$ on each good purchased.

for a discussion of commitment in the context of a Ramsey problem in monetary policy. In that chapter we introduce the game-theoretic concepts required to model the strategic interactions of the private sector and the government.

So our benevolent government will take the purchasing behavior of its citizens (in the form a representative household) in response to its announced set of taxes $\{\tau_t\}_{t=0}^{\infty}$ as given. It will seek to raise some exogenous, known, amount of money sufficient, in present-value terms, to pay for the stream of real government expenditures on goods and services, $\{G_t\}_{t=0}^{\infty}$. These expenditures will not affect the representative household's utility or output in a meaningful way: they will be used to fight a war, or, more succinctly, thrown into the ocean. Many sequences of taxes will pay for the government's stream of purchases. Our government will choose among them by finding the tax sequence that maximizes the representative household's indirect utility.

Indirect Utility

The technical definition of indirect utility is the utility function with the choice variables replaced by their optimal values. Consider for example the following two-good problem. The utility function is:

$$U(c_1, c_2) = \ln(c_1) + \gamma \ln(c_2), \tag{14.6}$$

where $\gamma > 0$, and the budget constraint is:

$$p_1 c_1 + p_2 c_2 \leq m.$$

The Lagrangian is:

$$\mathcal{L}(c_1, c_2, \lambda) = \ln(c_1) + \gamma \ln(c_2) + \lambda(m - p_1 c_1 - p_2 c_2).$$

The first-order conditions are thus:

$$\begin{aligned} \frac{1}{c_1} - \lambda p_1 &= 0, \\ \frac{\gamma}{c_2} - \lambda p_2 &= 0, \text{ and:} \\ p_1 c_1 + p_2 c_2 &= m. \end{aligned}$$

Combined with the budget constraint, these imply that:

$$c_2 = \gamma \frac{p_1}{p_2} c_1, \text{ so:}$$

$$c_1 = \frac{1}{1+\gamma} \frac{m}{p_1}, \text{ and:} \tag{14.7}$$

$$c_2 = \frac{\gamma}{1+\gamma} \frac{m}{p_2}. \tag{14.8}$$

To find the indirect utility function, we substitute the optimal policies in equations (14.7) and (14.8) into the utility function in equation (14.6). Call the indirect utility function $V(p_1, p_2, m)$. It is how much utility the household can achieve at prices p_1, p_2 and at income m when it is optimizing. Thus, in this case:

$$\begin{aligned} V(p_1, p_2, m) &= \ln\left(\frac{1}{1+\gamma}\frac{m}{p_1}\right) + \gamma \ln\left(\frac{\gamma}{1+\gamma}\frac{m}{p_2}\right) \\ &= (1+\gamma)\ln(m) - \ln(p_1) - \ln(p_2) - (1+\gamma)\ln(1+\gamma) + \gamma\ln(\gamma). \end{aligned} \tag{14.9}$$

So we can see immediately the effect on maximized utility of an increase in wealth m, or of an increase in the prices p_1 and p_2. As we expect, optimized utility is increasing in wealth and decreasing in the prices.

Annuities

In this chapter we will often characterize income streams in terms of an *annuity*. An annuity is one of the oldest financial instruments, and also one of the simplest. In essence an annuity is a constant payment each period in perpetuity. Thus if one has an annuity of \$100, one can be assured of a payment of \$100 each year for the rest of one's life, and one may also assign it to one's heirs after death.

Risk averse agents with a known but fluctuating income stream of $\{y_t\}_{t=0}^{\infty}$ may, depending on the interest rate and their discount factor, want to convert it to an annuity, paying a constant amount a each period, of the same present discounted value. Given a constant net interest rate of r, it is easy to determine what a must be. We call a the *annuity value* of the income stream $\{y_t\}_{t=0}^{\infty}$.

We begin by calculating the present discounted value $\mathbf{Y}$ of the income stream $\{y_t\}_{t=0}^{\infty}$:

$$\mathbf{Y} \equiv \sum_{t=0}^{\infty}\left(\frac{1}{1+r}\right)^t y_t.$$

We know that the present discounted value of an annuity of a is just:

$$\sum_{t=0}^{\infty}\left(\frac{1}{1+r}\right)^t a.$$

For the present discounted value of the income stream and the annuity to be equal, a must satisfy:

$$\mathbf{Y} = \sum_{t=0}^{\infty}\left(\frac{1}{1+r}\right)^t a = \frac{a}{1-1/(1+r)} = \frac{(a)(1+r)}{r}, \text{ so:}$$

$$a = \frac{r}{1+r}\mathbf{Y}.$$

A reasonable value of r is around 0.05, which means that $r/(1+r)$ is 1/21.

14.4 The Ramsey Optimal Tax Problem

The Household's Problem

Consider a household with a known endowment stream $\{y_t\}_{t=0}^{\infty}$. This household orders infinite sequences of consumption $\{c_t\}_{t=0}^{\infty}$ as:

$$U(\{c_t\}_{t=0}^{\infty}) = \sum_{t=0}^{\infty} \beta^t \ln(c_t), \tag{14.10}$$

where $0 < \beta < 1$. To get a nice closed form for consumption we are going to assume logarithmic preferences.

There is a perfect bond market on which the household may borrow and lend at a constant real interest rate r which we assume satisfies $1 + r = \beta^{-1}$.

The household faces a known sequence of *excise taxes* $\{\tau_t\}_{t=0}^{\infty}$ levied by the government (see the above discussion for a review of excise taxes). Since there is only one consumption good in each period, we can safely take the within-period price of the consumption good to be unity. Thus in some period t, if the household consumes c_t, expenditures must be $c_t + \tau_t c_t$ or more simply $(1 + \tau_t)c_t$. Hence the present discounted value of expenditures including the tax bill is:

$$\text{PDV}_{\text{expenditures}} = \sum_{t=0}^{\infty}(1+r)^{-t}c_t + \sum_{t=0}^{\infty}(1+r)^{-t}\tau_t c_t = \sum_{t=0}^{\infty}(1+r)^{-t}(1+\tau_t)c_t.$$

The household's present discounted value of expenditures must equal the present discounted value of the endowment stream. Hence its budget constraint is:

$$\sum_{t=0}^{\infty}(1+r)^{-t}(1+\tau_t)c_t \leq \sum_{t=0}^{\infty}(1+r)^{-t}y_t \equiv \mathbf{Y}. \tag{14.11}$$

Here I have defined the term $\mathbf{Y}$ to be the present discounted value of the income sequence $\{y_t\}_{t=0}^{\infty}$. This is merely for convenience.

Hence the household's Lagrangian is:

$$\mathcal{L}(\{c_t\}_{t=0}^{\infty}, \lambda) = \sum_{t=0}^{\infty} \beta^t \ln(c_t) + \lambda\left(\mathbf{Y} - \sum_{t=0}^{\infty}(1+r)^{-t}(1+\tau_t)c_t\right). \tag{14.12}$$

The first-order condition of equation (14.12) with respect to consumption in the arbitrary period j is:

$$\frac{\beta^j}{c_j} - \lambda(1+r)^{-j}(1+\tau_j) = 0, \text{ for all } j = 0, \ldots, \infty.$$

With the assumption that $1/(1+r) = \beta$, and with a certain amount of manipulation, we can write this as:

$$c_j(1+\tau_j) = 1/\lambda, \text{ for all } j = 0 \ldots \infty. \tag{14.13}$$

Notice that this last equation implies expenditures will be constant across all periods. In periods with relatively higher excise tax rates, consumption will decrease exactly enough to keep the dollar outlays exactly the same as in every other period. This is an artifact of log preferences and not a general property of this problem. However, it does simplify our job enormously.

The next step will be to substitute the optimal consumption plan in equation (14.13) into the budget constraint in equation (14.11) to determine how much the household spends each period. Substituting, we find:

$$\sum_{t=0}^{\infty}(1+r)^{-t}\frac{1}{\lambda} = \mathbf{Y}.$$

Taking out the $1/\lambda$ term (because it does not vary with t), we find that:

$$\frac{1}{\lambda} = \frac{\mathbf{Y}}{\sum_{t=0}^{\infty}(1+r)^{-t}} \equiv \mathbf{W}. \tag{14.14}$$

In other words, $1/\lambda$ is equal to the *annuity value* of the endowment stream (which I have named $\mathbf{W}$ for convenience). Of course from equation (14.13) we conclude that:

$$c_t^{\star} = \frac{\mathbf{W}}{1+\tau_t}, \text{ for all } t = 0, 1, \ldots, \infty. \tag{14.15}$$

Here I denote the optimal consumption decision as $c_t^{\star}$.

The Household's Indirect Utility

We are now ready to calculate the household's indirect utility function. Substituting the optimal policy in equation (14.15) into the preferences in equation (14.10) produces:

$$\begin{aligned}
V(\{\tau_t\}_{t=0}^{\infty}, \mathbf{W}) &= \sum_{t=0}^{\infty} \beta^t \ln(c_t^{\star}) \\
&= \sum_{t=0}^{\infty} \beta^t \ln\left(\frac{\mathbf{W}}{1+\tau_t}\right) \\
&= \sum_{t=0}^{\infty} \beta^t \left[\ln(\mathbf{W}) - \ln(1+\tau_t)\right] \\
&= -\sum_{t=0}^{\infty} \beta^t \ln(1+\tau_t) + \frac{\ln(\mathbf{W})}{1-\beta}. \qquad (14.16)
\end{aligned}$$

Here we are using some more convenient properties of the log function to simplify our result. Notice that $V(\cdot,\cdot)$ is decreasing in tax rates τ_t and increasing in the annuity value of wealth, $\mathbf{W}$.

The Government's Problem

The government faces a known unchangeable stream of period real expenditures $\{G_t\}_{t=0}^{\infty}$ and can borrow and lend at the same rate $1 + r = \beta^{-1}$ as the household. Define:

$$\mathbf{G} \equiv \sum_{t=0}^{\infty} (1+r)^{-t} G_t,$$

i.e., let $\mathbf{G}$ denote the PDV of government expenditures. These government expenditures do not affect the household's utility or output in any meaningful way.

The government realizes revenue only from the excise tax it levies on the household. Hence each period the tax system produces revenues of:

$$\mathcal{H}_t(c_t; \tau_t) = \tau_t c_t.$$

But of course consumption is itself a function of taxes. The government takes as given the household's decisions. From equation (14.15) we know we can rewrite this as:

$$\mathcal{T}_t(\tau_t) = \mathcal{H}_t\left[c_t^\star(\tau_t); \tau_t\right] = \tau_t c_t^\star(\tau_t) = \mathbf{W}\frac{\tau_t}{1+\tau_t}.$$

The government's present-value budget constraint is:

$$\sum_{t=0}^{\infty} (1+r)^{-t} \mathcal{T}_t \leq \mathbf{G}, \text{ or:}$$

$$\sum_{t=0}^{\infty} (1+r)^{-t} \mathbf{W}\frac{\tau_t}{1+\tau_t} \leq \mathbf{G},$$

which we will find it convenient to rewrite as:

$$\sum_{t=0}^{\infty} (1+r)^{-t} \frac{\tau_t}{1+\tau_t} \leq \frac{\mathbf{G}}{\mathbf{W}}. \tag{14.17}$$

We divide by $\mathbf{W}$ merely to keep the algebra clean later.

The government maximizes the representative household's indirect utility subject to its present-value budget constraint by choice of sequences of excise taxes. Hence the government's Lagrangian is:

$$\mathcal{L}(\{\tau_t\}_{t=0}^{\infty}, \mu) = -\sum_{t=0}^{\infty} \beta^t \ln(1+\tau_t) + \frac{\ln(\mathbf{W})}{1-\beta} + \mu\left(\sum_{t=0}^{\infty} (1+r)^{-t} \frac{\tau_t}{1+\tau_t} - \frac{\mathbf{G}}{\mathbf{W}}\right). \tag{14.18}$$

The government is choosing the sequence of tax rates $\{\tau_t\}_{t=0}^{\infty}$ which makes the household as happy as possible given that the government has to raise enough tax revenue to finance the war. Here μ is the multiplier on the government's budget constraint, the same way λ was the multiplier on the household's budget constraint previously. In some typical period j, where $j = 0, \dots, \infty$, the first-order condition with respect to the tax rate is as follows:

$$-\frac{\beta^j}{1+\tau_j} + \mu(1+r)^{-j}\left(\frac{1}{1+\tau_j} - \frac{\tau_j}{(1+\tau_j)^2}\right) = 0, \text{ for all } j = 0, 1, \dots, \infty.$$

Recall that we are assuming that $\beta = (1+r)^{-1}$. Hence we can manipulate this equation to produce:

$$\frac{1}{1+\tau_j} = \frac{1}{1+\tau_j}\mu\left(1 - \frac{\tau_j}{1+\tau_j}\right),$$

which reduces to:

$$\tau_j = \mu - 1.$$

This equation implies that the tax rate should not vary across periods (since μ is constant). This is one very important implication of our model: the optimal tax rate is constant. Thus we write:

$$\tau_t = \tau^\star, \text{ for all } t \geq 0.$$

Now let's find $\tau^\star$ by substituting into the government budget constraint in equation (14.17):

$$\frac{\mathbf{G}}{\mathbf{W}} = \sum_{t=0}^{\infty}(1+r)^{-t}\frac{\tau_t}{1+\tau_t} = \sum_{t=0}^{\infty}(1+r)^{-t}\frac{\tau^\star}{1+\tau^\star}.$$

We can rewrite this as:

$$\mathbf{W}\frac{\tau^\star}{1+\tau^\star} = \frac{\mathbf{G}}{\sum_{t=0}^{\infty}(1+r)^{-t}}. \tag{14.19}$$

Notice that this says that the amount of revenue collected by the government each period is constant and equal to the annuity value of government expenditures. Thus the government collects the same amount of revenue each period, running deficits when it has unusually high expenditures and surpluses when expenditures are low.

Implications for the Path of Debt.

Imagine a government that has to fight a war in period 0, and makes no other expenditures in all other periods. Let the cost of a war be unity. Thus government expenditures satisfy:

$$G_t = \begin{cases} 1, & t = 0 \\ 0, & t \geq 1. \end{cases}$$

Hence the present discounted value of government expenditures is:

$$\mathbf{G} = \sum_{t=0}^{\infty}(1+r)^{-t}G_t = (1+r)^0(1) + \sum_{t=1}^{\infty}(1+r)^{-t}(0) = 1.$$

We know from equation (14.19) that the optimal tax rate $\tau^\star$ satisfies:

$$\mathbf{W}\frac{\tau^\star}{1+\tau^\star} = \frac{1}{\sum_{t=0}^{\infty}(1+r)^{-t}} = 1 - \frac{1}{1+r} = \frac{r}{1+r}, \text{ so:}$$
$$\frac{\tau^\star}{1+\tau^\star} = \frac{r}{1+r}\frac{1}{\mathbf{W}}.$$

For the sake of argument, say that the household has a constant endowment $y_t = 1$ all $t \geq 0$. In other words, the government has to fight a war in the first period that costs as much as the total economy-wide wealth in that period. If this is the case then we can find $\mathbf{Y}$ and $\mathbf{W}$:

$$\mathbf{Y} = 1 \cdot \sum_{t=0}^{\infty}(1+r)^{-t} = \frac{1}{1-\frac{1}{1+r}} = \frac{1+r}{r}.$$
$$\mathbf{W} = \frac{\mathbf{Y}}{\sum_{t=0}^{\infty}(1+r)^{-t}} = \mathbf{Y}\left(\frac{1+r}{1+r-1}\right)^{-1} = 1.$$

This makes sense: the annuity value of $\mathbf{Y}$ is just the infinite flow of constant payments that equals $\mathbf{Y}$. But since $\mathbf{Y}$ is made up of the infinite flow of constant payments of $y_t = 1$ each period, then the annuity value must also be unity. Now we can find $\tau^\star$ from:

$$\frac{\tau^\star}{1+\tau^\star} = \frac{r}{1+r},$$

which means that $\tau^\star = r$. In other words, the optimal tax rate $\tau^\star$ is simply the interest rate r. Government tax revenues each period are:

$$\mathcal{T}_t = \frac{r}{1+r}, \text{ for all } t \geq 0.$$

The government is collecting this relatively small amount each period in our example.

Notice what this implies for the path of deficits (and hence debt). In period $t = 0$ the government collects $r/(1+r)$ and pays out 1 to fight its war. Hence the core deficit in period $t = 0$ is:

$$G_0 - \mathcal{T}_0 = 1 - \frac{r}{1+r} = \frac{1}{1+r}.$$

In all subsequent periods, the government spends nothing and collects the usual amount, hence running core surpluses (or negative core deficits) of:

$$G_t - \mathcal{T}_t = 0 - \frac{r}{1+r} = -\frac{r}{1+r}, \text{ for all } t = 1, 2, \ldots, \infty.$$

From the government's flow budget constraint we know that:

$$G_t + (1+r)B^g_{t-1} = \mathcal{T}_t + B^g_t, \text{ for all } t \geq 0.$$

Hence from period 1 onward, while the government is repaying its debt from period 0:

$$B^g_t = -\frac{r}{1+r} + (1+r)B^g_{t-1}, \text{ for all } t \geq 1.$$

From this is it easy to see that the government debt, after the war, is constant at:

$$B^g_t = \frac{1}{1+r}, \text{ for all } t = 1, 2, \ldots, \infty.$$

Each period, the government raises just enough revenue to pay the interest cost on this debt and roll it over for another period. Does this violate our transversality condition in equation (14.2)? No, because the government debt is not exploding, it is merely constant. Thus, from equation (14.2):

$$\lim_{t\to\infty} \left(\frac{1}{1+r}\right)^t B^g_t = \lim_{t\to\infty} \left(\frac{1}{1+r}\right)^t \left(\frac{1}{1+r}\right) = \lim_{t\to\infty} \left(\frac{1}{1+r}\right)^{t+1} = 0.$$

So the transversality condition is satisfied by the government's optimal debt plan.

Exercises

Exercise 14.1 (Easy)
Supply the following facts. Most can be found in the Barro textbook.

1. What was the ratio of nominal outstanding public debt to GNP for the U.S. in 1996?
2. About what was the highest debt/GNP ratio experienced by the U.S. since 1900? In what year?
3. According to Barro, about what was the highest marginal tax rate paid by the "average" American since 1900? In what year?
4. Illinois has a standard deduction of about $2000. Every dollar of income after that is taxed at a constant rate (a "flat" tax). What is that rate?

Exercise 14.2 (Moderate)
The government must raise a sum of G from the representative household using only excise taxes on the two goods in the economy. The household has preferences over the two goods of:

$$U(x_1, x_2) = \ln(x_1) + x_2.$$

Variable	Definition
B_t	Real end-of-period savings by the household sector.
B_t^g	Real par value of outstanding end-of-period government debt.
B_{-1}^g	Initial stock of government debt, $B_{-1}^g = 0$.
C_t, c_t	Aggregate (household) consumption at t.
G_t	Real government spending at t.
Y_t, y_t	Aggregate (household) income at t.
r_t	Real net return on end-of-period debt.
$\mathbf{C}$	Present discounted value of consumption stream, $\{C_t\}_{t=0}^{\infty}$.
$\mathbf{G}$	Present discounted value of government spending stream, $\{G_t\}_{t=0}^{\infty}$.
$\mathbf{T}$	Present discounted value of tax revenue stream, $\{\mathcal{T}_t\}_{t=0}^{\infty}$.
$\mathbf{W}$	Annuity value of income stream, $\{Y_t\}_{t=0}^{\infty}$.
$\mathbf{Y}$	Present discounted value of income stream, $\{Y_t\}_{t=0}^{\infty}$.
$\mathcal{H}(a;\psi)$	Government tax policy mapping household action a and a vector of tax parameters ψ into a tax bill (see Chapter 13).
$\mathcal{T}(\psi)$	Government revenue as a function of tax policy parameters ψ when the household chooses its best response, $a_{\max}$ (see Chapter 13).
$U(\{C_t\}_{t=0}^{\infty})$	Utility from a consumption stream.
$u(C_t)$	One-period utility from consumption of C_t in period t.
$V(\{\tau_t\}_{t=0}^{\infty})$	Household's indirect utility given a stream of excise taxes $\{\tau_t\}_{t=0}^{\infty}$ and an income stream $\{Y_t\}_{t=0}^{\infty}$ with annuity value $\mathbf{W}$.

Table 14.1: Notation for Chapter 14. Note that, with the assumption that $Y_t = N_t$, variables denoted as per-capita are also expressed as fractions of GDP.

The household has total wealth of M to divide between expenditures on the two goods. The two goods have prices p_1 and p_2, and the government levies excise taxes of t_1 and t_2. The government purchases are thrown into the sea, and do not affect the household's utility or decisions. Determine the government revenue function $\mathcal{T}(t_1, t_2; p_1, p_2, M)$. Determine the household's indirect utility function in the presence of excise taxes, $V(p_1+t_1, p_2+t_2, M)$.

Hence write down the problem to determine the distortion-minimizing set of excise taxes t_1 and t_2. *You do not have to find the optimal tax rates.*

Exercise 14.3 (Easy)
Assume that a household lives for two periods and has endowments in each period of $\{y_1 = 1, y_2 = 1 + r\}$. The household can save from period $t = 1$ to period $t = 2$ on a bond market at the constant net interest rate r. If the household wishes to borrow to finance consumption in period $t = 1$ and repay in period $t = 2$, it must pay a higher net interest rate of $r' > r$. The household has preferences over consumption in period t, c_t, of:

$$V(\{c_1, c_2\}) = u(c_1) + u(c_2).$$

Assume that $u' > 0$ and $u'' < 0$.

There is a government which levies lump sum taxes of $\mathcal{T}_t$ on the household in each period, $t = 1, 2$. The government can borrow and lend at the same interest rate, r. That is, the government does not have to pay a premium interest rate to borrow. The government must raise $\mathbf{G} = 1$ in present value from the household. That is:

$$\mathcal{T}_1 + \frac{\mathcal{T}_2}{1 + r} = \mathbf{G} = 1.$$

Answer the following questions:

1. Find the household's consumption in each period if it does not borrow or lend.
2. Assume that $\mathcal{T}_1 > 0$ and $\mathcal{T}_2 > 0$. Draw a set of axes. Put consumption in period $t = 2$, c_2, on the vertical axis and consumption in period $t = 1$ on the horizontal axis. Draw the household's budget set.
3. Now assume that $\mathcal{T}_1 = 0$ and $\mathcal{T}_2 = (1 + r)\mathbf{G}$. Draw another set of axes and repeat the previous exercise.
4. Now assume that $\mathcal{T}_1 = \mathbf{G}$ and $\mathcal{T}_2 = 0$. Draw another set of axes and repeat the previous exercise.
5. What tax sequence would a benevolent government choose? Why?

Exercise 14.4 (Easy)
Assume that the government can only finance deficits with debt (it cannot print money, as in the chapter). Assume further that there is a constant real interest rate on government debt of r, and that the government faces a known sequence of real expenditures $\{G_t\}_{t=0}^{\infty}$. The government chooses a sequence of taxes that produces revenues of $\{\mathcal{T}_t\}_{t=0}^{\infty}$ such that:

$$\begin{aligned}
&\mathcal{T}_0 = G_0 - 1, \text{ and:} \\
&\mathcal{T}_t = G_t, \text{ for all } t = 1, 2, \ldots, \infty.
\end{aligned}$$

Find the path of government debt implied by this fiscal policy. Does it satisfy the transversality condition? Why or why not?

Chapter 15

Comparative Advantage and Trade

Most people would rather have a job making computer chips rather than potato chips. This may be rational, but the speed with which people state their preference belies a common misconception. In fact, the one occupation is not necessarily more profitable than the other. Haitians, for example, can make more money per hour growing and harvesting peanuts than they could make building computers. Economists use the terms absolute advantage and comparative advantage in discussing such issues.

A worker (or a country of workers) has an *absolute advantage* in production of a particular good if that worker (or country) can produce the good using fewer inputs than the competition. For example, in producing a good that requires labor only, the worker who can make a unit of the good in the least amount of time has an absolute advantage in the production of that good. The United States has an absolute advantage in producing a number of goods, since its work force is extremely productive and its economy is very well organized. Guatemala has an absolute advantage in the production of bananas, because of the country's climate. Kuwait has an absolute advantage in the production of crude oil, since its plentiful reserves make it easier to extract oil.

The term *comparative advantage* dates back to Ricardo.[1] Suppose a worker (or a country of workers) can make some good x and sell it for p_x dollars per unit on an open market. Obviously, the worker would like to sell the good for as high a price p_x as possible. If p_x is low enough, the worker will switch to production of some other good. We say that the worker has a comparative advantage in the production of x if the worker (or country) will find it profitable to make x at lower p_x than that at which the competition will find it profitable. This will be made quite a bit clearer when we formalize the concept in Section 15.2.

For now, let's think about the production of a particular good: the amount you learn in

[1] Ricardo, David. *The Principles of Political Economy and Taxation*. 1817.

your macroeconomics course. Suppose you have a truly gifted instructor. This instructor could explain the textbook page by page, and teach that material better than the black and white textbook, i.e., the instructor has an absolute advantage over the textbook when it comes to expositing material linearly. Nonetheless, this would not be the best use of the instructor's time. He or she could be more productive by conducting in-class discussions and answering your questions. The textbook is terrible at answering your questions; your only hope is to look things up in the index and search through the text looking for an answer. The key point is that the instructor's time will be put to best use by doing the activity that he or she is relatively better at. This is the activity in which the instructor has a comparative advantage.

15.1 Two Workers under Autarky

We now move to a concrete model so as to be precise about the meaning of comparative advantage. There are two workers, Pat (P) and Chris (C), and two goods, wine (w) and beer (b). In this section we introduce the baseline case in which Pat and Chris live in autarky, i.e., they are not permitted to trade with each other. In the next section, we allow them to trade. It is the possibility of trade that raises the issue of comparative advantage.

Pat and Chris have H hours to devote to production each day. Use n_w^P to denote the number of hours that Pat needs to make a jug of wine. Similarly, use n_b^P for the number of hours that Pat needs to make a jug of beer. Replacing P with C gives us the time requirements of Chris. (Throughout this chapter, superscripts will denote whether the variable pertains to Pat or Chris, and subscripts will distinguish between variables for wine and beer.)

Pat's utility is: $U(c_w^P, c_b^P) = (c_w^P)^\gamma (c_b^P)^{1-\gamma}$, where c_w^P and c_b^P are Pat's consumption of wine and beer, respectively, and γ is some number between zero and one. Let h_w^P and h_b^P be the number of hours per day that Pat spends on production of wine and beer, respectively. That means that Pat will produce h_w^P / n_w^P jugs of wine each day. (For example, if $n_w^P = 4$, then it takes Pat 4 hours to make a jug of wine. If $h_w^P = 8$, then Pat spends 8 hours on wine production, so Pat makes 2 jugs of wine.)

Putting all this together, we get Pat's maximization problem:

$$\max_{c_w^P, c_b^P, h_w^P, h_b^P} \left\{ (c_w^P)^\gamma (c_b^P)^{1-\gamma} \right\}, \text{ such that:}$$

$$h_w^P + h_b^P = H,$$

$$c_w^P = \frac{h_w^P}{n_w^P}, \text{ and:}$$

$$c_b^P = \frac{h_b^P}{n_b^P}.$$

Substituting all the constraints into the objective yields:

$$\max_{h_b^P} \left\{ \left(\frac{H - h_b^P}{n_w^P} \right)^{\gamma} \left(\frac{h_b^P}{n_b^P} \right)^{1-\gamma} \right\}.$$

Taking the first-order condition with respect to h_b^P and solving yields: $h_b^{P\star} = (1 - \gamma)H$. Plugging this back into the time constraint gives us: $h_w^{P\star} = \gamma H$. (You should check to make sure you know how to derive these.) These optimal time allocations do not hinge on n_w^P and n_b^P because Pat's preferences are homothetic; it's not a general result.

We assume that Chris has the same preferences as Pat. All the math is the same; we just replace each instance of P with C. Chris's optimal choices are: $h_b^{C\star} = (1 - \gamma)H$ and $h_w^{C\star} = \gamma H$, just like for Pat.

Suppose $H = 12$; each day Pat and Chris have 12 hours in which to work. Further, assume $\gamma = 1/3$. This implies that $h_b^{P\star} = h_b^{C\star} = 8$ and $h_w^{P\star} = h_w^{C\star} = 4$. Now suppose Pat can make a jug of wine in 4 hours and a jug of beer in 6 hours. Chris can make a jug of wine in 3 hours and a jug of beer in 1 hour. Translating these values to our variables gives the values in Table 15.1. Since Chris can make a jug of in fewer hours than Pat, Chris has an absolute advantage in wine production. Chris also has an absolute advantage in beer production.

	Hours per Jug	
	Wine	Beer
Pat	$n_w^P = 4$	$n_b^P = 6$
Chris	$n_w^C = 3$	$n_b^C = 1$

Table 15.1: Time Requirements

Plugging these values into our formulae above, we get the consumptions of Pat and Chris. Namely, $c_w^{P\star} = h_w^{P\star}/n_w^P = 4/4 = 1$, etc. Table 15.2 contains the rest of the consumption values.

	Consumption	
	Wine	Beer
Pat	$c_w^{P\star} = 1$	$c_b^{P\star} = 4/3$
Chris	$c_w^{C\star} = 4/3$	$c_b^{C\star} = 8$

Table 15.2: Consumption under Autarky

15.2 Two Workers Who Can Trade

We now turn to a world in which Pat and Chris can trade with each other. The open-market prices of beer and wine are p_b and p_w, respectively. We assume that Pat and Chris take these prices as given. There are two pieces to the maximization problem that they face. First, given these prices, they want to choose a way to allocate their time to production of the two goods so as to maximize their income. Second, for any given income, they want to choose how to divide up their consumption between the two goods. Thankfully, we can consider the problems separately. We address production first.

Production under Trade

Recall, n_w^P is the number of hours that Pat needs to make a jug of wine. This means that Pat can make $1/n_w^P$ jugs of wine per hour. Similarly, n_b^P is the number of hours for Pat to make a jug of beer, so Pat can make $1/n_b^P$ jugs of beer per hour.

If Pat chooses to make wine, Pat's hourly wage will be the number of jugs times the price: $(1/n_w^P)(p_w)$. Similarly, if Pat makes beer, Pat's hourly wage will be: $(1/n_b^P)(p_b)$. Pat will choose which to produce based solely on which wage is highest. Accordingly, Pat makes wine if:

$$\left(\frac{1}{n_w^P}\right) p_w > \left(\frac{1}{n_b^P}\right) p_b \quad \text{or:}$$

$$\frac{p_w}{p_b} > \frac{n_w^P}{n_b^P}. \tag{15.1}$$

This makes sense. Pat is more inclined to make wine if the price of wine p_w is higher or if Pat is able to make more wine per hour (i.e., n_w^P is smaller). Pat is less inclined to make wine if the price of beer p_b is higher or if Pat is able to make more beer per hour (i.e., n_b^P is smaller).

We get a similar relation for Chris, who makes wine if:

$$\frac{p_w}{p_b} > \frac{n_w^C}{n_b^C}. \tag{15.2}$$

To figure out the aggregate supply decisions of Pat and Chris, we conduct a thought experiment. First we suppose that the relative price of wine p_w/p_b is very low. Then we ask what happens as the relative price rises. When p_w/p_b is very low, both will make beer because the return to making wine is low relative to the return to making beer. As the price rises, eventually one of the two workers will find it profitable to switch to wine production. Eventually the price will rise enough so that both will make wine.

The numbers on the right-hand sides of equations (15.1) and (15.2) are the relative efficiencies of Pat and Chris at wine production, respectively. Whoever has the smaller number on the right-hand side is said to have a comparative advantage in wine production. This is because that worker will find it profitable to make wine even at low relative prices for wine.

Using the numbers from the previous section, we see that Pat makes wine if:

$$\frac{p_w}{p_b} > \frac{4}{6},$$

and Chris makes wine if:

$$\frac{p_w}{p_b} > \frac{3}{1}.$$

Since $4/6 < 3$, Pat has a comparative advantage in the production of wine. Recall, Chris has an absolute advantage in the production of both wine and beer. The idea of comparative advantage is that Chris has a more significant absolute advantage in beer production. There will be relative prices levels at which Chris will not make wine even though Chris has an absolute advantage in that market, since it will be even more profitable for Chris to specialize in beer.

We could just as well have looked at production in terms of beer. This is just the flip side of the wine market; the roles of w and b are just reversed. Pat makes beer if:

$$\frac{p_b}{p_w} > \frac{n_b^P}{n_w^P},$$

and Chris makes beer if:

$$\frac{p_b}{p_w} > \frac{n_b^C}{n_w^C}.$$

Be sure you understand how these equations relate to equations (15.1) and (15.2). It turns out that in a market with two producers and two goods, if one producer has a comparative advantage in one market, then the other will have a comparative advantage in the other market. Using the example numbers above, Pat has a comparative advantage in wine production, so Chris has a comparative advantage in beer production.

Using our example numbers, we can construct the aggregate supply curve for wine produced by Pat and Chris. On the vertical axis we put the relative price of wine: p_w/p_b. On the horizontal axis we put the quantity of wine supplied. See Figure 15.1.

When the relative price of wine is really low, neither Pat nor Chris produce wine, so the quantity supplied is zero. As the relative price of wine rises above 4/6, it suddenly becomes profitable for Pat to make wine instead of beer. Since Pat works $H = 12$ hours and can make $1/n_w^P = 1/4$ jugs per hour, Pat's supply of wine is 3 jugs. (At a relative price of 4/6, Pat is indifferent between producing wine and beer, so Pat's wine production could be anything between 0 jugs and 3 jugs.) At prices between 4/6 and 3, Pat makes three jugs, and Chris still finds it profitable to make beer only.

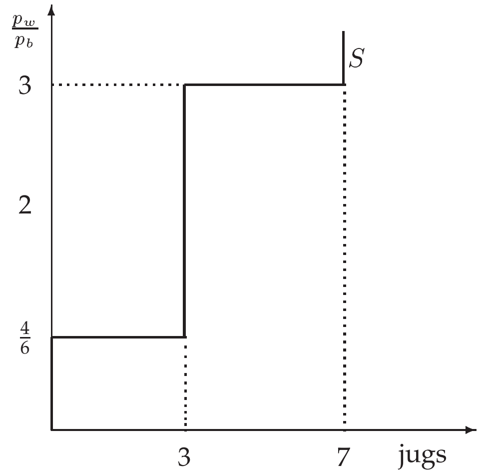

Figure 15.1: The Supply of Wine by Pat and Chris Together

When the relative price rises to 3, Chris finds it profitable to switch to beer production. Since Chris can make $(H)(1/n_w^C)$ = 4 jugs per day, aggregate production jumps up to 7 jugs. (At a relative price of 3, it is now Chris who is indifferent between producing wine and beer, so Chris could produce anything from 0 to 4 jugs of wine, and aggregate production could be anything from 3 to 7 jugs.) As the relative price continues to rise above 3, both Pat and Chris reap higher profits, but the quantity supplied does not change, since both have already switched to produce wine exclusively.

Consumption under Trade

We can use the production numbers and prices from the previous section to calculate the dollar incomes of both Pat and Chris. Let m^P and m^C be the incomes of Pat and Chris, respectively. For example, if p_w = 2 and p_b = 1, then p_w/p_b = 2, and Pat will make wine only. Since Pat can make 3 jugs of wine per day, Pat's income will be: m^P = (2)(3) = 6. Similarly, given our sample parameters, Chris makes beer only and earns an income of: m^C = 12.

In the general case, our task now is to determine the optimal choices of consumption for Pat and Chris when their incomes are m^P and m^C, respectively. This is just a standard consumer-choice problem. Pat's maximization problem is:

$$\max_{c_w^P, c_b^P} \left\{ (c_w^P)^\gamma (c_b^P)^{1-\gamma} \right\}, \text{ such that:}$$

$$c_w^P p_w + c_b^P p_b = m^P.$$

Pat's optimal choices are:

$$c_w^{P\star} = \frac{\gamma m^P}{p_w}, \text{ and:} \tag{15.3}$$

$$c_b^{P\star} = \frac{(1-\gamma)m^P}{p_b}. \tag{15.4}$$

(See Exercise 15.1 for the derivations.) The choices of Chris are analogous, with C replacing P.

Equilibrium under Trade

We now have all the pieces to determine the equilibrium prices $p_w^\star$ and $p_b^\star$. Given two candidate values for these prices, we use equations (15.1) and (15.2) to determine which goods Pat and Chris produce. Multiplying each worker's production by the prices gives each worker's income. We then use equations (15.3) and (15.4), and the equivalent versions for Chris, to determine what Pat and Chris will consume at those prices. If the sum of the their production equals the sum of their consumption for each good, then these candidate prices are an equilibrium.

Actually finding the correct candidate equilibrium prices is somewhat complicated. We consider possible prices in regions dictated by the supply curve. Consider the supply curve derived from the sample parameter values we have been using in this chapter. We might start by supposing p_w/p_b is between 0 and 4/6. It turns out that this would make supply of wine smaller than demand, so that can't be an equilibrium. Then we might suppose that $p_w/p_b = 4/6$. These prices too lead to excess demand.

It turns out that the equilibrium occurs at $p_w/p_b = 2$. For example, $p_w^\star = 2$ and $p_b^\star = 1$ is an equilibrium. At these prices: Pat makes 3 jugs of wine and no beer; and Chris makes 12 jugs of beer and no wine. Their optimal consumptions are in the Table 15.3. At these prices aggregate consumption of each good equals aggregate production of each good, so this is an equilibrium.

	Consumption	
	Wine	Beer
Pat	$c_w^{P\star} = 1$	$c_b^{P\star} = 4$
Chris	$c_w^{C\star} = 2$	$c_b^{C\star} = 8$

Table 15.3: Consumption under Trade

When we compare Table 15.3 with Table 15.2 we see that there are gains from trade. Pat's consumption of wine and Chris's consumption of beer are the same in each case, but under trade, Pat gets to consume 2/3 extra units of wine, and Chris gets to consume 2/3 extra units of wine. Accordingly, both Pat and Chris are made better off by trade. This is literally the single most important result from the theory of international trade. Free trade allows each worker (or country) to specialize in production of the good in which the worker (or country) has a comparative advantage. As a result, free trade is generally Pareto improving.

There is open debate among economists about just how often trade is, or can be, Pareto improving. The most difficult aspect of analysis along these lines is that citizens of a given country are not affected equally. Consider peanut exports from Haiti to the United States. Free trade in peanuts almost certainly makes just about every Haitian better off. Similarly, most consumers in the United States are made better off by free trade, because Haitian peanuts cost less than those produced in the United States, but U.S. peanut producers are almost certainly made worse off by unfettered imports of Haitian peanuts. Accordingly, free trade in peanuts would not be Pareto improving. That said, almost all economists agree that it could be made Pareto improving, simply by having U.S. consumers reimburse U.S. peanut growers for their losses due to imports. Such a move can be made Pareto improving because the benefits to consumers because of cheaper peanuts far outweigh the high profits U.S. peanut growers receive from blocking Haitian imports.

There are other arguments about whether free trade is Pareto improving. For example, many economists think that free trade can damage what are called "infant industries". If South Korean makers of automobiles had faced unfettered imports when they first started production, those auto makers might never have had enough time to learn how to make competitive products. By shielding their producers from competition when they were just getting started, South Korea may have allowed a productive and efficient industry to develop. Now that South Korea's auto industry is no longer "infant", free trade can almost certainly be made Pareto improving, but that industry might never have existed without some protection in the early years.

Variable	Definition
P	Designates that a variable pertains to Pat
C	Designates that a variable pertains to Chris
w	Designates that a variable relates to wine
b	Designates that a variable relates to beer
n_i^j	Hours worker j requires to make a jug of good i
H	Total hours available for work in a day
h_i^j	Hours worker j spends making good i
p_i	Price of a jug of good i
m^j	Dollar income of worker j
c_i^j	Consumption of good i by worker j

Table 15.4: Notation for Chapter 15

Exercises

Exercise 15.1 (Moderate)
Derive the equations for $c_w^{P\star}$, $c_b^{P\star}$, $c_w^{C\star}$, and $c_b^{C\star}$ from Section 15.2. See equations 15.3 and 15.4.

Exercise 15.2 (Moderate)
Pat and Chris work 8 hours each day. They each try to make as much money as possible in this time. Pat can make a jug of wine in 2 hours and a jug of beer in 1 hour. Chris can make a jug of wine in 6 hours and a jug of beer in 2 hours. Pat and Chris are the only producers of wine in this economy. The price of wine is p_w, and the price of beer is p_b. The daily demand for wine is:

$$Q_w^D = 11 - 2\left(\frac{p_w}{p_b}\right).$$

1. Graph the aggregate supply curve for wine.
2. Graph the demand curve for wine (on the same graph).
3. Determine the equilibrium relative price of wine (i.e., the value of p_w/p_b that causes supply to equal demand).
4. Calculate the equilibrium values of: (i) the amount of wine made by Pat; (ii) the amount of beer made by Pat; (iii) the amount of wine made by Chris; and (iv) the amount of beer made by Chris.
5. Does either Pat or Chris have an absolute advantage in wine production? If so, which does?
6. Does either Pat or Chris have a comparative advantage in wine production? If so, which does?

Chapter 17

Financial Intermediation

In this chapter we consider the problem of how to transport capital from agents who do not wish to use it directly in production to those who do. Some agents are relatively wealthy and already have all of the productive capital they need. Others accumulate capital for retirement, not production. In each case, the agents would want to lend their surplus capital to other agents, who would then use it in production. In the real world this lending takes the form of loans to individuals and businesses for the purpose of undertaking risky ventures.

Capital transportation of this form is known as *financial intermediation.* The institutions that stand between *savers* (those with surplus capital) and *borrowers* (those with less capital than they would like to use in their productive technology) are known as financial intermediaries. The most common financial intermediary is the bank, so the study of intermediation is sometimes also known as banking.

In this chapter we will examine how banks operate, starting from the bank's balance sheet, its role in matching lenders and borrowers and continuing on to its ability (or desire) to make loans. We will study competitive equilibria in banking. And we will consider the inherent instability built into many banking systems.

The balance sheet of a bank is a little unusual at first sight. The key fact to remember is that, for the bank, loans are assets, in exactly the same way that vault cash and government bonds are, while deposits are liabilities. Every financial instrument that is an asset on one balance sheet must be a corresponding liability on another balance sheet. For example, loans on a bank's balance sheet are assets, while those same loans are liabilities on the borrower's balance sheet.

We will consider a completely worked-out model with a competitive equilibrium in banking. This model, originally due to Williamson, provides several important insights. First, banks will act as aggregators of deposits, bundling together several small deposits to make one large loan, as we actually observe in practice. Second, some agents will be completely unable to get loans in equilibrium because the bank finds it too expensive to make loans to them. These agents are *credit rationed*, which acts something like a "credit crunch" in reality. Third, banks will use a *pure debt* contract with *default* in dealing with borrowers. That is, the bank will loan a borrower an amount, say $100 at an interest rate r, say 5%, and expect to be repaid $105 at the end of the period no matter how the borrower's finances have changed in that period. Even if the borrower's house burns down and employer goes bankrupt, the bank wants its $105. The borrower's only recourse is to declare bankruptcy, hand over all assets and consume nothing (or some very low amount). This is known as *defaulting*. The bank could have written a very different loan contract, something like, "Pay me $110 if your house doesn't burn down, but $10 if your house does burn down". Assuming that there is only a 5% chance that the borrower's house will burn down, the bank will in expected value get $105, and the borrower would vastly prefer such an insurance contract. The absence of such contracts must be explained. The explanation we use here is one of *moral hazard.* The bank has no way of knowing whether or not the borrower's house has really burned down without paying an *audit cost.* Thus the borrower always has an incentive to lie (hence the term "moral" hazard) and claim misfortune.

Next we turn our attention another model with moral hazard. In this model there are no audit costs, but agents will supply a secret amount of labor effort. Because labor effort is secret, lenders will not be able to directly contract on it, and borrowers will supply lower-than-optimal amounts of effort. Agents will also be of different wealth levels, which will allow us to think about how credit is provided to rich people as opposed to poor people. We will see that poorer agents will work less hard, pay a higher interest rate and default more frequently than richer agents. This effect will be so strong that certain very poor agents will be credit rationed. These very poor agents will save their meager assets.

Finally we will consider the celebrated model of *bank runs* by Diamond and Dybvig. Bank runs refer to financial panics in which depositors rush to their bank to liquidate their assets, usually because they doubt their bank's ability to make payments. The most famous bank runs happened during the Great Depression and indeed, according to Diamond and Dybvig, they might have greatly contributed to the economic collapse of that period. The model has continuing appeal because, although the American banking system has been somewhat insulated from panics, the banking systems of other countries continue to succumb to panics. It is an unfortunate fact of life that banking panics still plague us. In Diamond and Dybvig's model, agents will look out of their windows, see other agents running to the bank and be immediately compelled to also run. The first (luckier or fleeter) agents to the door of the bank withdraw all of their deposits, leaving nothing for the remaining (slower) agents.

17.1 Banking Basics

In this section we briefly review some important concepts in modern American banking. We begin with a discussion of accounting for banks. We then turn to the fractional reserve banking system, and examine how the government manipulates the money supply. We conclude this section with a discussion of an important banking reform proposal, that of narrow banking.

Assets and Liabilities of a Bank

Banks are businesses, and like all enterprises they incur liabilities and accumulate assets. The confusing thing about banks is that by accepting a deposit, which is after all an inflow of money to the institution, the bank has incurred a liability. Since a liability is an obligation to pay, the bank has, by accepting the deposit, promised to pay the depositor the amount of his deposit plus accumulated interest either on demand or at a particular time. In the same way, by making a loan, which is an outflow of money from the bank, the bank has accumulated an asset. An asset is a claim to payment, and by making the loan, the bank has a claim on some repayment schedule of principal plus interest. Not all loans are repaid, so the bank must estimate the expected value of loans that will not be repaid and count this against its assets. Thus bank balance sheets have an item marked "Outstanding Loans net of Loss Reserve". This loss reserve is a polite term for the expected value of loan defaults.

A bank also holds, by law, a certain proportion of its deposits in zero-interest accounts with the Federal reserve system. These are also assets, although pretty low-yield ones. Banks hold a very small amount of "vault cash" which is currency (notes and coins) held at the bank (usually in impressive safes). This is used to meet the cash needs of depositors day-to-day. Banks also directly hold securities like U.S. Government debt (bonds). The nature of these securities is limited by law, so in the U.S. banks are not big stock market players. Banks also often directly own property, such as the bank building itself.

Like other business, banks make operating profits or losses as the value of assets and liabilities fluctuate. If a bank makes a profit, so that assets exceed liabilities, a residual liability is added to balance assets and adjusted liabilities. Thus profit is a liability. The opposite is true for losses.[1]

In this chapter we are going to assume that banks are zero cost enterprises with no assets other than loans and no liabilities other than deposits. We will assume that banks make zero economic profit in expectation, so the expected return on loans must cover the amount owed to depositors.

[1]For more information on this and other oddities in banking, see Mathis Dewatripont and Jean Tirole *The Prudential Regulation of Banks*, 1994, MIT PRess: Cambridge, MA.

Fractional Reserve Banking and the Money Supply

In the U.S. today, banks are required to hold a certain fraction of their deposits in reserve, that is, not lend them out (these reserves are held on deposit by the retail banks at Federal Reserve member banks). Since the reserve requirement is not 100%, banks may lend out the portion of their deposits not held on reserve. Reserve requirements are indexed by the nature of deposits and loans, so banks have to hold extra reserves against riskier loans, but let us suppose for a moment that they are constant at 10%. Imagine that the government prints \$100 and gives it to household 0, so $H_0 = 100$. This household immediately places it in the bank or spends it. Any amount spent must go as a payment to some other household, which then faces the same choice: deposit or spending. And so on, until the banking system has absorbed the entire \$100 of cash.

In this way, the fractional reserve banking system can multiply an infusion of cash, augmenting the M1 money supply by more than the infusion. Reserves held at the Fed will also affect the money supply in the same way. For this reason, cash and reserves held at the Fed are often called *base money* or *high-powered money.* The sum of all cash and deposits held at the Fed is called the *monetary base.*

Assume for simplicity that each household simply deposits the cash. Assume also that there is only one common bank. Now since the reserve requirement is 10%, the bank places \$10 of its new deposits on reserve at the local Federal Reserve system member bank. The remaining \$90 it lends out again to some other household, household 1, so $H_1 = 90$. This household spends or deposits the money, as before, so a further \$90 of deposits appear in the bank. Now the bank sends \$9 to the Federal Reserve, and lends out \$81 to household 2, so $H_2 = 81$. This process continues until the bank is lending out, to household i an amount H_i:

$$H_i = \$100(1 - 0.10)^i.$$

The amount of *new money* created is just the sum of all loans made to households as a result of the original \$100 transfer, plus that \$100. That is:

$$M' - M = \sum_{i=0}^{\infty} H_i = \sum_{i=0}^{\infty} 100(1 - 0.10)^i = 100 \sum_{i=0}^{\infty} 0.90^i = 100 \frac{1}{1 - 0.90} = 1000.$$

Here M' is the new stock of money and M is old stock of money. Thus via the fractional reserve banking system, the \$100 initial cash transfer of the government grows to be \$1000 worth of new deposits in the banking system, which can then be used for payments purposes.

Targeting "The" Interest Rate

The U.S. government does not, as a rule, print money and then hand it out to agents (as we shall see in Chapter 18, however, other governments do exactly that). Instead, monetary policy is controlled by the Federal Reserve system (known as the “Fed”). The specific instrument of monetary policy most commonly used are *open-market operations* which the Fed uses to target “the” interest rate. The level of “the” interest rate is closely watched by industry and the media.

The interest rate that is at the center of all this fuss is called the “Fed funds rate”. It is associated with a very specific and abstract credit contract that is executed only between banks. No private individual has ever executed a Fed funds contract. Here is how the Fed funds market works: Banks are constantly gaining and losing deposits as individuals move, experience good or bad fortune or die. At the close of business each day, banks must meet their reserve requirement at the Fed. Imagine that a bank, say Hyde Park Bank, sees a sudden surge in deposits one day (say the first day of classes at the University of Chicago). Hyde Park Bank must place more funds on reserve at the Fed to meet its reserve requirement. Since deposits at the Fed earn no interest, it is likely that Hyde Park Bank does not have any excess reserves. Moreover, it takes time and money to make a deposit at the Fed. To meet the sudden shortfall in required reserves, Hyde Park Bank turns to other banks, who might have surplus reserves (caused by a sudden outflow of deposits). Thus Hyde Park Bank borrows reserves held at the Fed from other banks. Reserves held at the Fed are called Fed funds. The market for these reserves is called the Fed funds market. The interest rate on Fed funds (the interest rate that Hyde Park Bank pays) is the Fed funds rate. It is “the” interest rate.

Recall that deposits at a bank are liabilities. Thus Fed funds, an asset of banks, are a liability of the government. However, they are not an onerous liability since the government does not pay interest on these deposits. The Federal reserve system is the banker’s bank, and its books must also balance.

The Fed does indeed announce a target for the Fed funds rate. However, it affects the rate with more than just the moral suasion of an announced target: the Fed affects the supply of Fed funds directly through open market operations (OMOs). In an OMO *purchase* the Fed trades Fed funds for assets held at banks. Thus the Fed approaches a bank, say Hyde Park Bank, which holds millions of dollars of U.S. Government debt (bonds) and offers to buy a thousand dollars of these bonds. How will the Fed pay for these bonds? With “store credit”, that is, by creating $1000 of new reserves held at the Fed in the name of Hyde Park Bank. Hyde Park Bank is neither a net loser or a net gainer from this operation, since it has traded $1000 of one safe asset (government bonds) for $1000 of another (Fed funds). What will Hyde Park Bank do with the new Fed funds? It could in principal trade them for cash. For this reason, we say that cash is a *liability of the government* in exactly the same way that Fed funds are. However, Hyde Park Bank will probably just float them on the Fed funds market, driving down the Fed funds rate. Thus OMO purchases are associated with decreased Fed funds rates and an increased money supply (since Hyde Park Bank can trade its new Fed funds for cash). This is an *expansionary* monetary policy.

In the opposite way, an OMO *sale* tends to increase the Fed funds rate and decrease the money supply. In an OMO sale, the Fed approaches Hyde Park Bank and offers to sell it some government bonds which the Fed holds. The Fed accepts as payment Fed funds.

Thus the stock of Fed funds decreases and the Fed funds rate increases. Since Hyde Park Bank had to give up an asset that could be traded directly for cash (Fed funds), OMO sales cause a decrease in the money supply. This is a *contractionary* monetary policy.

Narrow Banking

In modern America, the banking system has two jobs. Intermediation is one job (and the primary focus of this chapter). The other job is acting as a *payments system.* That is, deposits at banks can be accepted as legal payment for goods and services, often in the form of checks (drafts on deposits). Although we are accustomed to thinking of these two jobs as linked, there is absolutely no inherent reason why they should be. In fact, true banking reform would institute so-called *narrow banking*, in which banks are not allowed to lend. They would have to back deposits completely with a riskless asset (usually government bonds). This is equivalent to a 100% reserve requirement. In this scenario, banks would concentrate only on providing low-cost payments services like checks, smart cards, direct deposit and direct withdrawal.

In a world of narrow banking, households seeking loans, for example a home mortgage, would not approach banks. Neither would households seeking a greater return on their investment than that provided by the narrow banks. Instead they would approach non-bank financial intermediaries. These intermediaries would be like mutual funds: savers would own an equity share of the mutual fund, which would have some constantly quoted price. Borrowers would negotiate with the mutual fund in much the same way that they would with a bank. The key difference is that the value of the mutual fund would not be fixed, like the value of deposits at a bank. In good economic times the mutual fund will perform better than in bad economic times. As we shall see in the section on bank runs below, this kind of arrangement effectively prevents financial panics.

Payments Systems

Before turning to models of intermediation, it is worth briefly discussing the payments system[2]. The payments system is the infrastructure, law and custom governing our ability to pay for goods and services. The most common methods of payment by households are cash, check and credit card. Of these three methods, only cash requires no up-front investment by the household. To have a checking account costs (on average) \$80 a year, and since credit cards are a form of unsecured debt, households must frequently pay some minimum amount each year to use one. Households too poor to afford either must use so-called "Currency Exchanges", which are ubiquitous features of working-class neighborhoods. These institutions will cash paychecks and welfare checks (for a fee), allow households to pay

[2]The numbers in this subsection are drawn from the paper "Retail Payments Instruments: Costs, Barriers, and Future Use" by David Humphrey and Lawrence Pulley, presented at the Conference on Bank Structure and Competition in Chicago, May 7, 1998.

their utility bills in cash (for a fee) and write "money orders" (a form of check), again for a fee. A year's worth of transactions at one of these exchanges can easily cost much more than $80, but the costs are directly contingent upon use, unlike checking accounts or credit cards, and do not require anything in the way of legal documentation.

Although there is about $400 billion in U.S. currency (cash) outstanding, only 10%-15% of this total cash (based on surveys of households) is used by households. Over 60% of total cash is thought to be outside the U.S. In the U.S., roughly 70% of all transactions are settled in cash, while the same figure is 78% in Holland, 83% in Finland, 86% in Germany and 90% in the United Kingdom. The average value of these transactions is small, below $10.

In the U.S., of all non-cash transactions, 75% are settled by check, and the average check is for $1,158. About 20% are settled by credit card, with an average value of $61. It is interesting to note that although less than 0.5% of transactions are settled by wire transfer, the average size of a wire transfer is over $4.2 million. The relatively large values of these figures stems from the fact that they are used by the government and other institutions. Households' non-cash transactions average below $50.

The total cost of the U.S. payments system to payors, banks and payees is $204 billion a year (about 3% of GDP) or roughly $1050 per adult. Each non-cash payment costs on average $2.60.

Finally, although the check system is deeply entrenched in the U.S., the new electronic payments means are substantially cheaper, costing a third to a half as much as checks.

Looking across countries, the average American made 326 non-cash transactions in 1996. The corresponding figure for Canada is 151, for Europe it is 126 and for Japan a paltry 40. Of these, checks accounted for 244 of the transactions in the U.S., 62 in Canada, 31 in Europe and only two in Japan. Thus check clearing is much more common the U.S. than even in Canada. These other countries rely more heavily on debit cards and other electronic means (although the total number of non-cash transactions abroad is still well below the level in the U.S.).

For most of the post-war period, the U.S. has led the world in payments system technology. In other countries it is still quite common to pay even very large bills with cash. However, the backbone of the domestic payments system, the check clearing system, is showing its age. Other countries are already using sophisticated electronic systems such as debit cards and electronic check presentation, that have yet to become common here. Although most economists do not consider payments systems directly relevant to intermediation and the conduct of monetary policy, there is no question that they have a large and direct impact on all households in an economy.

17.2 A Model with Costly Audits

In this section we will consider a model that replicates several of the important features of financial intermediation. This model is originally due to Williamson (1987). Intermediaries will arise endogenously to evaluate the credit worthiness of borrowers, to bundle small deposit together to make a large loan, to minimize the cost of monitoring borrowers and to spread default risk across several lenders, so they are fully insured against default by borrowers.

This model is exciting because we are, for the first time, going to calculate a competitive equilibrium in the capital market with heterogeneous agents. In addition, this capital market is going to suffer from a realistic problem: audit costs. As a result, in equilibrium, some agents will be unable to get loans no matter what interest rate they promise to pay, so there will be *credit rationing* in equilibrium. Also, the audit costs are going to produce very familiar credit contracts, the fixed-obligation loan.

Agents

In this model there will be two kinds of agents. Type-1 agents, who form a proportion α of the population, will be *workers.* Type-2 agents, who form the remaining $1 - \alpha$ proportion of the population, will be *entrepreneurs.*

There will be two periods. In the first period, type-1 agents work some amount, consume and then save. Intermediaries accept these savings as deposits and use them to make loans to type-2 agents. Type 2 agents will operate their risky technologies in the "morning" of the second period, realize their outcomes and repay the intermediaries. The intermediaries will then take these payments and repay their depositors. There is no way to store the consumption good between periods.

Type-2 agents are born with a number stenciled on their foreheads. This gives their *audit cost.* That is, if a type-2 agent is born with γ_i on his forehead, it costs γ_i for any other agent to observe the outcome of his technology (more on that below). Note that the audit cost γ is public. These audit costs are distributed uniformly on the interval $[0, 1]$. Thus 25% of all type 2 agents will be born with audit costs below 0.25, and the remaining 75% of type-2 agents will have audit costs greater than 0.25.

Preferences

Type-1 agents care about consumption in both periods. Their preferences over consumption c_0 and c_1 and labor effort exerted in the first period ℓ_0 are as follows:

$$U^1(c_0, \ell_0, c_1) = u(c_0, \ell_0) + c_1. \tag{17.1}$$

Here $u_1(\cdot,\cdot) > 0$, $u_2(\cdot,\cdot) < 0$ and $u_{22} < 0$.

Type-2 agents care only about consumption in the second period. Their preferences over consumption c_1 are simply:

$$U^2(c_1) = c_1. \tag{17.2}$$

So type-2 agents are risk neutral in consumption and supply no labor effort.

Technology

Type-1 agents may work up to h hours while young (in the first period of life), where $h < 1$. Labor effort is transformed directly into the consumption good (so that the implicit "wage rate" is just unity). The risk-free interest rate offered by the financial intermediaries is r, so type-1 agents face budget constraints of the form:

$$c_0 = \ell_0 - s, \text{ and:} \tag{17.3}$$

$$c_1 = (1 + r)s. \tag{17.4}$$

Where $\ell_0 \leq h$ and s is the savings of a type-1 agent.

All type-2 agents have access to the same technology, no matter what their audit cost is (that is, they differ only in their audit cost). In exchange for a capital input of $k = 1$ while young, type-2 agents will operate a risky technology that produces output of y, where, for a particular agent i, output y_i is:

$$y_i = 1 + \varepsilon_i.$$

The idiosyncratic shock term ε_i is distributed uniformly on the interval $[0, 1]$. So the mean or average shock is 0.5 and mean output is 1 + 0.5 = 1.5.

Notice that type-2 agents have no intrinsic wealth of their own–they must get a loan of size $k = 1$ to finance their projects. Since the absolute maximum that a type-1 agent can produce, by working full-time ($\ell_0 = h$) is h, which by assumption is less than unity, a type-2 agent's project can be financed only by a loan from several type-1 agents. Thus intermediaries are going to have to aggregate deposits from several savers to make a single productive loan, which matches the real world experience well.

Finally, the output y_i of a particular type-2 agent's technology is *private* to that agent. Only agent i may costlessly observe y_i. All other agents must pay the audit cost γ_i stenciled on agent i's forehead to observe y_i.

Intermediation

Although we will not explicitly model the industrial organization of banking in this paper (that is, we will not write down how banks are formed, who owns them and so on), it is easy to see what the optimal banking structure will look like. Since the only cost to forming a bank is auditing defaulting borrowers, any agent can declare himself a bank, accept deposits and make loans. The larger the bank, the more insurance that depositors will have against defaults. Imagine a bank that matched n depositors, each of whom deposits some amount $1/n$ with exactly one borrower. There is no insurance at all for those n depositors. If the borrower defaults, then all n depositors lose everything. Now imagine that a second bank opens across the street from this first bank. This bank seeks to match $2n$ depositors with two borrowers. If one borrower defaults, depositors may still recover something from the other borrower. The second bank will clearly attract more depositors than the first bank. In this way, larger and larger banks will form, until all the potential depositors go to the same bank, which also makes all of the loans.

Now let us turn our attention to the question of optimal loan contracts. The bank advances a type-2 agent a loan of $k = 1$. The agent then experiences output y, which is secret. The bank is interested in minimizing the expected audit cost it will have to pay. One contract would be for the borrower to report his output, and for the bank to take, say 10% of that output as payment for the loan, and the bank never audits. The problem with this contract is that borrowers will always announce that $y = 1$, the minimum possible output. Since the bank is not auditing, there is no way for it to dispute this claim.

The bank wishes to minimize the cost of audits while at the same time ensuring that borrowers tell the truth. It turns out that the best contract for this purpose is one in which the bank announces some required repayment level X. No matter what output is, the bank insists on getting its amount X. Any borrower who announces output $y < X$ is declared bankrupt and audited. Whatever output that agent had produced (no matter how large or small) is seized by the bank. Borrowers with outcomes $y < X$ have no incentive to announce anything other than the truth, while borrowers with outcomes $y > X$ know they must always pay X or be audited, so they merely announce y, pay X and consume the residual $y - X$.

It is interesting to note that is exactly the kind of debt contract that is actually written. Lenders make loans at some interest rate, and expect to be repaid the loan amount plus interest. If borrowers do not repay this amount, they are declared bankrupt and legal proceedings begin. Of course in our society bankrupt agents do not consume zero, and in fact their minimum consumption varies from state to state, but the underlying principle is the same. Risk averse agents would prefer more complicated credit contracts, ones that provided a certain amount of insurance along with the loan. For example, most homeowners would be willing to pay above-market interest rates on their mortgages if their bank agreed to cut their mortgage payments if something bad happened to the homeowner (loss of job, injury etc). The reason such contracts are not more common is because output is largely hidden. Lenders want to be assured of repayment without having to closely monitor their borrowers (an expensive proposition).

Let us return to the problem of a financial intermediary making a standard loan of size $k = 1$ to a particular borrower with known audit cost γ, requiring repayment X and auditing the borrower if output falls below X. Write X as $1 + x$, so x is the net interest rate paid by borrowers. In addition, recall that output y is always between 1 and 2, depending on the value of the shock term ε. The actual revenue of the intermediary conditional on the repayment amount x, the audit cost γ, and the shock term ε is:

$$\pi(x,\gamma,\varepsilon) = \begin{cases} 1+x, & \varepsilon \geq x \\ 1+\varepsilon-\gamma, & \varepsilon \leq x. \end{cases}$$

The bank knows that the output shock ε is distributed uniformly on the interval $[0,1]$. Assuming that $0 \leq x \leq 1$, the expected revenue of the intermediary conditional only on x and γ is:

$$\begin{aligned} \pi(x,\gamma) &= \int_0^1 \pi(x,\gamma,\varepsilon)\, d\varepsilon \\ &= 1 + \int_0^x (\varepsilon-\gamma)\, d\varepsilon + x\int_x^1 d\varepsilon \\ &= 1 + \left(\frac{x^2}{2} - \gamma x\right) + x(1-x) \\ &= 1 + (1-\gamma)x - \frac{x^2}{2}. \end{aligned}$$

It is easy to see that expected revenues $\pi(x,\gamma)$ as a function of x are a parabola with a peak at $x^\star(\gamma) = 1-\gamma$. No bank would ever charge a repayment amount x greater than $x^\star(\gamma)$, since revenue is declining in x beyond that point, and borrowers are worse off.

Let $\pi^\star(\gamma)$ be the absolute maximum amount of revenue that the bank can accumulate as a function of the audit cost, γ. That is:

$$\pi^\star(\gamma) = \pi\left[x^\star(\gamma),\gamma\right] = 1 + \frac{1}{2}(1-\gamma)^2.$$

Thus for agents with $\gamma = 1$, the bank's maximum revenue occurs when $x^\star(\gamma = 1) = 0$ and the bank never audits, producing a revenue of $\pi^\star(\gamma = 1) = 1$ with certainty. For agents with $\gamma = 0$, the maximum revenue that the bank can extract occurs when $x^\star(\gamma = 0) = 1$ and the bank's expected revenue is $\pi^\star(\gamma = 0) = 3/2$.

Now let us turn our attention to the liabilities of the bank. The bank owes its depositors an amount $1+r$ on a unit loan. Thus for each borrower of audit cost γ, banks will pick the lowest value of x such that $\pi(x,\gamma) = 1+r$. There will be some agents, with relatively high values of γ, for which banks will never be able to realize an expected return of $1+r$. That is, those agents with audit costs $\gamma \geq \gamma^\star(r)$ such that:

$$\pi^\star[\gamma^\star(r)] = 1+r.$$

will be *credit rationed.* Banks will never make them a loan, they will be squeezed out of the credit market and their projects will not be funded. It is easy to see that $\gamma^\star(r)$ is given by:

$$\gamma^\star(r) = 1 - \sqrt{\frac{r}{2}}.$$

Thus when $r = 0$, $\gamma^\star = 1$ and no agents are credit rationed. As r increases, more agents are credit rationed, and at the astronomical interest rate of $r = 2$, all agents are credit rationed.

Equilibrium

It is easy to see now how equilibrium in this economy will be achieved. There will be a demand for capital, which is decreasing in r, and a supply of capital, which is increasing in r. The demand for capital is given by the number of agents who are not credit rationed at the interest rate r. That is, the proportion of type-2 agents (remember, they make up $(1-\alpha)$ of the total) with $\gamma \le \gamma^\star(r)$. Thus the aggregate demand for capital is given by:

$$K^d(r) = (1-\alpha)\gamma^\star(r) = (1-\alpha)\left(1 - \sqrt{\frac{r}{2}}\right).$$

As r decreases, more and more type-2 agents may finance their projects. Each type-2 agent always wants exactly one unit of capital, so individual borrowing is constant but aggregate borrowing increases.

The supply of capital comes from the saving schedules of type-1 agents. As r increases, type-1, or worker, agents each save more. Say that the savings schedule of workers is given by $S(r)$. Thus the supply of capital is:

$$K^s(r) = \alpha S(r).$$

Equilibrium in the capital market will occur at the interest rate $r^\star$ at which:

$$K^s(r^\star) = K^d(r^\star).$$

That is, where the supply of capital from type-1 agents' savings equals the demand for capital from intermediaries lending to type-2 agents, so they can finance their projects.

An Example

Imagine that type-1 agents had preferences over consumption while young c_0, labor effort while young ℓ_0, and consumption while old c_1 given by:

$$U^1(c_0, \ell_0, c_1) = u(c_0, \ell_0) + c_1 = 2\sqrt{c_0} + 2\sqrt{h - \ell_0} + c_1.$$

Recall from the budget constraints, equations (17.3) and (17.4), that, by substituting in the savings term s, this becomes:

$$U^1(\ell_0, s) = 2\sqrt{\ell_0 - s} + 2\sqrt{h - \ell_0} + (1+r)s.$$

We take derivatives with respect to ℓ_0 and s to find the first-order necessary conditions for maximization:

$$\frac{1}{\sqrt{\ell_0 - s}} - \frac{1}{\sqrt{h - \ell_0}} = 0, \text{ and:}$$
$$-\frac{1}{\sqrt{\ell_0 - s}} + (1 + r) = 0.$$

We can solve these to find the savings schedule of type-1 agents:

$$s(r) = h - \frac{2}{(1 + r)^2}.$$

The aggregate supply of capital is thus:

$$K^s(r) = \alpha \left[h - \frac{2}{(1 + r)^2} \right].$$

Notice that $K^s(r_0) = 0$ where $r_0 = \sqrt{2/h} - 1$.

For equilibrium to occur with this specification of preferences, the interest rate at which there is zero demand for capital must exceed the interest rate at which there is zero supply. The zero-demand interest rate we know from above to be $r = 2$. The zero-supply interest rate is $r = r_0$. Thus for equilibrium:

$$\sqrt{\frac{2}{h}} - 1 < 2.$$

It is easy to see that this means that $h > 2/9$. Indeed, as h gets closer and closer to 2/9, the supply curve shifts upward. This in turn causes the equilibrium interest rate to rise and the equilibrium supply and demand of capital to fall, decreasing the number of projects that are undertaken.

17.3 A Model with Private Labor Effort

The model of the previous section was very useful in thinking about equilibrium in the credit market. However, since in the real world we do not necessarily recognize "type-1" and "type-2" agents from birth, and we do not readily observe different audit costs, it makes sense to consider a different model. In this model all agents will be identical except for their wealth level. Some agents will be richer, others poorer. There will be no audit cost, and output will be public. However, agents are going to have to work to make the project succeed, and the amount of their labor effort will be *private* (that is, hidden). It can never be known (cannot be audited). This is another example of moral hazard.

Technology, Endowments, Preferences

All agents will have access to a common "back-yard" technology which will map capital k and labor effort ℓ into a probability that the project succeeds. If the project succeeds, output is high, at q. If it fails, output is zero. Capital k can take on only two values: $k = 0$ or $k = 1$. If $k = 1$ then the probability of the high output is just given by the level of labor effort, ℓ. If $k = 0$ then the low output occurs with certainty, no matter what the effort level was.

Agents are all endowed with some level of wealth a. For a particular agent, if $a < 1$, then that agent must get a loan of size $1 - a$ to operate the technology. If $a > 1$, then the agent can finance the technology alone and lend the surplus $a - 1$.

Agents have preferences over consumption c and labor effort ℓ given by:

$$c - \frac{q}{\alpha}\frac{\ell^2}{2}.$$

Here $0 < \alpha < 1$ and q is just the high output level.

Rich Agents

Assume that there is some riskless rate of return r on wealth. A rich agent, one with $a > 1$, can finance the project from her own wealth and lend the remainder at this interest rate r. How much effort does she supply? If the project fails and output is zero, she consumes $c = 0+(1+r)(a-1)$. If the project succeeds and output is q, she consumes $c = q+(1+r)(a-1)$. The project succeeds with probability ℓ, her labor effort. Thus her maximization problem is:

$$\max_{\ell}\left\{\ell[q + (1 + r)(a - 1)] + (1 - \ell)[(1 + r)(a - 1)] - \frac{q}{\alpha}\frac{\ell^2}{2}\right\}.$$

The first-order condition for maximization with respect to ℓ is:

$$q - \frac{q}{\alpha}\ell = 0.$$

This implies that $\ell^\star(a) = \alpha$ for $a \geq 1$. That is, agents wealthy enough to finance the project out of their own funds all supply the same labor effort, α.

Given that all rich agents supply $\ell = \alpha$ regardless of their wealth, we can easily calculate their expected utility by plugging back in:

$$\alpha q + (1 + r)(a - 1) - \frac{q}{\alpha}\frac{\alpha^2}{2}, \text{ or:}$$
$$\frac{\alpha q}{2} + (1 + r)(a - 1).$$

Thus the value from operating the technology is $\alpha q/2$, and the value from lending any excess capital is $(1 + r)(a - 1)$. The opportunity cost of the capital used in the productive process is $(1 + r) \cdot 1$ (since production requires one unit of capital). Thus for any agent to undertake the project, it must the case that:

$$\frac{\alpha q}{2} \geq 1 + r.$$

For the rest of this section we are going to assume that the interest rate r is below $\alpha q/2 - 1$, so agents will want to use the productive technology.

Poor Agents

Now consider the much more interesting case of a poor agent, with wealth $a < 1$, who seeks a loan from a financial intermediary of size $1 - a$ in order to finance the project. For now we will assume that the intermediary charges the borrower an amount X if the project succeeds and zero if it fails. This is again the root of the moral hazard problem: the borrower only repays the bank if her project succeeds. This decreases the incentive for the borrower to work.

Now if the project succeeds, the agent consumes $c = q - X$, that is, output q net of repayment X. If the project fails, the agent consumes $c = 0$. Thus the agent is repaying only in the state when the project succeeds. The agent's choice of effort comes from the maximization problem:

$$\max_{\ell} \left\{ \ell(q - X) + (1 - \ell)(0) - \frac{q}{\alpha}\frac{\ell^2}{2} \right\}.$$

The first-order necessary condition for maximization with respect to effort ℓ is:

$$(q - X) - \frac{q}{\alpha}\ell = 0.$$

Solving for $\ell^{\star}(X)$ (effort as a function of repayment) produces:

$$\ell^{\star}(X) = \alpha\left(1 - \frac{X}{q}\right).$$

Notice immediately that $\ell^{\star}(0) = \alpha$ and that $\ell^{\star}(X)$ is decreasing in X. Hence poor agents, borrowers, will work less hard than rich agents, lenders.

Intermediaries

Now consider the problem of the intermediary raising deposits and making loans to poor agents. If this intermediary makes a loan of size $1 - a$, it must pay its depositors an amount

$(1+r)(1-a)$ on this loan. Thus its expected return on the loan must be equal to $(1+r)(1-a)$, the risk-free cost of capital. If the project succeeds the bank makes X, if it fails, the bank makes 0. The probability of success is ℓ, which cannot be observed or controlled directly by the bank. The bank must take the agent's effort choice as a function of repayment $\ell^\star(X)$ as given. So the bank's zero profit condition is:

$$\ell^\star(X) \cdot X + [1 - \ell^\star(X)] \cdot 0 - (1+r)(1-a) = 0.$$

Substituting the borrower's choice of labor effort conditional on repayment, $\ell^\star(X)$, in from above, the bank's zero-profit condition becomes:

$$X\alpha\left(1 - \frac{X}{q}\right) - (1+r)(1-a) = 0.$$

This is a quadratic equation in X and may be written as:

$$X^2 - qX + \left(\frac{q}{\alpha}\right)(1+r)(1-a) = 0.$$

Using the quadratic formula, we obtain two possible values for $X(a)$, that is, repayment as a function of wealth:

$$X(a) = \frac{q \pm \sqrt{q^2 - 4\left(\frac{q}{\alpha}\right)(1+r)(1-a)}}{2}$$

Since competitive pressures will force intermediaries to charge the lowest possible value of X, we concentrate on the lower branch. Notice that $X(1) = 0$.

Consider the term inside the radical in the definition of $X(a)$ above:

$$q^2 - 4\left(\frac{q}{\alpha}\right)(1+r)(1-a).$$

Notice that if a is small and r large, then this term might be negative. This means that, for poor borrowers, there is no value of X for which the intermediary can recover the cost of making the loan. Call this critical wealth $a^\star(r)$. Notice that $a^\star(r)$ satisfies:

$$a^\star(r) = 1 - \frac{\alpha q}{4(1+r)}.$$

Agents with wealth below $a^\star(r)$ are *credit rationed.* What will these agents do? They will become so-called *poor savers* and invest their meager funds in the economy-wide mutual fund.

Thus there will be three classes of agents in this model: the poor, who wish to borrow but cannot, and so save; the middle-class, who cannot self-finance but can borrow; and the rich, who self-finance and save the remainder of their wealth. Thus the demand for loans comes entirely from the middle-class, while the supply of funds comes from both the rich and the very poor. Increases in the interest rate will help the rich and the poor, and hurt the middle class.

17.4 A Model of Bank Runs

Now we turn our attention to the celebrated model of bank runs from the paper "Bank Runs, Deposit Insurance, and Liquidity" by Diamond and Dybvig. As Diamond and Dybvig say:

> Bank runs are a common feature of the extreme crises that have played a prominent role in monetary history. During a bank run, depositors rush to withdraw their deposits because they expect the bank to fail. In fact, the sudden withdrawals can force the bank to liquidate many of its assets at a loss and to fail. In a panic with many bank failures, there is a disruption of the monetary system and a reduction in production.

This neatly sums up the important features of the discussion. Bank runs can be a self-reinforcing phenomenon: if one agent sees others running for the bank, she must also join the run or face the certain loss of her wealth if the bank should fail. This sudden demand for cash (also called the demand for *liquidity*) causes the bank to sell (or *liquidate*) its assets (loan portfolio) at a loss, so it may be unable to satisfy the demands of its depositors. The real loss caused by premature liquidation is the fundamental reason why bank runs are bad.

How can bank runs be stopped? The authors consider two possibilities. First, *suspension of convertibility*, in which a bank temporarily refuses to cash out deposits. This is also known as a *bank holiday* and was quite common in the financial panics of the Great Depression. We shall see that bank holidays will only work under special conditions. Second, the authors consider *deposit insurance*, in which the government taxes all agents in order to honor banks' obligations. Deposit insurance will prevent bank runs even under very general conditions, and so we conclude that they are a more robust way of preventing bank runs. The authors do not consider the possibility of replacing the bank with a *mutual fund*, but, as we shall see, this too would prevent bank runs.

In this section we will consider a very different reason for banks to exist than in the previous two sections. Earlier, we viewed banks as institutions for getting capital from rich savers to poor borrowers (roughly speaking). In this section, all agents will have identical wealth and productive opportunities, but they will differ in the timing of their demands for consumption. Some agents will be content to consume later, while others will want to consume immediately. Agents will not know their type when they invest. Think of it this way: all agents are perfectly identical, except that some have cancer. Cancer diagnoses are announced only after all agents have made their investment decisions. When they are diagnosed, the cancerous agents want to consume immediately, while the noncancerous agents are content to wait until the following period to consume. The productive technology is (as we shall see) illiquid, so the cancerous agents are forced to prematurely liquidate. Banks will convert their illiquid assets into liquid liabilities. In doing so, the bank will leave itself open to the possibility of a bank run. If there is a bank run, then all assets are prematurely liquidated and there is real economic harm done.

As in the previous two models that we have considered, there is an informational problem. Here, there will be no way for banks to distinguish agents who truly have urgent consumption needs (our "cancerous" agents) and those who do not (the "noncancerous" agents).

Technology, Endowments and Preferences

There are three periods, $t = 0, 1, 2$. There is a single homogeneous good, and agents are endowed in period $t = 0$ with one unit of this good. There is a common, riskless, technology which converts a unit invested in $t = 0$ to $F > 1$ units in period $t = 2$. If the technology is interrupted in the middle period, $t = 1$, the *salvage value* is just the unit again. Think of this as a "growing turnip" technology. All agents are endowed with a turnip at birth in period $t = 0$, which they plant. If they uproot the turnip in the second period of life, $t = 1$, they just get their original turnip back. If they leave the turnip in the ground all the way to the harvest date of $t = 2$, it will have grown to $F > 1$ turnips.

Agents have no desire to consume in period $t = 0$. Let c_1 be consumption in period $t = 1$ and c_2 be consumption in period $t = 2$. Agents will have preferences of the form:

$$U(c_1, c_2; \Theta) = \begin{cases} \ln(c_1), & \text{if } \Theta = 1 \\ Q\ln(c_1 + c_2), & \text{if } \Theta = 2. \end{cases}$$

Here $1 \geq Q > F^{-1}$. The term Q is less than one. As a result, agents with $\Theta = 2$ (noncancerous) have a lower marginal utility than agents with $\Theta = 1$ (cancerous). So not only do agents with $\Theta = 1$ have to consume in period $t = 1$, they have a high marginal utility to boot. This is sometimes known as being "urgent to consume".

Agents of type $\Theta = 1$ have no desire to consume in period $t = 2$ at all, while agents of type $\Theta = 2$ are indifferent between consumption in $t = 1$ and $t = 2$. An agent with $\Theta = 1$ has cancer and must consume while young while one with $\Theta = 2$ does not have cancer and is indifferent between consumption while young or old.

Assume that there is some probability θ of having $\Theta = 1$ (that is, θ is the probability of getting cancer). With probability $1 - \theta$, $\Theta = 2$. Assume further that there is a continuum of agents, so a *proportion* θ will get cancer and the remaining $1 - \theta$ will not.

Optimal Insurance Contracts

Let c_t^i be the consumption of an agent of type i in period t. Without banks, because one's type is private, there can be no insurance contracts, so all agents with $\Theta = 1$ uproot their turnip in $t = 1$ and consume $c_1^1 = 1, c_2^1 = 0$. All agents with $\Theta = 2$ leave the turnip in the ground until $t = 2$ and consume $c_1^2 = 0, c_2^2 = F$.

Since agents are risk averse, they would prefer to be assured of some consumption between the low level of 1 and the high level of F. If Θ were public (that is, commonly observed), zero-cost insurance companies would provide agents with insurance contracts. Optimal insurance contracts would have the feature that $c_2^1 = c_1^2 = 0$ since agents with $\Theta = 2$ are content to wait. The budget constraint of the insurance company (equation 1c in Diamond and Dybvig) is a little hard to understand at first glance. Think of it like this: a proportion θ of the population will get c_1^1. This leaves $1 - \theta c_1^1$ in the ground between period $t = 1$ and $t = 2$, where it grows to $F(1 - \theta c_1^1)$. This is then spread between the remaining $1 - \theta$ of the population. Thus c_2^2 must satisfy:

$$c_2^2 \leq F \frac{1 - \theta c_1^1}{1 - \theta}.$$

This is equivalent to:

$$\theta c_1^1 + \frac{(1 - \theta) c_2^2}{F} \leq 1. \tag{17.5}$$

The insurance companies' Lagrangian is:

$$\mathcal{L}(c_1^1, c_2^2, \lambda) = \theta \ln(c_1^1) + (1 - \theta) Q \ln(c_2^2) + \lambda \left(1 - \theta c_1^1 - \frac{(1 - \theta) c_2^2}{F} \right).$$

This has first-order conditions of:

$$\begin{aligned} \frac{\theta}{c_1^1} - \lambda\theta &= 0, \\ Q \frac{1 - \theta}{c_2^2} - \lambda \frac{1 - \theta}{F} &= 0, \text{ and:} \\ \theta c_1^1 + \frac{(1 - \theta) c_2^2}{F} &= 1. \end{aligned}$$

We can solve these equations for the *optimal consumptions,* call them $c_1^{1\star}$ and $c_2^{2\star}$. We find:

$$c_1^{1\star} = \frac{1}{\theta + Q(1 - \theta)}, \text{ and:} \tag{17.6}$$

$$c_2^{2\star} = \frac{QF}{\theta + Q(1 - \theta)}. \tag{17.7}$$

Note that by assumption $QF > 1$, so $c_2^{2\star} > c_1^{1\star}$. In a perfect insurance contract, type-2 agents (agents with $\Theta = 2$) consume more than type-1 agents (unlucky agents with $\Theta = 1$).

Demand Deposit Contracts

The optimal insurance contract can be recaptured by a demand deposit contract provided by banks. Banks will accept deposits from agents in period $t = 0$. At period $t = 1$, some depositors will be of type 1, and will approach the bank to withdraw their deposits early. Their turnips have not matured, but the bank will rip up other agents' turnips to provide type-1 agents with more than just their unit turnip in return. In period $t = 2$, the remaining (type-2) depositors will get whatever is left over.

The bank has no way of telling which agents are type 1 and which are type 2, so it structures contracts like this: Deposits placed at $t = 0$ will earn an interest rate of r_1 if withdrawn in period $t = 1$ and r_2 if withdrawn in period $t = 2$. Of course, since $r_1 > 0$, it is technologically impossible for the bank to pay off all agents the amount $1 + r_1$ in period $t = 1$, since at that time no turnips have actually matured. However, the bank has a technical legal liability to pay off $1 + r_1$ to any depositor who appears at its door in period $t = 1$.

The banking system as a whole faces a *sequential service constraint*. This constraint is fundamental to the operation of banks in this model. It requires that depositors be honored in the order in which they show up at the bank. Even though the bank can look out the window and see a line that clearly exceeds its capacity to pay, it must pay out $1 + r_1$ on a first-come, first-served basis.

These demand deposit contracts will have two equilibria. The first equilibrium will be the "good" equilibrium and will not feature bank runs. The second will be the "bad" equilibrium and will feature a bank run.

Begin with the first equilibrium. Agents of type 1 (and only those agents) appear at the bank in period $t = 1$ requesting their deposits plus interest, withdrawing $1 + r_1$ each. In period $t = 2$, the remaining agents split what is left (remember that all the turnips left in the ground from period $t = 1$ to $t = 2$ will have grown by a factor of F). Thus c_1^1 and c_2^2 are related by:

$$c_1^1 = 1 + r_1, \text{ and:}$$
$$c_2^2 = F\frac{1 - \theta(1 + r_1)}{1 - \theta}.$$

Notice that although the bank announces the interest rates $\{r_1, r_2\}$, they are not independent. Choosing a value of r_1 automatically fixes r_2. For the rest of this section we will not calculate r_2 explicitly.

For this equilibrium to work, type-2 agents must not prefer the contract offered to type-1 agents. Since, for type-2 agents, consumption in periods $t = 1$ and $t = 2$ are perfect substitutes, it must be the case that, for the banking equilibrium to work:

$$c_2^2 \geq c_1^1.$$

Substituting in from above, this translates to:

$$F\frac{1-\theta(1+r_1)}{1-\theta} \geq 1+r_1.$$

We can rearrange this find:

$$\theta(1+r_1) + \frac{(1-\theta)(1+r_1)}{F} \leq 1. \tag{17.8}$$

Notice that this is exactly the same as equation (17.5) above, the budget constraint on the optimal insurance contract.

How do banks set r_1? They maximize the expected utility of their borrowers:

$$\theta \ln(1+r_1) + (1-\theta)Q\ln\left(F\frac{1-\theta(1+r_1)}{1-\theta}\right). \tag{17.9}$$

But this is exactly the same objective function as we used above, in the optimal insurance contract. Thus $1+r_1$ will satisfy:

$$1+r_1 = c_1^{1\star}.$$

Here $c_1^{1\star}$ is from equation (17.6) above. Thus demand deposits can replicate perfectly the optimal insurance contract.

If all type-2 agents stay at home in period $t = 1$, everything works perfectly. Type-2 agents will only be willing *not* to go to the bank in period $t = 1$ if they are certain that no other type-2 agent will be withdrawing in period $t = 1$.

This is the key to the bad equilibrium in this model, the bank run. If type-2 agents suspect that other type-2 agents are withdrawing from the bank in period $t = 1$, their consumption in period $t = 2$ will diminish. If enough type-2 agents attempt to withdraw in period $t = 1$, there will be nothing left in period $t = 2$. Thus if a type-2 agent believes that other type-2 agents are going to the bank to withdraw in period $t = 1$, their optimal response is also to withdraw in period $t = 1$. In a bank run equilibrium, the entire population appears at the bank in period $t = 1$ demanding $1 + r_1$. Since $1 + r_1 > 1$, and there is only one unit of the consumption good present in the bank, the bank liquidates its entire stock of consumption good to satisfy the first $1/(1 + r_1)$ proportion of agents in line. All other agents get nothing in $t = 1$ and, of course, nothing in $t = 0$.

Notice that the bank run has a real economic cost: by liquidating the turnip crop in period $t = 1$, none is left to grow in period $t = 2$ and total economy-wide output goes down. Moreover, in a bank run, some agents lose the entire value of their endowment. Although Diamond and Dybvig don't model it, we would expect this outcome to lead to social unrest.

Suspension of Convertibility

During the financial panics of the of the Great Depression, banks would often close their doors when faced with a bank run. Remember that in those days deposits were entirely uninsured, so depositors were desperate to realize any part of their deposits. As more and more banks refused to honor their deposits, the Federal government declared several *bank holidays,* during which no banks (solvent or insolvent) could open their doors. The banks were, in effect, suspending the ability of their depositors to convert their deposits to cash.

Diamond and Dybvig's model provides us a way to think about how suspension of convertibility works. It turns out to be an effective deterrent against bank runs only if θ is known in advance.

The complete derivation of this result is beyond the scope of this section, but we can sketch it out here. Imagine that θ is known with certainty. The bank announces that only the first θ depositors in line in period $t = 1$ will be served. A type-2 agent faces no penalty for staying at home in period $t = 1$ even if other type-2 agents are going to the bank. He is secure that there will be no excessive liquidation, and that his deposits will mature as expected in the next period. Indeed, it is to his benefit to have other type-2 agents withdraw early, in period $t = 1$, since the total number of withdrawals is capped at θ, the more type-2 agents who withdraw early, the fewer type-2 agents will be left in period $t = 2$ to share the value of the remaining deposits.

What if θ is not known with certainty? The first thing to establish that, in principal, nothing is different. Imagine that are two possible values of θ: high, with $\theta = \theta_1$ and low, with $\theta = \theta_0$. Say that the high-θ outcome occurs with probability ζ.[3] Then the expected utility of an agent who consumes c^1 if type 1 and c^2 if type 2, is:

$$\zeta\left[\theta_1 u(c^1) + (1-\theta_1)Qu(c^2)\right] + (1-\zeta)\left[\theta_0 u(c^1) + (1-\theta_0)Qu(c^2)\right].$$

This can be rearranged as:

$$[\zeta\theta_1 + (1-\zeta)\theta_0]u(c^1) + [\zeta(1-\theta_1) + (1-\zeta)(1-\theta_0)]Qu(c^2).$$

Define $\bar{\theta}$ to be the expected value of θ:

$$\bar{\theta} = \zeta\theta_1 + (1-\zeta)\theta_0.$$

The expected utility may be rewritten using $\bar{\theta}$ as:

$$\bar{\theta}u(c^1) + (1-\bar{\theta})Qu(c^2).$$

When forming expectations, agents use the expected value of θ.

Imagine that θ_1 is quite high, approaching one. The bank cannot suspend convertibility at any proportion below θ_1, because it cannot know the true value of θ. In fact, no one in the

[3] The Greek letter ζ is called "zeta".

economy knows the true value of θ. Imagine the plight of a type-2 agent watching agents in line in front of the bank. Is this a bank run? Are there type-2 agents in that line? Or is it simply the case that the high-θ outcome has been realized? If there are type-2 agents in that line, the optimal response for the type-2 agent is also to get in line, since there is a probability that in fact the low-θ outcome has been realized, and real economic damage is being done. This story bears a striking resemblance to what actually happened during the financial panics of the Great Depression: in the midst of confusion about the true state of liquidity demand, banks kept their doors open, forcing other agents to run to the bank.

Deposit Insurance

Deposit insurance will completely cure bank runs, even if θ is not known. In this model, deposit insurance is more than a promise by the government to honor all deposits. Since the stock of turnips is limited, the government must also tax agents to honor deposits. Deposit insurance works like this: in period $t = 1$ a certain number of agents apply to withdraw their deposits and realize $1 + r_1$. If the banks can honor these deposits and still invest enough between $t = 1$ and $t = 2$ to honor the remaining deposits, the government does nothing. If there is excess demand for withdrawals, the government begins taxing depositors to honor all the demand deposits in $t = 1$ and to ensure that deposits are honored in period $t = 2$. Agents (of both types) who withdraw their deposits in $t = 1$ will, if there is a bank run, consume less than $1 + r_1$, because of the taxes used to finance the deposit insurance.

From the point of view of a type-2 agent, even if other type-2 agents are running to withdraw in period $t = 1$, he is assured that there will be enough invested to honor his deposit in period $t = 2$. Thus there is no benefit to joining in the run. Indeed, because of the excess demand for withdrawals in period $t = 1$ precipitated by a bank run, all agents (type 1 and type 2) who rush to cash out their deposits in period $t = 1$ will realize less than the $1 + r_1$ they are owed because they are taxed by the government.

Mutual Funds

The multiple equilibria in this model of banking depend critically on the presence of the sequential service constraint. By relaxing this constraint, we can overcome the bad equilibrium.

A sequential service constraint is an integral part of a banking system with fixed-obligation deposit contracts. That is, if a bank is going to promise $1 + r_1$ to anyone who walks through the door in period $t = 1$, it is bound to serve its customers sequentially. Doing away with the sequential service constraint means doing away with banking entirely.

As an alternate system, consider a mutual fund. This is exactly the kind of institution that would replace banks in a narrow-banking system. All agents trade their turnips in period $t = 0$ for a single share in the mutual fund. In period $t = 1$ there will be a market for shares in the mutual fund: agents will be able to cash them out at some price p_1 for the consumption good. If all agents decide to cash out their shares, this price will be unity. In period $t = 2$, the remaining shareholders will split the remaining assets of the mutual fund. If some proportion α of the population wish to trade in their shares at some price p_1, the remaining proportion of population will consume p_2 in period $t = 2$, where p_2 is given by the familiar equation:

$$p_2 = F\frac{1-\alpha p_1}{1-\alpha}.$$

It must be the case that $p_2 \geq p_1$ or no agents (not even type-2 agents) will be willing to hold on to the mutual fund until period $t = 2$. This can be rewritten as:

$$\alpha p_1 + \frac{(1-\alpha)p_1}{F} \leq 1.$$

The competitive equilibrium in mutual fund shares will have the highest possible value for p_1, but p_2 will still be greater than p_1. As a result, only type-1 agents will sell out in period $t = 1$ and all type-2 agents will wait until period $t = 2$ to consume. This arrangement is not susceptible to runs. Imagine a type-2 agent in period $t = 1$ when other type-2 agents are "running" (in this case, selling out early). Since there must always be a competitive equilibrium, p_1 falls, and p_2 is always greater than p_1. As a result, our type-2 agent sees no benefit in joining the run, waits until period $t = 2$ and consumes $p_2 \geq p_1$. The key is that the sequential service constraint has been replaced by a competitive market in shares. Unusually high demand for consumption in period $t = 1$ is met by an unusually low price for shares in that period, p_1. In all cases, $p_2 \geq p_1$.

Variable	Definition
H_i	Cash loan or transfer to household i.
M', M	The new stock of money, the old stock of money.

Table 17.1: Notation for Chapter 17.

Exercises

Exercise 17.1 (Easy)
In the model of bank runs, explicitly calculate the interest rate on deposits held until period $t = 2$, r_2, when the interest rate on deposits held until period $t = 1$ is r_1.

Exercise 17.2 (Moderate)
For this problem, assume that there are only two types of potential borrowers: Safe (who comprise α of the population) and Risky (who comprise the remaining $1-\alpha$ of the population). Banks cannot tell the difference between them, and with probability α, a borrower is safe and probability $1-\alpha$ a borrower is risky. Safe borrowers have access to safe projects, which pay off π_S if they succeed and 0 if they fail. Safe projects succeed with probability p_S. Risky borrowers have access to risky projects, which pay off π_R if they succeed and zero if they fail. Risky projects succeed with probability p_R.

Risky and safe projects have the same expected payoff:

$$p_S \pi_S = p_R \pi_R,$$

but the probability of success is lower for risky projects, so $p_R < p_S$, and the payoff from succeeding is greater, so $\pi_R > \pi_S$. Both risky and safe projects have public failure, that is, there is no need to audit agents who claim that their project failed.

To finance the projects borrowers need a unit of capital from a bank. The bank in turn announces a repayment amount x in the event that the borrower's project does not fail. If the project fails, borrowers owe nothing (they declare bankruptcy). If the project succeeds, borrowers consume their output minus x, if the project fails, borrowers consume zero. Assume that borrowers are *risk neutral* so that their utility function is just their expected consumption.

There is a risk-free interest rate of r that banks must pay to their depositors (thus they have to realize at least $1+r$ in expected value on their loan to meet their deposit liability).

1. Write down a bank's balance sheet (in terms of x, r, p_S, and p_R) assuming that, with probability α the borrower is safe and with probability $1-\alpha$ the borrower is risky.
2. Assume that banks compete by offering the lowest value of x that gives them non-negative profits in expectation. Determine the equilibrium interest rate $x^\star(r,\alpha)$ as a function of the interest rate r and the proportion of safe agents α.
3. Find the expected utility of a safe agent who borrows, $V_S(r)$, as a function of the interest rate r when x is given by $x^\star(r,\alpha)$. Repeat for a risky agent.
4. Agents stop borrowing if the expected utility of being a borrower falls below zero. Show that if a safe agent decides to borrow, a risky agents will too. Find the critical interest rate $r^\star$ at which safe agents stop borrowing. At interest rates greater than or equal to this critical value, $r \geq r^\star$ all safe agents leave the pool, so $\alpha = 0$. What happens to the equilibrium payment x?

Exercise 17.3 (Moderate)
Consider the model of costly audits again. Now suppose that intermediaries gain access to a technology which allows them to extract more from each borrower (that is, for each value of announced repayment x and audit cost γ, suppose $\pi(x,\gamma)$ shifts up). What happens to

the demand schedule of capital? What happens to the supply schedule of capital? What happens to the equilibrium interest rate? What happens to equilibrium economy-wide output? Are agents made better off or worse off?

Exercise 17.4 (Moderate)
Yale University costs 1 dollars to attend. After graduation, Yalies (that is, graduates of Yale) either land good jobs paying w or no job at all, paying nothing. The probability of landing the good job is π where π is hidden effort exerted by the Yalie. Yalies are born with wealth $a \geq 0$, and those Yalies born with wealth $a < 1$ must have a loan of $1 - a$ to attend. Yale University will act as a lender to those students. Yale must borrow at the risk-free gross interest rate $r > 1$ to finance the loans. Student borrowers who get the good job must repay Yale University some amount x out of their wages w. Student borrowers who do not land the good job pay nothing. All students have preferences over lifetime expected consumption $E(c)$ and private labor effort π of:

$$V(E(c), \pi) = E(c) - \frac{w}{\alpha}\frac{\pi^2}{2}.$$

Assume $0 < \alpha < 1$.

1. Start with a rich Yalie, with $a > 1$. Show that her optimal effort $\pi^\star$ is α.

2. Now consider poor Yalies, with $a < 1$, who must borrow to finance their education. Calculate a borrower's optimal effort $\pi(x)$ as a function of x.

3. Write down Yale University's expected profit on a loan to a student with wealth $a < 1$ as a function of x, assuming that Yale University knows $\pi(x)$ from Exercise (2).

4. Assume that Yale University operates a "fair lending policy" in which borrowers of wealth a must repay an amount $x(a) = r(1 - a)/\alpha$ if they get the good job. What is "fair" about this lending policy? Given this policy and your answer to Exercise (2) above, calculate a borrower's optimal effort as a function of their wealth. That is, write down $\pi[x(a)]$, and call it $\pi(a)$.

5. Show that, given Yale University's "fair lending policy", all Yalie borrowers exert less effort than rich Yalies, that is, for Yalies with wealth $0 \leq a < 1$, show that $\pi(a) < \pi^\star$ and that $\pi(a = 1) = \pi^\star$, where $\pi^\star$ is from Exercise (1) above.

6. Finally, show that given its fair lending policy, that Yale loses money on student loans, and that the loss is increasing in loan size. Why does the fair lending policy cost Yale money?

Variable	Definition
α	Proportion of population who are type-1 workers.
γ_i	The audit cost of agent i.
c_0	Consumption in the first period of life (type-1 agents only).
ℓ_0	Labor effort in the first period of life (type-1 agents only).
c_1	Consumption in the second and last period of life (both types of agent).
$U^1(c_0, \ell_0, c_1)$	Preferences of type-1 agents.
$U^2(c_1)$	Preferences of type-2 agents (risk neutral).
k	Capital input to type-2 agent's project, can take on only two values, $k = 0$ or $k = 1$.
y_i, y	Output of agent i's project, or just output.
ε	Shock to output, distributed uniformly on $[0, 1]$.
h	Maximum labor effort by a type-1 agent, $h < 1$.
$s, S(r)$	Savings of a type-1 agent, or of the representative type-1 agent (aggregate supply of capital).
r	Economy-wide equilibrium interest rate on capital.
X, x	Repayment amount, $X = 1 + x$.
$\pi(x, \gamma, \varepsilon)$	Revenue of bank on a loan (gross of the borrowing cost r) to an agent with audit cost γ, when the repayment amount is x and the production shock is ε.
$\pi(x, \gamma)$	Expected revenue (gross of the borrowing cost r) on a loan to an agent with audit cost γ, when the repayment amount is x. Expectation taken over the production shock ε.
$\pi^\star(\gamma)$	The highest possible expected revenue (gross of the borrowing cost r) on a loan to an agent with an audit cost of γ.
$x^\star(\gamma)$	The repayment amount that results in the highest revenue to the bank on a loan to an agent of audit cost γ.
$\gamma^\star(r)$	Largest value of the audit cost γ at which the bank can make enough revenues to cover the cost of borrowing, r.
$K^d(r)$	Aggregate demand for capital.

Table 17.2: Notation for the model of audit costs in Section 17.2

Variable	Definition
k	Capital input to productive technology, can take only two variables, $k = 0$ or $k = 1$.
a	Wealth of agent.
ℓ	Private labor effort of agent.
c	Consumption of agent.
q	High output of technology (the low output is zero).
α	Disutility of effort in agent's preferences.
r	Economy-wide risk-free rate on capital.
X	Repayment amount.
$a^{\star}(r)$	Threshold credit rationing wealth.

Table 17.3: Notation for model with moral hazard in Section 17.3

Variable	Definition
F	Technology parameter: growth of asset between periods $t = 1$ and $t = 2$.
$U(c_1, c_2; \Theta)$	Utility function over consumption in period $t = 1$, $t = 2$ and shock term Θ.
Θ, θ	Shock term: $\Theta = 1$ means that the agent is urgent to consume (probability θ).
Q	Preference parameter: marginal utility of type-2 agents, $Q < 1$.
r_1, r_2	Interest rate promised by the bank on deposits held until period $t = 1$ (r_1) or period $t = 2$ (r_2).
c_j^i	Consumption by agent to type i in period j.
$c_1^{1\star}, c_2^{2\star}$	Optimal consumption in period $t = 1$ by type-1 agents and the optimal consumption in period $t = 2$ by type-2 agents.

Table 17.4: Notation for model of bank runs in Section 17.4

Chapter 18

Fiscal and Monetary Policy

In Chapter 14 we described how the government changes its outstanding debt over time so as to match its revenues and expenditures. In that framework, there was nothing intrinsically harmful about government debt. Now we turn our attention to the inflationary effects of persistent government budget deficits. This will give us a theory about the interplay between fiscal and monetary policies.

Imagine a government comprised of two competing authorities: a fiscal authority (in the U.S., the Congress and the President) and a monetary authority (in the U.S., the Federal Reserve System). The fiscal authority wants to finance government spending, while the monetary authority wants to keep inflation low. But inflation produces revenue for the government through a process known as *seigniorage.* If the monetary authority is dominant, it simply announces a sequence of inflation rates, which in turn implies a sequence of seigniorage revenues, and the fiscal authority takes this as given when making spending decisions. Completely dominant central banks are, however, extremely rare. Even the U.S. Federal Reserve System is statutorily a creature of the Congress and the Executive, and must, by law, balance the twin goals of fighting inflation and targeting full employment.

The case of a central bank (i.e., a monetary authority) that is not fully dominant is much more interesting. Note that this does not mean that the fiscal authority controls the money supply absolutely, merely that the fiscal authority does not have to credibly commit to a sequence of taxes sufficient to finance, in present value, its spending. In particular, we are going to assume that there is some limit on the debt-to-GDP ratio. That is, investors will only accept government debt up to some ceiling, defined in proportion to output. The monetary authority will control the money supply until this ceiling is reached, and thereafter it will fully accommodate government spending with seignorage revenue. This is the fashion in which government deficits are inflationary.

After presenting the theory, we will discuss the evidence. In a study of post-WWI hyperinflations in Germany, Austria, Hungary and Poland titled "The Ends of Four Big Inflations", Thomas Sargent illustrates this effect clearly. More recent monetary disturbances in some the successor states of the old Soviet empire can also be traced to persistent government budget deficits. This will provide us with a practical guide on how to end hyperinflations.

18.1 Are Government Budget Deficits Inflationary?

The model for this section is taken from a paper by Thomas Sargent and Neil Wallace, "Some Unpleasant Monetarist Arithmetic". The interested reader is advised to read the original paper, since it doesn't require very much math and is, despite the title, actually quite pleasant.

Government Budget Constraint

We will consider the problem of a government which must cover a sequence of real core deficits $\{D_t\}_{t=0}^{\infty}$:

$$D_t = G_t - \mathcal{T}_t, \text{ for all } t = 0, 1, \ldots, \infty,$$

where G_t is the real value of government expenditures and $\mathcal{T}_t$ is the real value of government revenues in period t. Notice that interest payments on the debt are not included in D_t (see Chapter 14 for more on the government budget constraint).

The government has some amount B_{t-1}^g of real debt outstanding at the beginning of each period t. The government must pay its creditors a real amount $(1 + r)B_{t-1}^g$ in period t. Hence the total real excess spending of the government on goods and services and debt service, net of tax revenue, is:

$$\text{borrowing demand} = D_t + (1 + r)B_{t-1}^g.$$

The government will finance this in two ways: (1) By issuing more bonds, dated end-of-period t (call these bonds B_t^g) and (2) By printing money and realizing the *seigniorage* revenue (more on what that is in a second). Hence government borrowing is:

$$\text{borrowing supply} = B_t^g + \frac{M_t - M_{t-1}}{P_t}.$$

Here M_t is the end-of-period quantity of pieces of paper with the words "Federal Reserve Note" and "In God We Trust" printed on them, also known as fiat currency. Take M_t to be strictly high-powered money, or the monetary base, which is under the control of the government.

For the government's books to balance it must borrow as much as it needs to, so:

$$D_t + (1+r)B^g_{t-1} = B^g_t + \frac{M_t - M_{t-1}}{P_t}, \text{ for all } t = 0, 1, \ldots, \infty. \tag{18.1}$$

Another way to write this is:

$$D_t + rB^g_{t-1} = (B^g_t - B^g_{t-1}) + \frac{M_t - M_{t-1}}{P_t}, \text{ for all } t = 0, 1, \ldots, \infty.$$

This form says that the government's real deficit plus the interest on the debt may be paid for by net new bonds ($B^g_t - B^g_{t-1}$) or seigniorage.

Seigniorage

The government has a monopoly on issuing pieces of paper with the words "Federal Reserve Note" written them. People want this stuff for transactions purposes, so they hold it even though it pays zero interest. As a result, the government can print more of the stuff and trade it for goods and services. We will not model the precise way in which the government does this. The effectiveness of this practice depends on how the general price level P_t responds to an increase in M_t.

Although seigniorage revenue in developed countries like the United States is currently very low, developing countries or countries in turmoil use it heavily. Internal bond markets and tax collection systems are often the first instruments of state power to vanish in turbulent times. Governments also often find direct taxation to be unpalatable for domestic political reasons, but are unable to sell bonds on international markets.

Consider the case of Zaire, an African country which is now called the Democratic Republic of the Congo. This government practiced a very bald form of seigniorage in which it would introduce a new denomination of the currency (the zaire), print up a bunch of notes and pack some of the print run into suitcases which were then distributed among government ministers. These ministers would then use the notes to purchase foreign currency on the black market as well as domestic goods and services. In the waning days of the rule of former president Mobuto Sese Seko, the government introduced the 500 zaire note and the 1000 zaire note. These were used, in part, to finance the president's cancer treatments in France. The population derisively termed the notes "prostates" and refused to accept them as payment in any transaction. The government's seigniorage revenue fell to zero and it succumbed to the rebels shortly thereafter.

More formally, the value of the seigniorage revenue in our model is the real value of net new notes:

$$\text{seigniorage} = \frac{M_t - M_{t-1}}{P_t}.$$

Notice that we will have to take a stand on how P_t varies with M_t to fully determine the seigniorage revenue.

Model Assumptions

To make this model work, we will have to specify a rule for output, population growth, how the price level is determined and what limits there are on borrowing. I list all of the model's assumptions here for convenience:

1. Output per capita y_t is constant and $y_t = 1$, but population N_t grows at the constant rate n, so $N_t = (1+n)N_{t-1}$, where $N_0 > 0$ is given. So total GDP each period $Y_t = y_t N_t$ is just equal to population.
2. The real interest rate on government debt is constant at $r_t = r$, and the government never defaults on its debt. This includes default by unexpected inflation when bonds are denominated in dollars. Thus we are dealing with inflation-indexed bonds. We also make the very important assumption that $r > n$. Without this assumption, most of the "arithmetic" is not so "unpleasant".
3. A stark monetarist Quantity Theory of Money relation with a constant velocity, $v = 1$:

   $$P_t Y_t = v M_t. \tag{18.2}$$

   Combine this with the definition of Y_t in Assumption (1) above to find the price level in period t, P_t, is:

   $$P_t = \frac{M_t}{N_t}.$$

4. There is an upper bound on per-capita bond holdings by the public of $\bar{b}$. That is, $B_t^g / N_t \leq \bar{b}$.

In addition, to make life easier, we will specify that the government's fiscal policy, which is a sequence of deficits $\{D_t\}_{t=0}^{\infty}$, is simply a constant per-capita deficit of d. Thus:

$$\frac{D_t}{N_t} = d, \text{ for all } t = 0, 1, \ldots, \infty.$$

Define b_t^g to be the level of per-capita bond-holdings $b_t^g \equiv B_t^g / N_t$. Assumption 4 states that $b_t^g \leq \bar{b}$ for some $\bar{b}$. Notice that with the assumption that the constant per-capita output level is $y_t = 1$, b_t^g is also the *debt-to-GDP ratio.* Also, D_t / N_t becomes the *deficit-to-GDP ratio.*

Monetary Policy

The monetary authority (in the U.S., the Fed) produces a sequence of money stocks $\{M_t\}_{t=0}^{\infty}$. These then feed through the quantity theory of money relation (18.2) to produce a sequence of inflation rates. A *monetary policy* will be a choice for the growth rate of money. If the stock of debt is growing, eventually the bond ceiling will be reached and the Fed will no longer be able to pick an inflation rate, it will be forced to provide enough seigniorage revenue to cover the government's reported deficit. We call this the *catastrophe.* The catastrophe happens at date T. The catastrophe date T is itself as a function of choice made by the government.

Given money supply growth, the gross inflation rate in period t is:

$$\frac{P_t}{P_{t-1}} = \frac{M_t}{N_t}\frac{N_{t-1}}{M_{t-1}} = \frac{1}{1+n}\frac{M_t}{M_{t-1}}.$$

The net inflation rate is defined as $P_t/P_{t-1} - 1$. For simplicity, assume (with Sargent and Wallace) that the Fed picks a constant growth rate for money, θ, in the periods before the catastrophe. Thus:

$$\frac{M_t}{M_{t-1}} = 1 + \theta, \text{ for all } t = 0, 1, \ldots, T.$$

This implies that inflation is:

$$\frac{P_t}{P_{t-1}} = \frac{1+\theta}{1+n}, \text{ for all } t = 0, 1, \ldots, T.$$

For $\theta > n$, the net inflation rate will be strictly positive. If the Fed dislikes inflation, it will seek to minimize the growth rate of money M_t/M_{t-1} by picking a low θ. Such a policy will decrease seigniorage revenue in the short run (until period T), forcing the fiscal authority to rely more on bond finance of deficits, bringing closer the catastrophe date T at which $b^g_T = \bar{b}$ and no more bonds may be sold. From period T on, the money supply expands to produce enough revenue to satisfy the government budget constraint.

Analysis

Our goal is to determine the time path of per-capita bond holdings b^g_t and to determine when (if ever) the limit of $\bar{b}$ is reached. Table (18.1) lists all of the variables and their meanings. In addition, let's list again all of the equations we know about this model:

(Gov. Budget Constraint) $$D_t = B^g_t - (1+r)B^g_{t-1} + \left(\frac{M_t - M_{t-1}}{P_t}\right).$$

(Fiscal Policy Rule) $$D_t/N_t = d, \text{ for all } t = 0, 1, \ldots, \infty.$$

(Monetary Policy Rule) $$M_t = (1+\theta)M_{t-1}, \ t = 0, 1, \ldots, T.$$

(Population Growth Rate) $$N_t = (1+n)N_{t-1}.$$

(Quantity Theory of Money) $$P_t = M_t/N_t.$$

Begin by dividing the government budget constraint (18.1) by N_t on both sides to produce:

$$\frac{D_t}{N_t} = \frac{B_t^g}{N_t} - (1+r)\frac{B_{t-1}^g}{N_t} + \frac{1}{N_t}\frac{M_t - M_{t-1}}{P_t}, \text{ for all } t = 0, 1, \dots, \infty.$$

Now, we use the fact that $1/P_t = N_t/M_t$ to write this as:

$$\begin{aligned} d &= b_t^g - (1+r)\frac{B_{t-1}^g}{N_{t-1}}\frac{N_{t-1}}{N_t} + \frac{1}{N_t}(M_t - M_{t-1})\frac{N_t}{M_t} \\ &= b_t^g - \frac{1+r}{1+n}b_{t-1}^g + \frac{M_t - M_{t-1}}{M_t} \\ &= b_t^g - \frac{1+r}{1+n}b_{t-1}^g + \left(1 - \frac{M_{t-1}}{M_t}\right), \text{ for all } t = 0, 1, \dots, \infty. \end{aligned}$$

Solving for b_t^g yields:

$$b_t^g = \frac{1+r}{1+n}b_{t-1}^g + d - \left(1 - \frac{M_{t-1}}{M_t}\right), \text{ for all } t = 0, 1, \dots, \infty. \tag{18.3}$$

Notice that the evolution of per-capita borrowing b_t^g determined in equation (18.3) holds in all periods, including those after the catastrophe period T. Before period T the monetary policy specifies a growth rate of money, $M_t/M_{t-1} = 1 + \theta$, so seigniorage is constant and potentially low. The remaining borrowing is done by issuing bonds. After the catastrophe date T, monetary policy must produce enough seigniorage revenue to completely meet the government's borrowing needs, and per-capita bonds are constant at $b_T^g = b_{T+1}^g = \cdots = \bar{b}$.

After the catastrophe the evolution of the money supply is determined by the post catastrophe government budget constraint, so we replace b_t^g with $\bar{b}$:

$$\bar{b} = \frac{1+r}{1+n}\bar{b} + d - \left(1 - \frac{M_{t-1}}{M_t}\right), \text{ for all } t \geq T+1.$$

We manipulate this equation to solve for the growth rate of money:

$$\frac{M_t}{M_{t-1}} = \frac{1}{1 - d - \left(\frac{r-n}{1+n}\right)\bar{b}}, \text{ for all } t \geq T+1. \tag{18.4}$$

Notice that after period T, money supply growth is increasing in the terms d and $\bar{b}$. Not only does the Fed have to pay for the deficit d entirely out of seigniorage, it also has to pay the carrying costs on the public debt $\bar{b}$.

Thus the money stock must evolve as:

$$\frac{M_t}{M_{t-1}} = \begin{cases} 1+\theta, & t = 1, \dots, T \\ \left\{1 - d - \bar{b}\left(\frac{1+r}{1+n} - 1\right)\right\}^{-1}, & t = T+1, T+2, \dots, \infty. \end{cases} \tag{18.5}$$

Equation (18.5) gives us the evolution of the money supply in all periods, including those after the catastrophe. Notice that the money supply growth rate after T is not affected by the value of T. In other words, after the catastrophe hits, the inflation rate will be the same, no matter when it hit.

How much seigniorage revenue does the government raise, given θ, each period prior to the catastrophe? That is, what happens when we substitute in the Fed's monetary policy into equation (18.3)? From equation (18.3):

$$b_t^g = \frac{1+r}{1+n} b_{t-1}^g + d - \left(1 - \frac{M_{t-1}}{M_t}\right), \text{ for all } t = 1, 2, \dots, \infty.$$

But in the periods before the catastrophe, money growth is simply θ, so:

$$b_t^g = \frac{1+r}{1+n} b_{t-1}^g + d - \left(1 - \frac{1}{1+\theta}\right) = \frac{1+r}{1+n} b_{t-1}^g + d - \frac{\theta}{1+\theta}, \text{ for all: } t = 1, 2, \dots, T. \tag{18.6}$$

Notice this interesting result: Before period T, the government takes as seigniorage a fraction $\theta/(1+\theta)$ of GDP. Any remaining portion of the per-capita deficit d must be raised by net new bonds.

Finally, let's calculate b_t^g without reference to b_{t-1}^g. We can do this with recursive substitution from equation (18.6), using the assumption that $b_0^g = 0$:

$$\begin{aligned}
b_1^g &= \frac{1+r}{1+n} b_0^g + d - \frac{\theta}{1+\theta} = d - \frac{\theta}{1+\theta}. \\
b_2^g &= \frac{1+r}{1+n} b_1^g + \left(d - \frac{\theta}{1+\theta}\right) = \left(1 + \frac{1+r}{1+n}\right)\left(d - \frac{\theta}{1+\theta}\right). \\
b_3^g &= \frac{1+r}{1+n} b_2^g + \left(d - \frac{\theta}{1+\theta}\right) = \left[1 + \frac{1+r}{1+n} + \left(\frac{1+r}{1+n}\right)^2\right]\left(d - \frac{\theta}{1+\theta}\right).
\end{aligned}$$

And so on. The pattern should be clear from these first terms. In general:

$$b_t^g = \left[d - \frac{\theta}{1+\theta}\right] \sum_{i=1}^{t} \left(\frac{1+r}{1+n}\right)^{i-1}, \text{ for all } t = 0, \dots, T. \tag{18.7}$$

Recall that $r > n$, hence the summation term is explosive.

Equation (18.7) neatly captures the Fed's dilemma in this model. By setting a low value for θ, the Fed trades low inflation today for an earlier onset of the hyperinflationary catastrophe. On the other hand, by choosing a relatively high value for θ the Fed suffers high inflation today but staves off the catastrophe point. Indeed, if:

$$\frac{\theta}{1+\theta} \geq d,$$

then there will be no catastrophe.

Determining The Catastrophe Date T

Given the time path for debt in equation (18.7), we can determine roughly in which period T the catastrophe hits. I say "roughly" because to keep the algebra neat we are going to assume that, at the monetary policy θ, end-of-period T debt b_T^g is perfectly equal to $\bar{b}$. You can see that it is easy to imagine cases in which b_T^g is slightly less than $\bar{b}$, in which case in period $T+1$ a residual amount of borrowing is allowed. However if T is large, this effect is unimportant. Thus at the end of period T:

$$\left[d - \frac{\theta}{1+\theta}\right] \sum_{i=1}^{T} \left(\frac{1+r}{1+n}\right)^{i-1} = \bar{b}.$$

For notational convenience, let $\gamma \equiv (1+r)/(1+n)$. Thus:

$$\sum_{j=0}^{T-1} \gamma^j = \frac{\bar{b}}{d - \frac{\theta}{1+\theta}}.$$

Recall that the sum on the left hand side of this equation is equal to $(1-\gamma^T)/(1-\gamma)$. Thus:

$$\frac{1-\gamma^T}{1-\gamma} = \frac{\bar{b}}{d - \frac{\theta}{1+\theta}} \equiv J,$$

where I have introduced J to keep the notation down. Manipulation produces:

$$\gamma^T = 1 - (1-\gamma)J.$$

Taking logarithms of both sides produces:

$$\begin{aligned} T\ln(\gamma) &= \ln\left(1 - (1-\gamma)J\right), \text{ so:} \\ T(\theta, \bar{b}) &= \frac{\ln\left(1-(1-\gamma)J\right)}{\ln(\gamma)}, \text{ where:} \\ \gamma &= \frac{1+r}{1+n}, \text{ and:} \\ J &= \frac{\bar{b}}{d - \frac{\theta}{1+\theta}}. \end{aligned} \tag{18.8}$$

Notice that T is increasing in θ and $\bar{b}$. Indeed, for T to be finite, we must have:

$$\frac{\theta}{1+\theta} < d,$$

so that the government must resort to bond financing.

Some Examples

In Figure (18.1) we present the time path of debt, b_t^g, under two different values of θ, $\theta_1 = 0.03$ and $\theta_2 = 0.10$. In this model $n = 0.02, r = 0.05, d = 0.10$ and $\bar{b} = 1.5$. That is, the government is trying to finance a persistent core deficit of 10% of GDP and the maximum value of total debt is 150% of GDP. The government does not have to pay a very high real interest rate on its debt, but output is growing at the relatively low rate of 2% a year. With the tight monetary policy ($\theta_1 = 0.03$), the government hits the catastrophe 16 years into the policy, while with the loose monetary policy ($\theta_2 = 0.10$), the catastrophe occurs 61 years in the future.

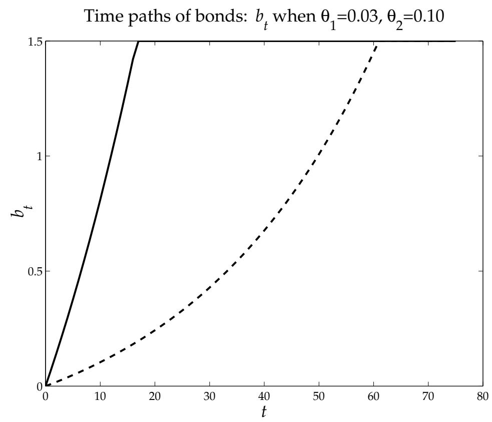

Figure 18.1: Evolution of the stock of per-capita debt holdings b_t under two monetary policies: the solid line under the tight money ($\theta = 0.03$) policy and the dotted line under the loose money ($\theta = 0.10$) policy.

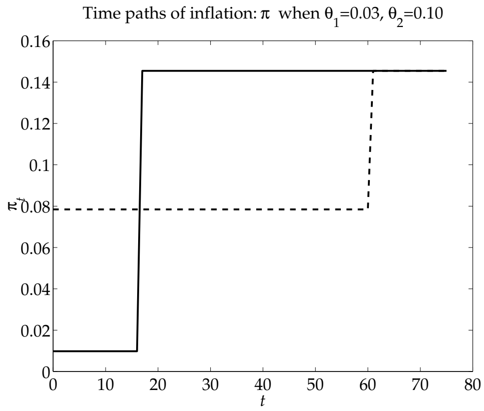

Figure 18.2: Evolution of the inflation rate π_t under two monetary policies: the solid line under the tight money ($\theta = 0.03$) policy and the dotted line under the loose money ($\theta = 0.10$) policy.

In Figure (18.2) we plot the inflation rates over time associated with the two monetary policies. Notice that the inflation rate π_t does not quite equal the growth rate of money since:

$$1 + \pi_t \equiv \frac{P_t}{P_{t-1}} = \frac{1}{1+n}\frac{M_t}{M_{t-1}}.$$

Before the catastrophe date T, inflation is constant at π_θ where:

$$1 + \pi_\theta = \frac{1+\theta}{1+n}, \text{ for all } t = 0, 1, \ldots, T., \text{ so:}$$

$$\pi_\theta = \frac{1+\theta}{1+n} - 1 = \frac{\theta - n}{1+n}.$$

Note that π_θ will not vary with the deficit d or the maximum debt load $\bar{b}$. On the other hand, the catastrophe date T and the post-catastrophe inflation rate will vary with d and $\bar{b}$. After the catastrophe, inflation $\pi_T(d, \bar{b})$ will not vary with the pre-catastrophe monetary policy θ. We can calculate $\pi_T(d, \bar{b})$ from the evolution of the money supply, equation (18.5). Thus:

$$1 + \pi_T(d, \bar{b}) = \frac{1}{1+n} \frac{1}{1 - d - \bar{b}\left(\frac{1+r}{1+n} - 1\right)}, \text{ so:}$$

$$\pi_T(d, \bar{b}) = \frac{1}{1+n} \frac{1}{1 - d - \bar{b}\left(\frac{1+r}{1+n} - 1\right)} - 1 = \frac{r\bar{b} + d - n(\bar{b} + 1 - d)}{(1-d)(1+n) - \bar{b}(r-n)}. \tag{18.9}$$

The tight monetary policy is associated with very low inflation initially, $\pi_{\theta_1} = 0.0098$ but, as noted above, the catastrophe happens relatively early. The loose monetary policy is associated with a relatively high inflation rate initially, $\pi_{\theta_2} = 0.0784$ but the catastrophe is staved off for over 60 years. After the catastrophe the inflation rate is $\pi_T = 0.1455$, or about twice the rate with the loose monetary policy.

Application: Optimal Inflationary Policies

In this section we consider the trade-off between two monetary policies: (1) A policy of high inflation in which the catastrophe never occurs and (2) A low-inflation policy which brings forward the catastrophe date.

Notice from equation (18.7) that if the government sets $\theta = \theta^\star$, where:

$$\theta^\star \equiv \frac{d}{1-d},$$

then each period's seigniorage revenue is:

$$\frac{\theta^\star}{1+\theta^\star} = d.$$

That is, with the money supply growth rule set to $\theta^\star$ as defined above, the government raises enough seigniorage revenue to completely finance the real deficit each period. As a result the government never resorts to bond finance, so $b_t^g = 0$ all $t = 0, 1, 2, \dots, \infty$ and the catastrophe never happens. When $\theta = \theta^\star$ inflation satisfies:

$$\pi_{\theta^\star} = \frac{\theta^\star - n}{1+n} = \frac{\alpha}{1-\alpha},$$

where $\alpha = d - n + nd$. Notice that if $d = n/(1+n)$ then $\pi_{\theta^\star} = 0$. That is, the government can pay for the real deficit entirely with seigniorage revenue and have zero inflation.

On the other hand, for any monetary policy $\theta < \theta^\star$, the government must resort to persistent debt financing and eventually face the catastrophe. We know from equation (18.9) above that after the catastrophe, inflation is $\pi_T(d, \bar{b})$. By examination, we see that:

$$\pi_T(d, \bar{b}) > \pi_{\theta^*}.$$

Intuitively, by waiting until period T to begin financing excess government spending by printing money the monetary authority has allowed the fiscal authority to borrow up to its limit. The Fed then has to repay creditors out of seigniorage as well.

If the Fed dislikes inflation, it has an unpleasant choice: suffer inflation of π_{θ^*} now or π_T at some future date T. As you can see, the Fed's choice of which policy to pursue depends in large part on how T varies with θ.

Variable	Definition
G_t	Real government spending at t
T_t	Real government tax revenues at t
D_t	Real government core deficit at t
Y_t	GDP at t, $Y_t = N_t$
N_t	Population at t
n	Constant population growth rate
r	Constant real net return on debt, $r > n$
M_t	end-of-period stock of money at t
P_t	exchange rate of money for goods at t
B_t^g	real par value of outstanding end-of-period debt
b_t^g	per-capita debt, $b_t^g = B_t^g/N_t$
$\bar{b}$	maximum possible value of b_t^g
B_{-1}^g	initial stock of debt $B_{-1}^g = 0$
T	"catastrophe date"–when $b_T^g = \bar{b}$
θ	Money supply growth rate before catastrophe
d	constant per-capita deficit (fiscal policy)

Table 18.1: Notation for Chapter 18. Note that, with the assumption that $Y_t = N_t$, variables denoted as per-capita are also expressed as fractions of GDP.

18.2 The Ends of Four Big Inflations

The most dramatic evidence of the validity of the Sargent-Wallace argument comes from the post-WWI hyperinflations in Germany and the successor states to the Austro-Hungarian Empire in a paper by Sargent, "The Ends of Four Big Inflations". What makes that case so special is that, not only was there a deficit-driven hyperinflation, once the fiscal authorities had made credible commitments to back all government debt by tax revenues, the inflation stopped (even though the printing presses were still running). These histories are valuable also because the problems facing the four nations in question bear much in common with those facing some of the successor states of the old Soviet empire.

The post-war central European inflations of 1919-1924 were a new, and deeply unpleasant, experience for its citizens. It is a commonplace to ascribe modern Germany's strong commitment to low inflation to a national horror of repeating those days. Yet it was not the abstract experience of seeing prices (and wages) climb to 10^{12} times their pre-war level that was so traumatic, nor was it the mild "shoeleather cost" studied in Chapters 4 and 8. As a result of the inflation, there were tremendous social dislocations as creditors were impoverished, as enterprises failed, as speculation flourished and as households hoarded illiquid assets rather than trading them for a currency whose value was essentially unknown. These were new phenomena at the time, but unfortunately since then they have been consistent hallmarks of monetary crises to the present day.

Sargent identifies four characteristics common to the hyperinflationary experiences in Poland, Hungary, Austria and Germany:

1. All four countries ran very large budget deficits.
2. All four countries took very similar, very dramatic, monetary and fiscal steps to end the hyperinflations.
3. In all cases, the inflation stopped very quickly.
4. After the inflationary episodes, there was a large and persistent rise in the level of "high-powered" money.

Governments ran deficits because, in the aftermath of the war, they made payments to the large numbers of unemployed workers, because state monopolies (such as railroads) kept prices artificially low and lost money, because governments subsidized basic necessities such as food and housing, and, in certain cases, because they had been ordered to pay war reparations of unknown amounts.

Sargent draws a clear distinction between government *actions* and government *regimes.* An action takes the form of a one-period decision of the government (cutting the subsidy on heating oil for one month, for example), with no credible assurance that the action will be repeated. In contrast a regime is a credible commitment to a sequence of actions, for example selling off the state railroad or making the central bank independent.

The solution to the hyperinflations, in all cases, was a switch in regime: governments abandoned deficits and seigniorage financing in favor of balanced budgets and independent central banks. In many cases, at least part of the credibility of the new regimes derived from international obligations. For example, in August of 1922 Austria signed agreements with the League of Nations binding her to fiscal balance and monetary stability.

In Germany, where the inflation was most dramatic, the largest single fiscal liability was the bill for war reparations. In the original treaty negotiations at Versailles, the Great Powers had been unable to fix a firm value on Germany's war reparations. In theory all of Germany was mortgaged for reparations, and indeed, in 1923 France occupied the Ruhr to drive home this point. In October of that year Germany issued a new currency, the rentenmark, whose initial value was 10^{12} reichsmarks.

Yet in 1924 the catastrophic German inflation stopped. Sargent reports several deliberate, permanent actions that constituted a regime shift. Among these, the government fired 25% of its workforce and cut employment in the state railroad system by about 180,000. Germany also negotiated a fixed, reasonable, value for its reparations bill with the treaty powers.

In all of these inflations, at some point the central banks were called upon to purchase almost all of the net new national debt issues. A common reform was to prohibit the central bank from purchasing government debt. This was a statutory commitment to fiscal discipline, and a good example of the difference between regimes and actions.

Once households were assured that the hyperinflationary regime was over, their holdings of currency rebounded remarkably. Thus even after the inflations had ended, governments continued to issue large quantities of new base money. This money was absorbed by households which had economized dramatically on their currency holdings during the hyperinflation.

The parallel to present-day countries such as Ukraine and Russia is clear. These too are new states without a history to guide investors, with bloated public sectors and inefficient systems of tax collection. In contrast with the earlier examples, they are committed, as much as possible, to fiscal discipline, although in some cases this required defaulting on some of the government's obligations (for example, Russian government employees must often wait months for paychecks). International organizations such as the IMF and foreign governments, just as in the 1920s, have acted as commitment devices to prevent the Russian government from using the printing press to meet its obligations. However, until either government spending obligations diminish or tax collections increase, there will be a persistent possibility of hyperinflation, with its attendant social dislocations.

Exercises

Exercise 18.1 (Easy)
The answers to these exercises can be found in Friedman and Schwartz (1963) *A Monetary History of the United States*, Chapter 7, entitled, "The Great Contraction"; the Barro textbook Chapters 4 and 18 but especially Chapters 7, 8, and 17; the article by Richard D. Porter, "The Location of U.S. Currency: How Much is Abroad?"; and Sargent *Rational Expectations and Inflation* (either the 1st or 2nd edition), Chapter 3, entitled "The Ends of Four Big Inflations",

and Chapter 5, entitled, "Some Unpleasant Monetarist Arithmetic".

1. What was the path of the money stock in the U.S. from January 1929 to March 1933? How did household's holdings of currency change over the same period?
2. What was the path of real income in the U.S. from January 1929 to March 1933? How did prices change over the same period?
3. In the period 1948-1991 have American real interest rates ever been negative? In the same period, has the U.S. inflation rate ever been negative? If so, when?
4. What is the evidence that inflation, in Milton Friedman's words, "is always and everywhere a monetary phenomenon"? In the long run? In the short run?
5. If you take the population of the U.S. to be 260 millions, roughly how many dollars of *currency* were in circulation for every U.S. citizen at the end of 1995? How much currency are you carrying right now? How do you account for this discrepancy?
6. Why do people hold currency and keep part of their wealth in low interest bearing accounts (like the Hyde Park Bank's zero interest checking account)?
7. Explain how a rational expectations view of agent's behavior (as defined by Sargent) can explain why inflation *seems* to have momentum, while in fact it does not.
8. What is *seignorage*? How much money did the U.S. raise via seignorage in 1991?
9. What is the Quantity Theory of Money? Explain the sense in which it is "just" an accounting identity.
10. What is a gold standard? True or false: Under a gold standard the quantity of money is fixed.

Exercise 18.2 (Easy)
Evaluate this statement: Government austerity programs cause civil unrest.

Exercise 18.3 (Moderate)
In each period t the government raises real tax revenue of T_t and spends (in real terms) G_t. Let $D_t \equiv G_t - T_t$ be the *real deficit* at time t. At the suggestion of a revered elder whose initials are M.F., the government is allowed to finance this deficit only by issuing fiat currency and obtaining the seignorage revenue. The government's budget constraint is thus:

$$D_t = \frac{M_t - M_{t-1}}{P_t},$$

where D_t is the real government deficit at t, M_t is the stock of money at t and P_t is the price level at t. Prices are related to the money supply by the Quantity Theory of Money relation with a constant velocity $v = 1$:

$$P_t Y_t = M_t.$$

Output Y_t satisfies $Y_t = N_t$, where N_t is population at t, and evolves according to:

$$N_t = (1 + n)N_{t-1},$$

with $N_0 = 1$. The government runs a constant per capita real deficit of d, so $D_t = dN_t$ for all t. Answer the following questions:

1. How must M_t evolve given M_{t-1} and d?
2. For what value of d is the inflation rate zero? That is, for what value of d will $P_t = P_{t-1}$?
3. A reasonable estimate for n is about 0.03. At this value, how large a deficit, expressed as a fraction of GDP, can the government cover by printing money and still not cause inflation?
4. Assume $d = 0$. What happens to prices?

Exercise 18.4 (Fun)
Through a map-making error in 1992 the Absolutely Autonomous People's Republic of Kolyastan (hereafter known as Kolyastan) was created out of the more rubbishy bits of neighboring successor states to the Soviet Union. The Kolyastani central bank is run part-time by a popular local weatherman on the state-run television station. The market for Kolyastan's chief export, really really big statues of Lenin, seems to have collapsed. Most of its citizens continue to work in the enormous state-run Lenin Memorial Lenin Memorial factory, which is currently producing no revenue at all. The government subsidizes consumption of bread and *kirghiz light* (the local liquor) by paying merchants to keep their prices artificially low. The Kolyastani currency, the neoruble, is made up of old Soviet rubles with the top left corner cut off. Inflation is currently running at 400% per month. Although the Kolyastani government claims to be financing most of its big budget deficits through bond sales, most of these bond sales, it turns out, are to the central bank. In desperation the Kolyastani government have turned to you, a University of Chicago undergraduate, for economic advice. Briefly outline your plan for Kolyastan's recovery. Be specific. How can the Kolyastani people be certain that the reforms proposed by the government will be maintained after you graduate?

Chapter 19

Optimal Monetary Policy

As we have discussed, expansionary monetary policies include decreases in the Fed funds rate and unexpected growth in the money supply. In the U.S., such expansionary monetary policies have tended to produce real expansions in output and increases in inflation. Conversely, contractionary monetary policies have tended to produce real contractions in output and decreases in inflation. In Chapter 18 Barro claims that these effects have been quite moderate, but recent empirical work lends support to the opposite view, that monetary shocks can have large effects on real variables in the short run.

Everyone agrees that expansionary monetary policies tend to lead to increases in inflation, while contractionary policies produce decreases in inflation. At this broad level, monetary policy would appear to be a matter of trading off inflation and output. Since unemployment tends to decrease as output increases, this is often cast a choice between inflation and unemployment. The empirical relationship between the two is called the *Phillips curve.*

In the U.S., as in most countries, monetary policy is under the control of the government. This immediately raises the question of how best to conduct monetary policy. As we shall see, this not so much a question of when and how to time expansions and contractions of the money supply, as economists used to think, as it is a question of what the private sector predicts the government will do and how the government can influence those predictions.

Before we can think fruitfully about monetary policy, we will have to have a reasonable model of how monetary shocks can influence the real economy. Our model will be a simplification of the seminal paper by Robert E. Lucas, Jr, "Expectations and the Neutrality of Money". In that model, the private sector is divided into different industries (called "islands") which observe only the price for their own product. This price is made up of a general price level (unobserved) and an industry-specific shock (also unobserved). The private sector has some forecast about inflation (never mind for the moment its origin) and uses this to derive an estimate of the industry specific shock it faces. If the estimated shock is high, the private sector increases production. If it is low, the private sector decreases production. The government chooses an inflation rate. An unexpected monetary expansion will produce a temporary increase in output. Thus Lucas's model highlights the role of *expectations* in the conduct of monetary policy.

In Lucas's model, only unanticipated changes in the price level have real effects. If a monetary expansion is completely expected, it has no real effects. This points to something quite important in the real conduct of monetary policy: only surprises matter. Moreover, the private sector does not enjoy being surprised, even if the monetary surprise produced a temporary boom. An older tradition in macroeconomics holds that governments should try to manipulate the money supply to cushion supply and demand disruptions. The central lesson from Lucas's research is that governments should instead strive to minimize the uncertainty surrounding monetary policy.

We then move away from the specific form of the Phillips curve derived from Lucas's model and start using a simple generalization in which inflation, inflationary expectations and unemployment are all related by a very simple formula. The government will have some preferences (and thus indifference curves) over unemployment and inflation (both will be bad), and monetary policy, if we ignore how expectations are formed, can be seen as a simple choice of unemployment and inflation.

Once we begin modeling the formation of expectations, we will see that the ability of the government to commit credibly to a particular inflationary path is critical. We will model explicitly a two-person game between the private sector and the government. With a so-called commitment device, the government will be able to play the *Ramsey strategy* and realize the *Ramsey outcome.* Recall the Ramsey optimal tax problem from Chapter 14. In that chapter we assumed that the government could commit to a particular tax sequence, hence the term "Ramsey". We did not consider what would happen if the government could not commit to a particular tax sequence. In this chapter we will see that without a commitment device, the government and the private sector will play *Nash strategies* and achieve the *Nash outcome.* The fundamental result of this chapter is that Ramsey is better than Nash. Both the government and the private sector are better off in the Ramsey outcome than in the Nash outcome. Indeed, under certain circumstances, the Nash outcome involves (temporarily) high inflation and high unemployment, the so-called "stagflationary" episode of the 1970s. At the time, stagflation was blamed on an oil price shock. We have to reconsider, and say that possibly it was the result of a lack of credible commitment by the government.

The theory in this chapter will give us an explanation for the "pain" associated with fighting inflation. There is a powerful maintained assumption in the media that policies that are anti-inflationary require some sacrifice of real output. As we shall see, when the private sector has formed strong expectations about continued high inflation, confounding those expectations with sudden, unexpected, low inflation can have a severe cost in terms of real output. This is not a reason to oppose anti-inflationary policies, it is a reason to campaign for a credible commitment to low inflation.

Finally, it is worth noting that, in this chapter, we will ignore the government budget constraint. In Chapter 18 we were very concerned about the relationship between persistent government budget deficits and inflation. In this chapter we will assume the government budget is more or less in balance, and that the government does not particularly need the seignorage revenue generated by high inflation. This is a safe assumption when thinking about inflation in the U.S. In thinking about inflation across different countries, though, the analysis of Chapter 18 is probably more appropriate in countries, like Brazil, that experience persistent inflation and large budget deficits. The "pain" of fighting inflation in those countries is the pain of raising direct taxes and decreasing government spending.

Both this chapter and Chapter 18 highlight the importance of credible government policies. In Chapter 18, to stop hyperinflations the government had to credibly commit to balancing its fiscal books. In this chapter, to prevent milder inflations, the government will have to credibly commit to keep its hands off of the monetary spigot. In both cases there is a role for international institutions as commitment devices.

19.1 The Model of Lucas (1972)

In this section we consider a simplified version of the important model of Lucas. We are going to get a relationship between the anticipated price level, the actual price level and something that looks like unemployment. We will use this relationship to argue for a particular functional form for the Phillips curve. We will not derive precisely a Phillips curve since our model is going to be static, to keep the exposition simple. The dynamic generalization is very elegant, and the interested reader is referred directly to the Lucas paper.

This model turns on the decisions made by many separated industries in the private sector. These industries cannot communicate with one another about prices. They will hire labor according to their estimate of the true state of demand for their product.

Let Q_i be output in industry i. Assume that all industries use only one input, labor. Let L_i be the number of workers hired in industry i. Assume that all industries have the common production function:

$$Q_i = L_i^{\alpha},$$

where the technology parameter α satisfies $0 < \alpha < 1$. Assume that all workers are paid the common wage of unity for their unit of labor supplied. To produce an output Q_i therefore requires labor input (and total costs) of $L_i^{1/\alpha}$. Thus the cost function in industry i is:

$$\text{Total Cost}(Q_i) = L_i^{\frac{1}{\alpha}}.$$

In industry i there will be a price P_i for that industry's output. It is known that this price is made up of two parts: a general price level P, common to all industries, and a shock term Z_i specific to industry i. These terms are related by the price equation:

$$P_i = PZ_i. \tag{19.1}$$

The shock term Z_i gives the *real price* of output in industry i. The general price level P will not be revealed until the end of the period, since the industries are on islands and cannot communicate during production.

All private-sector industries begin the period with a common *forecast* of P, which we denote by P^e. Thus an industry i's best estimate of its real price Z_i is:

$$Z_i^e = \frac{P_i}{P^e}. \tag{19.2}$$

Recall that industry i only observes P_i.

Equilibrium in industry i, assuming that it is competitive, requires that marginal cost equal estimated real price Z_i^e. Since we know the total cost curve, marginal cost must just be its derivative with respect to output Q_i. That is, equilibrium requires:

$$\frac{1}{\alpha} Q_i^{\frac{1}{\alpha}-1} = Z_i^e.$$

We can solve this to produce the equilibrium demand for labor conditional on the estimated shock Z_i^e:

$$L_i = (\alpha Z_i^e)^{\frac{1}{1-\alpha}}. \tag{19.3}$$

As expected, industries will demand more labor if they estimate that demand for their product is unusually strong (if Z_i^e is large).

The estimated shock Z_i^e is comprised of two parts: the known estimate of the price level P^e and industry-specific price level P_i, related by equation (19.2). Thus we can substitute from that equation into equation (19.3) to find the industry-specific demand for labor conditional on P^e and P_i:

$$L_i = \left(\alpha \frac{P_i}{P^e}\right)^{\frac{1}{1-\alpha}}.$$

Now we take logarithms of both sides. From now on, let lower-case variables denote logarithms. Thus $\ell_i = \ln(L_i)$ is given by:

$$\ell_i = \frac{1}{1-\alpha} \ln\left(\alpha \frac{P_i}{P^e}\right) = \frac{1}{1-\alpha}\ln(\alpha) + \frac{1}{1-\alpha}(p_i - p^e),$$

substitute $z_i + p$ for p_i from equation (19.1) above, and let $A = [1/(1-\alpha)]\ln(\alpha)$ to produce:

$$\ell_i = A + \frac{1}{1-\alpha}(z_i + p - p^e). \tag{19.4}$$

Equation (19.4) captures the log of labor demand as a function of the (log of the) shock, the common price level p and the common price forecast p^e.

Define u to be the "not employed rate" (not quite the unemployment rate, but something close).[1] If N is the total workforce, and $n = \ln(N)$, then define u as:

$$u = n - \sum_i \ell_i.$$

Assume for a moment that there are only two industries. Now:

$$u = n - 2A + \frac{2}{1-\alpha}(p^e - p) - \frac{2}{1-\alpha}(z_1 + z_2).$$

Define further:

$$u^\star = n - 2A,$$

where $u^\star$ is something like the natural rate of not-employment,

$$\begin{aligned} \varepsilon &= -\frac{2}{1-\alpha}(z_1 + z_2), \text{ and:} \\ \gamma &= \frac{2}{1-\alpha}. \end{aligned}$$

Now we can write the aggregate not-employment rate as:

$$u = u^\star + \gamma(p^e - p) + \varepsilon. \tag{19.5}$$

We will use some version of this equation throughout this chapter.

From the point of view of the government, the common price level p is a control variable. The government picks a level for p with monetary policy. Notice what equation (19.5) says about the relationship of unemployment (or not-employment), the price level and the forecast price level: unemployment is decreasing in the price level p but increasing in the forecast price level p^e. From the point of view of private industry, if the actual price level exceeds the forecast price level, $p > p^e$, the industry has produced too much and suffers losses as a result. From the point of view of the government, if $p > p^e$, it can stimulate a one-period boom in which unemployment is below its natural level.

19.2 Monetary Policy and the Phillips Curve

For the rest of this chapter we will be using a modified version of equation (19.5). Assume that:

$$u = u^\star + \gamma(\pi^e - \pi). \tag{19.6}$$

[1] The unemployment rate is $1 - (1/N)\sum_i L_i$ which doesn't translate well into logarithms.

Here u is the unemployment rate, $u^\star$ is the "natural rate" of unemployment, π^e is the *expected inflation rate* and π is the *actual inflation rate.* The natural rate of unemployment is the level of unemployment when inflation is perfectly anticipated, so no industries are fooled into thinking that relative demand is unusually high or low. The slope of this Phillips curve is $-\gamma$ where we assume $\gamma > 0$ (monetary expansions reduce unemployment). If we think that there is uncertainty about the state of the real economy, we can add a mean zero shock term, ε, to produce:

$$u = u^\star + \gamma(\pi^e - \pi) + \varepsilon.$$

For the most part we will assume that the monetary authority knows the state of the real economy with certainty.

In Figure (19.1) we plot Phillips curves with two different values of expected inflation π^e, a low value in which the expected inflation rate is zero, and a high value, in which the expected inflation rate is 8.3%. The dotted line gives the natural rate of unemployment (here $u^\star$ = 5%), and $\gamma = 0.3$. Notice that when inflationary expectations are high, to achieve any given unemployment rate requires a higher inflation rate, and to achieve zero inflation requires an unemployment rate well above the natural rate.

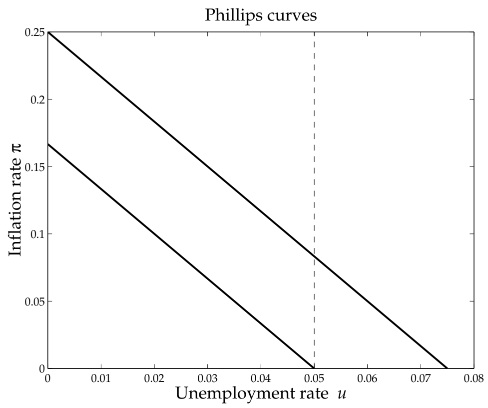

Figure 19.1: Phillips curves under two different expectations about inflation. The bottom curve assumes $\pi^e = 0$ and the top curve assumes $\pi^e = 0.0833$. The dotted line gives the natural rate of unemployment.

Monetary Policy with Fixed Expectations

Assume that the government (or which ever arm of the government controls monetary policy) has a utility function over unemployment and inflation of $V^g(u, \pi)$ given by:

$$V^g(u, \pi) = -u^2 - \pi^2.$$

That is, the government dislikes unemployment and inflation equally. We will assume this form for V for the rest of the chapter, so it's worth mentioning that the Federal Reserve Board is, by law, supposed to balance the twin goals of full employment and price stability. Thus this utility function seems to be written in law.

If we assume that π^e is given exogenously and fixed, we can substitute the Phillips curve in equation (19.6) into the government's utility function above to produce a maximization problem. Thus if the private sector has fixed expectations about the inflation rate given by π^e, then the government's optimal choice of inflation π is given by:

$$\max_{\pi} \left\{ -[u^\star + \gamma(\pi^e - \pi)]^2 - \pi^2 \right\}.$$

The first-order condition with respect to inflation π is:

$$2\gamma[u^\star + \gamma(\pi^e - \pi)] - 2\pi = 0.$$

We can solve this for π to get the optimal inflation choice when expected inflation is fixed at π^e and the natural rate is $u^\star$ (call it $\pi^\star(\pi^e)$):

$$\pi^\star(\pi^e) = \frac{\gamma}{1 + \gamma^2}(u^\star + \gamma\pi^e).$$

We can plug $\pi^\star(\pi^e)$ into the Phillips curve in equation (19.6) to produce the associated unemployment rate, $u_0(\pi^e)$:

$$u_0(\pi^e) = \frac{1}{1 + \gamma^2}u^\star + \frac{\gamma}{1 + \gamma^2}\pi^e.$$

Notice that if π^e is "small" that $u_0(\pi^e)$ will lie below $u^\star$. The government trades off some inflation for a lower unemployment rate.

We plot $\pi^\star(\pi^e)$ in Figure (19.2) below. Notice that for low values of expected inflation, π^e, the government chooses inflation rates above expectations and for high values of π^e, the government chooses inflation rates below expectations. At one unique expected inflation rate, the government's best response is to choose an actual inflation rate exactly equal to the expected inflation rate. This will play a special role, as we shall see.

Two Stories About Inflationary Expectations

We are not yet ready to discuss the strategic interactions between the private sector and the government that determine inflationary expectations. However, we can study the outcomes under two different stories about inflationary expectations. These will help us to

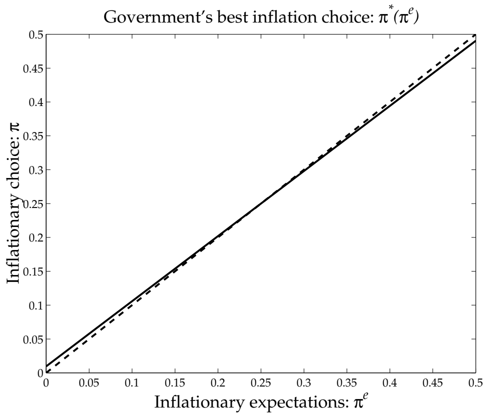

Figure 19.2: Government's optimal choice of inflation $\pi^\star$ as a function of different expectations of the private sector (solid) and the 45-degree line (dotted). Where the two lines cross is the inflationary expectation at which the government's best response is to meet expected inflation. In other words: the Nash inflation level.

think about the government's problem. First, we will assume that expectations are fixed, but that the private sector knows the government's maximization problem. If this is the case, then the private sector will set expectations to a unique value such that the government chooses to set inflation at exactly the same value the private sector anticipated. Second, we will assume that expected inflation exactly equals actual inflation in all cases. The private sector has a crystal ball (or a spy) which informs it precisely of the government's inflationary plan, no matter what the government picks.

Imagine for a moment that the private sector understands the government's maximization problem and correctly anticipates inflation. That is, assume that inflationary expectations satisfy:

$$\pi^e = \pi^\star(\pi^e).$$

From Figure (19.2) below, we see that there is exactly one such expected inflation rate. Expanding produces:

$$\pi^e = \frac{\gamma}{1+\gamma^2}(u^\star + \gamma\pi^e).$$

We can solve for this special value of π^e, call it π_1, to get:

$$\pi_1 = \gamma u^\star, \tag{19.7}$$

where π_1 is unique inflation rate such that when expectations satisfy $\pi^e = \pi_1$, the government's inflation target is also π_1. The associated unemployment rate is:

$$u_1 = u^\star,$$

since $\pi^e = \pi$. Thus the unemployment rate is at the natural rate $u^\star$ and inflation is relatively high at π_1. This will be the *Nash equilibrium* in inflation (as we shall see below).

Now imagine that the government is forced by law to correctly announce its inflation target each period. The private sector anticipates this and sets $\pi^e = \pi$. Thus the Phillips curve in equation (19.6) becomes:

$$u = u^\star + \gamma(\pi - \pi) = u^\star.$$

In other words, inflation does not affect output. If this is the case, the government chooses an inflation rate of zero (since inflation is costly and now provides no benefit), and the unemployment rate again goes to the natural rate. This will turn out to be the *Ramsey equilibrium* as we shall see below.

Contrast the Ramsey and the Nash equilibria. Both produced the natural rate of unemployment, but the Nash equilibrium also had a high inflation rate. Thus the government and the population are better off if the government is able to announce the inflation rate and be believed. As we shall see below, unfortunately, when the private sector expects inflation to be low, there is a temptation for the government to inflate.

Ramsey Monetary Policy

This last example was the Ramsey problem. If the government can credibly commit to a particular inflation rate, the private sector responds by setting inflationary expectations to the announced inflation target. As a result, the government announces an inflation target of zero, and the result is the natural rate of unemployment. What are some commitment devices? By making the monetary authority completely independent of the fiscal authority it can be insulated from political pressure. Further, if the central banker has a reputation for being an unpleasant misanthrope who cares only about defeating inflation, the private sector can become convinced over time that in fact the central bank will set $\pi = 0$ for all time.

Indeed, one reading of the deeply unpleasant recession in the early 1980s is that the private sector had to be convinced of the new central banker's commitment to low inflation. Paul Volcker arrived as Chairman of Federal Reserve Board at a time of high inflation and high unemployment. He announced that there would be low inflation in the future. The private sector did not adjust its expectations, but Volcker followed through on his promise. The result was the unusual case in which inflationary expectations exceeded actual inflation, that is $\pi^e > \pi$. As a result, unemployment shot above its natural rate in one of the deeper recessions of the century. After two years of this treatment, the private sector adjusted its expectations, convinced that Mr. Volcker was committed to low inflation.

Other countries, without the benefit of the tradition of anti-inflationary policies of the Fed to reassure the private sector, will completely let go of the reins of monetary policy. In Hong Kong, for example, the local currency is pegged to the U.S. dollar in an arrangement known as a *currency board.* For every 7.8 Hong Kong dollars issued, one U.S. dollar must be placed on deposit, so the currency is fully backed. The Hong Kong government cannot print money. Thus the exchange rate is immutably fixed, and there can be no depreciation of the local currency against the U.S. dollar. Countries will go to great lengths to convince the private sector that they are really committed to low inflation. They have to work so hard at it, we shall see, precisely because, if expectations are low, there is always a temptation to inflate.

19.3 Optimal Monetary Policy without Commitment: The Nash Problem

In this section we will explicitly model the strategic interaction between the private sector and the government when forming inflationary expectations. We will force the government to choose from only two possible inflation levels, and the private sector to pick from only two possible inflationary expectations. The results we derive here generalize to the case in which both choose from continuous distributions.

Inflation π can only take on one of two values: $\{0, \pi_1\}$. That is, inflation can be zero or the high level we derived in equation (19.7). The private sector expects π^e which can also only take on the values $\{0, \pi_1\}$, since it wouldn't make sense for the private sector to anticipate inflation rates that the government can't pick.

There are four possible combinations of expected and actual inflation, $\{\pi^e, \pi\}$. At each one of these four combinations we will specify the payoff to the private sector and to the government. These payoffs will be known by both players. We will look for a *Nash equilibrium*, which is simply a pair of choices (one for the private sector, one for the government) such that, given the other player's choice, no player can do better.

We now consider each of the four possible combinations. Let $V^g(\pi^e, \pi)$ be the payoff to the government and $V^p(\pi^e, \pi)$ be the payoff to the private sector at each possible $\{\pi^e, \pi\}$ combination. We will assume that the private sector suffers a penalty of -1 if it does not correctly forecast the inflation rate and gets a payoff of zero otherwise (this is just a normalization). We assume that the baseline government payoff (at zero inflation and the natural rate of unemployment) is 0, and that otherwise the government dislikes inflation and unemployment. At each of the four possible outcomes, the payoffs of the two players are:

Private Sector	Government	
	$\pi = 0$	$\pi = \pi_1$
$\pi^e = 0$	$V^g = 0, V^p = 0$	$V^g = 1, V^p = -1$
$\pi^e = \pi_1$	$V^g = -1, V^p = -1$	$V^g = -0.5, V^p = 0$

Notice that the government really dislikes $\{\pi^e = \pi_1, \pi = 0\}$; this corresponds to the Volcker play of low inflation when expectations are high. The result is unemployment above the natural rate. Also, the government dislikes (but not as much) $\{\pi^e = \pi_1, \pi = \pi_1\}$; here inflation is high, but unemployment is at the natural rate. The government would prefer to be at $\{\pi^e = 0, \pi = \pi_1\}$; here inflation is unexpectedly high, so unemployment is below the natural rate.

Now let us work through these payoffs to find the Nash equilibrium. If the household plays $\pi^e = 0$, the best response of the government is to set $\pi = \pi_1$. If the household plays $\pi^e = \pi_1$, the best response of the government is to play $\pi = \pi_1$. If the government plays $\pi = 0$ the best response of the household is $\pi^e = 0$, but this is not a Nash equilibrium since, if the household does play $\pi^e = 0$, we saw that the government will want to deviate to $\pi = \pi_1$. If the government plays $\pi = \pi_1$ then the household's best response is to play $\pi^e = \pi$. Since $\pi = \pi_1$ is the government's best response to a household play of $\pi^e = \pi_1$, this is the only Nash equilibrium in this example.

The Nash equilibrium then is unemployment at the natural rate combined with high inflation. Compare this to the Ramsey outcome of unemployment at the natural rate and inflation of zero.

19.4 Optimal Nominal Interest Rate Targets

In this section we will consider the government's optimal choice of nominal interest rates. We will consider the real cost of inflation, whereas previously we had simply taken it as given that the government disliked inflation. We will use the simple inventory model of cash holdings from Chapter 4 to show that households are best off when the nominal interest rate is zero. This is a form of what is known as the *Friedman rule.* It appears frequently in monetary economics.

Recall that the nominal interest rate R, the real interest rate r and the expected inflation rate π^e are related by the *Fisher formula*: $R = r + \pi^e$. For this discussion we will take the real interest rate as fixed and beyond the control of the government. Furthermore, we will assume that the government cannot directly manipulate inflationary expectations, and that the private sector correctly forecasts inflation. That is: $\pi^e = \pi$. Thus the government influences the nominal interest rate only through its choice of the actual inflation rate, π. The intuition behind the Fisher formula is quite compelling: households demand a premium of π for holding assets denominated in money, which is losing value at the rate of inflation.

In our model there will be no production. Households own a stock of interest bearing assets, which earn a nominal rate of return of R, and a stock of zero interest money. Money must be used for transactions. There is a fixed cost of μ of converting the interest bearing assets into money, which must be paid every time the household goes to the bank to replenish its cash inventory. The household has real consumption at a rate c per period which it does not vary.

The household goes to the bank x times in one year, so it goes $1/x$ of a year between trips to the bank. To have enough cash on hand to meet its consumption requirement c per period over those $1/x$ periods, the household has to withdraw a real amount c/x at every trip. Thus average real cash holdings over the entire year are $c/(2x)$. Those cash balances could have been invested in interest bearing assets earning an amount R over the year. Thus the foregone interest cost is: $(Rc)/(2x)$. Each time the household goes to the bank to replenish its cash inventory, it incurs a real cost of μ. Thus, transactions costs are: μx. Total costs for a particular policy x are:

$$\mu x + R\frac{c}{2x}.$$

The household minimizes total costs. The minimization problem has a first order condition of:

$$-\frac{cR}{2}\frac{1}{x^2} + \mu = 0.$$

Solving for x produces Baumol and Tobin's famous square-root rule for trips to the bank:

$$x = \sqrt{\frac{cR}{2\mu}}.$$

We can plug the household's decision x back into its cost function to determine the household's annual cash management costs, $\omega(\mu, c, R)$:

$$\omega(\mu, c, R) = \sqrt{2\mu c R}.$$

It is increasing in the fixed charge of going to the bank, μ, the rate of consumption, c and the nominal interest rate R.

A benevolent government that wishes to minimize the household's costs by choice of R would clearly choose to set $R = 0$. At this interest rate, the household goes to the bank only once in its lifetime and incurs no interest penalty for holding money. This is because money also earns in a real interest rate of r. This can only be the case if inflation is negative. From the Fisher formula $R = r + \pi$, we see that $R = 0$ implies that $\pi = -r$. So if $R = 0$, money is a perfect substitute for bonds. Holding a dollar isn't so bad, because next year the household will be able to purchase more with that dollar than it can now.

Although this model is quite limited, it points to one of the important real costs of inflation. Inflation causes households to engage in privately useful but socially useless activities. In times of high inflation, households find it in their interest to spend time and real resources economizing on cash balances. Notice that this is the first indication we have had this chapter that perfectly anticipated inflation is harmful.

If the idea of negative inflation rates seems outlandish, think of the Friedman rule instead as advocating paying interest on money. It is difficult (though not impossible) to pay interest on cash holdings (C. A. E. Goodhart, an English central banker, suggested having a lottery based on cash serial numbers), it is quite easy to pay interest on demand deposits.

Exercises

Exercise 19.1 (Easy)
True, False, or Uncertain (and explain):

1. The Consumer Price Index overstates increases in the "true" cost-of-living index.
2. Inflation is bad because to fight it the Fed increases interest rates, which hurts Americans.

Exercise 19.2 (Easy)
Do governments prefer Phillips curves that are relatively flat (low value of γ) or relatively steep (high values of γ)?

Exercise 19.3 (Moderate)
Assume that the government has a payoff over inflation π and unemployment u of:

$$V^g(u, \pi) = -\phi u^2 - \pi^2.$$

Here $\phi > 0$. The larger ϕ is, the nicer the central banker (that is, the more the central banker cares about the unemployed. Assume that there is a Phillips curve of the form in equation (19.6). Answer the following questions:

1. Assume that inflationary expectations are fixed at π^e. Find the optimal inflation rate choice of the government, $\pi_0(\phi)$.
2. For fixed inflationary expectations, find the corresponding choice of unemployment rate, $u_0(\phi)$.
3. Now assume that the private sector is aware of the government's maximization problem and knows ϕ perfectly. Find the inflation rate π_1 at which expectations are met. What is the associated unemployment rate, u_1?

Variable	Definition
Q_i, q_i	Output (and its log) in industry i.
L_i, ℓ_i	Employment (and its log) in industry i.
α	Production parameter common across industries.
P, p	Common but unobserved general price level (and its log).
P_i, p_i	Observed price (and its log) in industry i.
Z_i, z_i	Shock (and its log) specific to industry i, also industry i's relative price.
P^e, p^e	Common price forecast (and its log).
Z_i^e, z_i^e	Estimated industry-specific shock and its log.
N, n	Population and its log.
A	Parameters (used to make notation neat).
u	In Lucas model: the "not-employment" rate, elsewhere, the unemployment rate.
$u^\star$	The natural rate of unemployment, that is, the rate of unemployment when all industries correctly estimate their specific shocks.
γ	Slope of Phillips curve.
π	Actual inflation (chosen by government).
π^e	Expected inflation (chosen by private sector).
$\pi^\star(\pi^e)$	Government's optimal choice of inflation when inflationary expectations are π^e.
$V^g(u, \pi)$	Government's preferences over unemployment and inflation (it dislikes both equally).
V^p	Private sector payoff: industries dislike making errors in estimating the inflation rate.
π_1	Inflation rate at which expected and chosen inflation coincide, the Nash equilibrium.
u_1	Unemployment rate at Nash, just equal to $u^\star$.

Table 19.1: Notation for Chapter 19

4. Would you prefer to live in a country whose government has a high value of ϕ or a low value of ϕ?

Exercise 19.4 (Moderate)
For this exercise, we will consider what happens when the government and the private sector repeatedly interact. Unemployment in period t u_t, inflation π_t and inflationary ex-

pectations π_t^e are related by the simple Phillips curve:

$$u_t = u_t^\star + \gamma(\pi_t^e - \pi_t), \text{ for all } t = 0, 1, \ldots, \infty.$$

The parameter γ is fixed over time. The government knows about the Phillips curve, but the private sector does not. The government has preferences over unemployment and inflation in period t of:

$$V_t^g(\pi_t, u_t) = -{u_t}^2 - {\pi_t}^2, \text{ for all } t = 0, 1, \ldots, \infty.$$

The private sector sets inflationary expectations based on last period's inflation. This is known as *adaptive expectations.* As a result, π_t^e is given by:

$$\pi_t^e = \pi_{t-1}, \text{ for all } t = 1, 2, \ldots, \infty.$$

Assume that $\pi_0^e = 0$, that is, the private sector begins by believing that inflation will be zero. Answer the following questions:

1. Assume that the government takes as given expectations in a period π_t^e and picks the inflation rate π_t which gives it the highest payoff in period t. Find the government's choice rule $\pi_t^\star(\pi_t^e)$.
2. If the government sets inflation $\pi_t = \pi^\star(\pi_t^e)$, how do expectations evolve over time? Thus right down a law of motion for inflation, $\pi_t(\pi_{t-1})$.
3. What do the trajectories of inflation and unemployment look like over time? Are they rising or falling? Do they settle down? If so, where?
4. How would your answer have been different if, instead of the initial expected inflation being zero, it had been some very large number instead?
5. Now assume that the Phillips curve is augmented with a mean zero shock term, ε, so:

$$u_t = u^\star + \gamma(\pi_t^e - \pi_t) + \varepsilon_t.$$

   Assume that the government knows the value of ε_t and reacts appropriately. Now what happens?

Exercise 19.5 (Easy)
To answer this exercise, you need to answer Exercise 19.4 above. Imagine that the private sector has adaptive expectations about the government's inflationary policy over time, but that part of expected inflation is the government's *announced* inflation target. This announced inflation target is merely an announcement and has nothing to do with reality. If π_t^a is the announced target for period-t inflation, expectations satisfy:

$$\pi_t^e = \delta\pi_{t-1} + (1 - \delta)\pi_t^a, \text{ for all } t = 1, 2, \ldots, \infty.$$

Here $0 < \delta < 1$ is a parameter indexing how much weight the private sector puts on the government's announced inflation rate target. Assume that the government lies constantly, and announces $\pi_t^a = 0$ always. Assume that the government, as in Exercise 19.4, always chooses the inflation rate that maximizes its one-period payoff. Find the steady-state levels of inflation and unemployment.

Solutions to Exercises

Exercise 1.1
You may have noticed that this question glosses over the compounding issue. You were intended to assume that the APR was quoted as a simple interest rate. Accordingly, the daily interest rate is just:

$$R = \frac{16.8\%}{365} \approx 0.046027\%.$$

Exercise 1.2
This exercise glossed over the compounding issue again. Assuming no compounding over the week, the interest rate is:

$$R = \left(\frac{25}{1,000}\right)(52) = 1.3 = 130\%.$$

Exercise 1.3
The key to this question is that the units you use to measure time in the exponent are the same units of time for the resulting interest rate. For example, if you measure n in years, then solving for R gives you an annual interest rate. If you measure n in "quarters", then R will be a quarterly interest rate.

Since this question asks you to annualize the answer, you want to measure n in years. The time interval is 3 months, which is 1/4 of a year. Accordingly:

$$157.8 = \left[e^{(R)(1/4)}\right](156.7), \text{ so:}$$
$$R = 4[\ln(157.8) - \ln(156.7)] \approx 0.02798104 = 2.798104\%.$$

Exercise 1.4
You do *not* want to annualize these interest rates, so you measure n in quarters, i.e., $n = 1$:

1st quarter: $R_1 = \ln(156.7) - \ln(155.7) \approx 0.00640207 = 0.640207\%$.
2nd quarter: $R_2 = \ln(157.8) - \ln(156.7) \approx 0.00699526 = 0.699526\%$.
3rd quarter: $R_3 = \ln(158.6) - \ln(157.8) \approx 0.00505690 = 0.505690\%$.
4th quarter: $R_4 = \ln(160.0) - \ln(158.6) \approx 0.00878851 = 0.878851\%$.

You can see that by adding these four lines together, all but two terms cancel, leaving: $R = \ln(160.0) - \ln(155.7) \approx 0.02724274 = 2.724274\%$. And of course, this is precisely the formula for the annual growth using a continuous interest rate.

Exercise 1.5

$$(2)(\text{GDP}) = \left[e^{(0.02)(n)}\right](\text{GDP}), \text{ so:}$$
$$n = \frac{\ln(2)}{0.02} = 34.66 \text{ years.}$$

The Rule of 72 says that it should take about $72/2 = 36$ years, which is pretty close (it is off by about 3.9%). Of course, you are smart enough to look at:

$$n = \frac{\ln(2)}{0.02}$$

and notice that a better rule would be the "Rule of 69", but nobody is very good at dividing into 69 in their head.

Exercise 1.6
The first thing you need to do is calculate the number of whole (i.e., undivided) years this investment will require. There is some number n of years such that:

$$(1 + 0.065)^n(\$10,000) < \$15,000, \quad \text{but where:}$$
$$(1 + 0.065)^{n+1}(\$10,000) > \$15,000.$$

This implies that $n = 6$. After the 6th year, the investment has grown to:

$$(1 + 0.065)^6(\$10,000) = \$14,591.42.$$

That becomes the principal of the investment in the 7th year, since interest was able to compound at the end of the 6th year. Now you need to figure out the number of days of simple interest the investment will need in the 7th year. It is short of \$15,000 by \$408.58, and each day the investment earns:

$$\left(\frac{0.065}{365}\right)(\$14,591.42) = \$2.60.$$

You use these facts to calculate the required number of days:

$$\frac{\$408.58}{\$2.60} = 157.23.$$

Since interest only accrues after a full day, the investment would not earn the interest from the last 0.23 days until 158 days had passed. All in all then, the investment would require 6 years and 158 days.

Exercise 1.7
Let F_t be the acreage of forest in year t. Then:

$$F_n = F_0(1 - 0.046)^n.$$

You are looking for the n such that $F_n = (0.5)F_0$. Plugging this into the equation and taking logs of both sides yields: $n \approx 14.72$, so half will be cleared in 15 years.

Exercise 1.8

1. The relevant formula is:

$$679e^{100R} = 954.$$

   so $R \approx 3.4 \times 10^{-3}$.

2. Letting x be the number of years, the relevant formula is:

$$2e^{xR} = 679,000,000.$$

   You use R from the previous part. Solving yields $x \approx 5777$, so they would have left the Garden of Eden in about 4075 BCE. (Population growth was probably slower in the past, so this is likely not early enough.)

Exercise 1.9

1. You are solving for n in the following equation:

$$\$10,600(1.047)^n = \$15,400(1.017)^n.$$

   This implies that $n \approx 12.85$ years, so the incomes would be the same sometime in 1996. (Nb: They were not. Japan's income was still less at that date.)

2. Just plug in the value of n:

$$\$15,400(1.017)^{12.85} \approx \$19,124.$$

Exercise 2.1

$$\max_{c,l,n_s} \{\ln(c) + \ln(l)\}, \text{ such that:}$$
$$c = 4n_s^{0.5} + (24 - l - n_s)w.$$

Exercise 2.2
First, write out the maximization problem:

$$\max_{c,l} \left\{c^\gamma (1-l)^{1-\gamma}\right\}, \text{ such that:}$$
$$c = y = Al^\alpha.$$

Plug the constraint into the objective to get an unconstrained maximization problem:

$$\max_{l} \left\{ (Al^{\alpha})^{\gamma}(1-l)^{1-\gamma} \right\}.$$

There is only one first-order condition:

$$\text{(FOC } l\text{)} \qquad \gamma[A(l^{\star})^{\alpha}]^{\gamma-1}[A\alpha(l^{\star})^{\alpha-1}](1-l)^{1-\gamma} = [A(l^{\star})^{\alpha}]^{\gamma}(1-\gamma)(1-l^{\star})^{-\gamma}.$$

After a bunch of canceling and rearranging, this reduces to:

$$l^{\star} = \frac{\gamma\alpha}{1-\gamma+\gamma\alpha}.$$

Plugging this back into the $c = y = f(l)$ constraint yields:

$$c^{\star} = A\left(\frac{\alpha\gamma}{1-\gamma+\gamma\alpha}\right)^{\alpha}.$$

Exercise 3.1

1. The marginal period utility is:

$$u'(c_t) = \frac{1}{2}c_t^{-\frac{1}{2}}.$$

   Plugging this into equation (3.7) yields:

$$\frac{\frac{1}{2}(c_1^{\star})^{-\frac{1}{2}}}{\frac{1}{2}(c_2^{\star})^{-\frac{1}{2}}} = \beta(1+R), \text{ or:}$$
$$\left(\frac{c_2^{\star}}{c_1^{\star}}\right)^{\frac{1}{2}} = \beta(1+R).$$

2. We have three unknowns: $c_1^{\star}$, $c_2^{\star}$, and $b_1^{\star}$. The three equations relating them are: the Euler equation above and the two budget equations. Solving these is an unpleasant exercise in algebra. Solve the Euler equation for $c_2^{\star}$:

$$c_2^{\star} = \beta^2(1+R)^2c_1^{\star}.$$

   Use this to remove $c_2^{\star}$ from the second-period budget:

$$Py_2 + b_1^{\star}(1+R) = P\left[\beta^2(1+R)^2c_1^{\star}\right].$$

   Solve this for $c_1^{\star}$, and plug the result into the first-period budget:

$$Py_1 = P\left[\frac{Py_2 + b_1^{\star}(1+R)}{P\beta^2(1+R)^2}\right] + b_1^{\star}.$$

This looks awful, but it reduces to:

$$b_1^\star = Py_1 - \frac{P[y_2 + y_1(1+R)]}{[1+\beta^2(1+R)](1+R)},$$

which is the answer for the household's choice of b_1. Plugging this back into the first-period budget gives the optimal c_1:

$$c_1^\star = \frac{y_2 + y_1(1+R)}{[1+\beta^2(1+R)](1+R)}.$$

Finally, we plug the answer for $c_1^\star$ into the second-period budget equation to get:

$$c_2^\star = \frac{\beta^2(1+R)[y_2 + y_1(1+R)]}{1+\beta^2(1+R)}.$$

3. In equilibrium, $b_1^\star = 0$, so:

$$Py_1 = \frac{P[y_2 + y_1(1+R^\star)]}{[1+\beta^2(1+R^\star)](1+R^\star)},$$

Solving for $R^\star$ yields:

$$R^\star = \left(\frac{y_2}{\beta^2 y_1}\right)^{\frac{1}{2}} - 1.$$

4. From the above equation, we see that an equal percentage increase in y_1 and y_2 will have no effect on the equilibrium interest rate $R^\star$, just like under logarithmic preferences.

Exercise 3.2

1.

$$\frac{\partial R^\star}{\partial \beta} = -\frac{y_2}{\beta^2 y_1} < 0.$$

Greater impatience means β decreases (say, from 0.95 to 0.9), and $R^\star$ moves in the opposite direction, so the equilibrium interest rate increases.

2.

$$\frac{\partial R^\star}{\partial y_1} = -\frac{y_2}{\beta y_1^2} < 0,$$

so smaller first-period income causes the equilibrium interest rate to increase.

Exercise 3.3

1.

$$\max_{c_1,c_2,s} \left\{\ln(c_1) + \beta\ln(c_2)\right\}, \quad \text{subject to:}$$

$$c_1 + s = e_1, \quad \text{and:}$$

$$c_2 = e_2 + (1-\delta)s.$$

2. The Lagrangean for this problem is:

$$\mathcal{L} = \ln(c_1) + \beta \ln(c_2) + \lambda_1[e_1 - c_1 - s] + \lambda_2[e_2 + (1-\delta)s - c_2].$$

The first-order conditions are:

$$\text{(FOC } c_1\text{)} \qquad \frac{1}{c_1} - \lambda_1 = 0,$$

$$\text{(FOC } c_2\text{)} \qquad \frac{\beta}{c_2} - \lambda_2 = 0, \quad \text{and:}$$

$$\text{(FOC } s\text{)} \qquad -\lambda_1 + \lambda_2(1-\delta) = 0.$$

We also have two first-order conditions for the Lagrange multipliers, but we leave those off, since they just reproduce the constraints. We can quickly solve the above equations to remove the Lagrange multipliers, giving us:

$$\frac{c_2}{c_1} = \beta(1-\delta).$$

We combine this with our two constraints to get:

$$s = \frac{\beta(1-\delta)e_1 - e_2}{(1-\delta)(1+\beta)},$$
$$c_1 = e_1 - \frac{\beta(1-\delta)e_1 - e_2}{(1-\delta)(1+\beta)}, \quad \text{and:}$$
$$c_2 = e_2 + \frac{\beta(1-\delta)e_1 - e_2}{1+\beta}.$$

3. We just take the derivatives of the above answers with respect to δ:

$$\frac{\partial c_1}{\partial \delta} = -\frac{e_2(1+\beta)}{(1-\delta)^2(1+\beta)^2} < 0,$$
$$\frac{\partial c_1}{\partial \delta} = -\frac{\beta e_1}{1+\beta} < 0, \quad \text{and:}$$
$$\frac{\partial s}{\partial \delta} = -\frac{e_2(1+\beta)}{(1-\delta)^2(1+\beta)^2} < 0.$$

Here, Maxine has learned how to defend against rats, so we are interested in δ going down. The negative derivatives above imply that all the choices change in the opposite direction, so consumption in both periods and first-period saving all increase.

Exercise 3.4

$$\max_{s_0,\dots,s_4} \left\{ \sum_{t=0}^{4} \beta^t \ln(c_t) \right\}, \text{ such that:}$$
$$c_t = (1-s_t)x_t, \text{ for } t = 0,\dots,4,$$
$$x_{t+1} = (1+\alpha)s x_t, \text{ for } t = 0,\dots,4, \text{ and:}$$
$$x_0 \text{ is some given constant.}$$

Exercise 4.1

1. Take the derivative of real money demand with respect to the interest rate R:

$$\frac{\partial\phi(R,c,\gamma/P)}{\partial R} = \left(\frac{1}{2}\right) c \left(\frac{2\gamma}{PRc}\right)^{-\frac{1}{2}} \left(-\frac{2\gamma}{PR^2c}\right) < 0,$$

so the interest rate and real money holdings move in opposite directions. An increase in the interest rate causes the consumer to hold less real money.

2. Differentiate with respect to c:

$$\begin{aligned}\frac{\partial\phi(R,c,\gamma/P)}{\partial c} &= \left(\frac{1}{2}\right) c \left(\frac{1}{2}\right) \left(\frac{2\gamma}{PRc}\right)^{-\frac{1}{2}} \left(-\frac{2\gamma}{PRc^2}\right) + \left(\frac{2\gamma}{PRc}\right)^{\frac{1}{2}} \left(\frac{1}{2}\right) \\ &= \left(\frac{1}{2}\right) \left(\frac{2\gamma}{PRc}\right)^{\frac{1}{2}} - \left(\frac{1}{4}\right) c \left(\frac{2\gamma}{PRc}\right)^{\frac{1}{2}} \left(\frac{1}{c}\right) \\ &= \left(\frac{1}{4}\right) \left(\frac{2\gamma}{PRc}\right)^{\frac{1}{2}} > 0,\end{aligned}$$

so consumption and real money holdings move in the same direction. If the consumer consumes more, then the consumer will hold more real money.

3. First, we replace γ/P with α, giving us:

$$\phi(R,c,\alpha) = \left(\frac{1}{2}\right) c \left(\frac{2\alpha}{Rc}\right)^{\frac{1}{2}}.$$

Taking the derivative with respect to α gives us:

$$\frac{\partial\phi(R,c,\gamma/P)}{\partial\alpha} = \left(\frac{1}{2}\right) c \left(\frac{1}{2}\right) \left(\frac{2\alpha}{Rc}\right)^{-\frac{1}{2}} \left(\frac{2}{Rc}\right) > 0,$$

so real money holdings and real transactions costs move in the same direction. If the consumer faces higher real transactions costs, the consumer will hold more real money.

Exercise 5.1

The budget constraint for the first period was given by:

$$Py_1 = Pc_1 + b_1. \tag{3.2}$$

The condition for clearing the goods market in the first period was:

$$Ny_1 = Nc_2. \tag{3.10}$$

This implies $y_1 = c_1$. Plugging this into (3.2) gives:

$$\begin{aligned}Pc_1 &= Pc_1 + b_1, \text{ or:} \\ b_1 &= 0,\end{aligned}$$

which is the market-clearing constraint for bonds.

Exercise 5.2

The price for a good can only be zero if all consumers are satiated with that good, that is, if they cannot increase their utility by consuming more of it. In our model this is ruled out because all utility functions are strictly increasing in all arguments. This implies that the consumers always prefer to consume more of each good. If the price for a good were zero, they would demand infinite amounts, which would violate market-clearing. Therefore, with strictly increasing utility functions, all prices are positive. If utility is not strictly increasing, zero prices are possible. In that case, Walras' Law might not hold, because total demand by consumers can be less than the total endowment. The proof of Walras' Law fails once we use the fact that the price of each good is positive. On the other hand, the First Welfare Theorem still goes through, since it does not rest on the assumption of positive prices.

Exercise 6.1

1. The first-order condition with respect to l_d^a is:

   $$(0.5)(l_d^{a\star})^{-0.5} - w = 0. \qquad \text{(FOC } l_d^a\text{)}$$

   Solving for $l_d^{a\star}$ yields: $l_d^{a\star} = \frac{1}{4w^2}$. The farm's profit is:

   $$(l_d^{a\star})^{0.5} - wl_d^{a\star} = \left(\frac{1}{4w^2}\right)^{0.5} - \frac{w}{4w^2} = \frac{1}{4w} = \pi^{a\star}.$$

2. The first-order condition with respect to l_d^b is:

   $$(0.5)(2)(l_d^{b\star})^{-0.5} - w = 0. \qquad \text{(FOC } l_d^b\text{)}$$

   Solving for $l_d^{b\star}$ yields: $l_b^{d\star} = \frac{1}{w^2}$. The farm's profit is:

   $$2(l_b^{d\star})^{0.5} - wl_b^{d\star} = 2\left(\frac{1}{w^2}\right)^{0.5} - \frac{w}{w^2} = \frac{1}{w} = \pi^{b\star}.$$

3. For economy, we will work this out for an unspecified $\pi^{j\star}$, where j is either a or b. We'll plug those in later. Substitute the constraint into to objective in order to eliminate c^j. This gives us:

   $$\max_{l_s^j} \left\{\ln(wl_s^j + \pi^{j\star}) + \ln(24 - l_s^j)\right\}.$$

   We carry out the maximization:

   $$\frac{w}{wl_s^j + \pi^{j\star}} + \frac{-1}{24 - l^{j\star}} = 0, \text{ so:} \qquad \text{(FOC } l_s^j\text{)}$$
   $$(24 - l_s^{j\star})w = wl_s^{j\star} + \pi^{j\star}.$$

Solving for $l_s^{j\star}$ yields:

$$l_s^{j\star} = \frac{24w - \pi^{j\star}}{2w}. \tag{S.8}$$

For this exercise, we are using $\pi^{a\star}$, so we plug that in to get:

$$l_s^{a\star} = \frac{24w - \frac{1}{4w}}{2w} = 12 - \frac{1}{8w^2}.$$

4. We can just re-use equation (S.8), but this time we plug in $\pi^{b\star}$, yielding:

$$l_s^{b\star} = \frac{24w - \frac{1}{w}}{2w} = 12 - \frac{1}{2w^2}.$$

5.

$$l_d^\star = 400 l_d^{a\star} + 700 l_d^{b\star} = \frac{400}{4w^2} + \frac{700}{w^2} = \frac{800}{w^2}.$$

6.

$$\begin{aligned} l_s^\star &= 400 l_s^{a\star} + 700 l_s^{b\star} \\ &= (400)\left(12 - \frac{1}{8w^2}\right) + (700)\left(12 - \frac{1}{2w^2}\right) \\ &= (1,100)(12) - \frac{400}{w^2}. \end{aligned}$$

7. We want to set $l_d^\star = l_s^\star$, or:

$$\frac{800}{w^2} = (1,100)(12) - \frac{400}{w^2},$$

which reduces to $w^2 = 1/11$, or $w \approx 0.3015$.

Exercise 6.2

1. We take the first-order condition of equation (6.9) with respect to l_d:

$$\frac{\partial \pi}{\partial l_d} = \left(\frac{7}{10}\right) A k^{\frac{3}{10}} (l_d^\star)^{-\frac{3}{10}} - w = 0 \tag{FOC l_d}$$

When we solve that for $l_d^\star$, we get:

$$l_d^\star = \left(\frac{7A}{10w}\right)^{\frac{10}{3}} k$$

2. The result is as follows:

$$\begin{aligned}
\pi^\star &= Ak^{\frac{3}{10}}(l_d^\star)^{\frac{7}{10}} - wl_d^\star \\
&= Ak^{\frac{3}{10}}\left[\left(\frac{7A}{10w}\right)^{\frac{10}{3}} k\right]^{\frac{7}{10}} - w\left(\frac{7A}{10w}\right)^{\frac{10}{3}} k \\
&= Ak^{\frac{3}{10}}\left(\frac{7A}{10w}\right)^{\frac{7}{3}} k^{\frac{7}{10}} - w\left(\frac{7A}{10w}\right)^{\frac{10}{3}} k \\
&= k\left(\frac{7A}{10w}\right)^{\frac{7}{3}}\left[A - w\left(\frac{7A}{10w}\right)\right] \\
&= \left(\frac{3A}{10}\right)\left(\frac{7A}{10w}\right)^{\frac{7}{3}} k.
\end{aligned}$$

3.

$$\max_{c,l_s}\{c^{\frac{1}{2}}(1-l_s)^{\frac{1}{2}}\}, \text{ subject to:}$$
$$c = \pi^\star + wl_s.$$

4. To begin, we can leave the $\pi^\star$ term in the Lagrangean:

$$\mathcal{L} = c^{\frac{1}{2}}(1-l_s)^{\frac{1}{2}} + \lambda[\pi^\star + wl_s - c].$$

Our first-order conditions are:

(FOC c) $$\left(\frac{1}{2}\right)(c^\star)^{-\frac{1}{2}}(1-l_s^\star)^{\frac{1}{2}} + \lambda^\star[-1] = 0, \text{ and:}$$

(FOC l_s) $$(c^\star)^{\frac{1}{2}}\left(\frac{1}{2}\right)(1-l_s^\star)^{-\frac{1}{2}}(-1) + \lambda^\star[w] = 0.$$

We leave off the FOC for λ. Combining the above FOCs to get rid of $\lambda^\star$ yields:

(S.9) $$c^\star = w(1-l_s^\star).$$

We plug this result and our expression for $\pi^\star$ into the budget equation $\pi + wl_s = c$, yielding:

$$\left(\frac{3A}{10}\right)\left(\frac{7A}{10w}\right)^{\frac{7}{3}} k = w(1-l_s^\star) - wl_s^\star.$$

Solving this for $l_s^\star$ gives us:

$$l_s^\star = \frac{1}{2}\left[1 - \frac{3}{7}\left(\frac{7A}{10w}\right)^{\frac{10}{3}} k\right].$$

When we plug this value of $l_s^\star$ back into equation (S.9), we get the optimal consumption $c^\star$:

$$c^\star = \frac{w}{2}\left[1 + \frac{3}{7}\left(\frac{7A}{10w}\right)^{\frac{10}{3}} k\right].$$

5. We just set $l_s^\star = l_d^\star$, solve for w, and call the result $w^\star$:

$$\left(\frac{7A}{10w^\star}\right)^{\frac{10}{3}} k = \frac{1}{2}\left[1 - \frac{3}{7}\left(\frac{7A}{10w^\star}\right)^{\frac{10}{3}} k\right].$$

   After a bunch of algebra, we get:

$$w^\star = \left(\frac{7A}{10}\right)\left(\frac{17k^{\frac{3}{10}}}{7}\right).$$

6. We are interested in the derivative if $w^\star$ with respect to k:

$$\frac{\partial w^\star}{\partial k} = \left(\frac{7A}{10}\right)\left(\frac{3}{10}\right)\left(\frac{17k}{7}\right)^{-\frac{7}{10}}\left(\frac{17}{7}\right) = \left(\frac{51A}{100}\right)\left(\frac{17k}{7}\right)^{-\frac{7}{10}}.$$

   Since this derivative is positive, $w^\star$ increases as k does.

7. The U.S. has a much larger stock of capital (per capita) than Mexico does. According to this model, that difference alone causes wages to be higher in the U.S. From the equation for the equilibrium wage $w^\star$, we see that increasing the per-capita capital stock k by a factor of two causes the wage to increase, but by less than a factor of two. Hence, wages between the two countries differ by less (in percentage terms) than their per-capita capital stocks.

   Of course, owners of capital try to export it to wherever labor is cheapest. In this case, the households in the U.S. try to send some of their capital to Mexico in order to take advantage of lower wages there. If this movement of capital is restricted, then the wage difference will persist, and there will be an incentive for workers to move to the country with more capital. In this case, Mexican workers will see higher wages across the border and will immigrate to the U.S. where they will earn more.

Exercise 8.1
According to the quantity theory, the inflation rate is approximately equal to the difference between the growth rate of money supply and the growth rate of output. Since the question assumes that velocity is constant, the quantity theory applies. The annual rate of inflation is therefore two percent.

Exercise 8.2
In Chapter 8 we determined that velocity is inversely related to the time spent between two trips to the bank. In Chapter 4 we saw that the time between two trips to the bank

decreases when the nominal interest rate increases. Therefore velocity and the nominal interest rate are positively related. In Section 8.3 we found out that inflation and nominal interest rates are positively related. Therefore, a high inflation rate results in high nominal interest rates and high velocity. This is also true in the real world: velocity is much higher in countries with high inflation than in countries with moderate inflation. Intuitively, high inflation means that money quickly loses value. It is therefore not attractive to hold a lot of money, so money circulates quickly. In countries with hyperinflation, wages are often paid daily, and workers usually spend wages the same day they receive them.

Exercise 9.1
Of course, the solution depends on the country you pick. As an example, Figure S.3 displays GDP and its trend for Germany. You can see that the trend does not look that much smoother than the actual series. This shows that our method of computing the trend is not especially good.

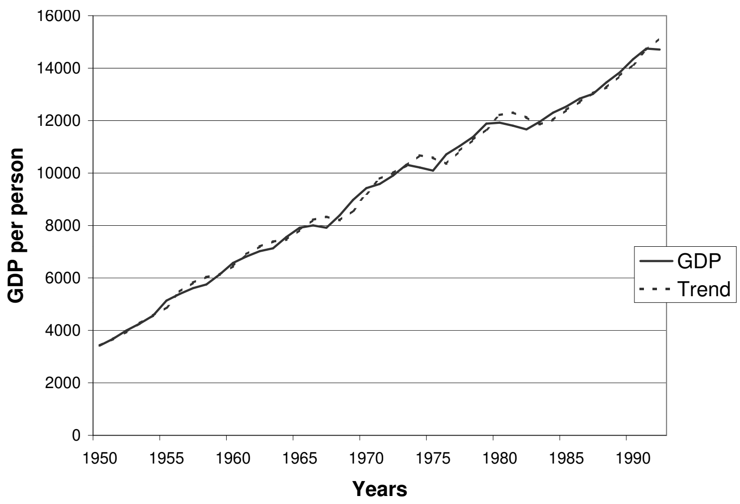

Figure S.3: GDP and Trend

Exercise 9.2
Figure S.4 shows the cyclical component for Germany. Your business cycle should look similar, unless your country is a former member of the communist block. Those countries either had radically different business cycles, or, more likely, they adjusted their statistics in order to get nice, smooth figures.

Exercise 9.3
For Germany, there are ten peaks in the cyclical component. The duration of a full cycle is between three and six years, with the average slightly above four years. The overall

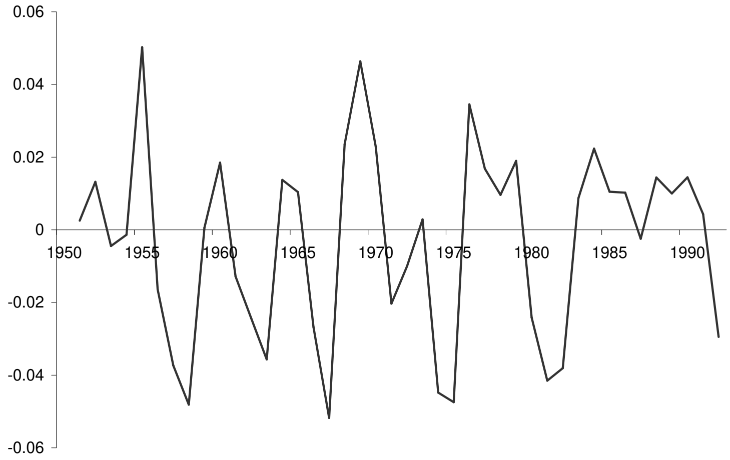

Figure S.4: The Cyclical Component

amplitude of the cycles is relatively stable. Although there are some general similarities, the cycles are of quite different shape. The process generating the cycles seems not to have changed much, however. The cycles in the fifties and sixties are not much different from those in the eighties and nineties.

Exercise 9.4
By using the resource constraints, we can write the problem as:

$$\max \ln(\sqrt{Bk_t} + \epsilon_t - i_t) + A\ln((1-\delta)k_t + i_t).$$

The first-order condition is:

$$0 = -\frac{1}{\sqrt{Bk_t} + \epsilon_t - i_t} + \frac{A}{(1-\delta)k_t + i_t}.$$

Solving for i_t, we get:

$$i_t = \frac{A[\sqrt{Bk_t} + \epsilon_t] - (1-\delta)k_t}{1+A}.$$

Using the resource constraint for the first period, we can solve for c_t:

$$c_t = \frac{\sqrt{Bk_t} + \epsilon_t + (1-\delta)k_t}{1+A}.$$

Exercise 9.5
The derivatives are:

$$\frac{\partial i_t}{\partial \epsilon_t} = \frac{A}{1+A} = 0.8, \text{ and:}$$
$$\frac{\partial c_t}{\partial \epsilon_t} = \frac{1}{1+A} = 0.2.$$

The numbers correspond to the value $A = 4$ that is used for the simulations. Investment reacts much stronger to shocks than consumption does, just as we observe in real-world data.

Exercise 9.6
Figure S.5 shows consumption and investment, and Figure S.6 is GDP. Investment is much more volatile than consumption. The relative volatility of consumption and investment is comparable to what we find in real data. We simulated the economy over 43 periods, because there were also 43 years of data for German GDP. In the simulation there are nine peaks, which is close to the ten peaks we found in the data. The length of the cycle varies from four to seven years. The average length is a little less than five periods, while the German cycles lasted a little more than four years on average.

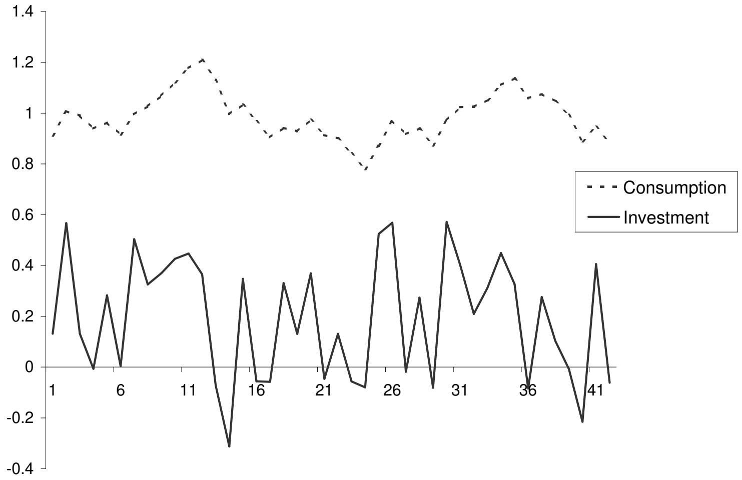

Figure S.5: Simulated Consumption and Investment

Exercise 9.7
The aim of real business cycle research is to gain a better understanding of business cycles. The theory differs from other approaches mainly by the methods that are applied. Real business cycle models are fully specified stochastic equilibrium models. That means that the microfoundations are laid out in detail. There are consumers with preferences, firms

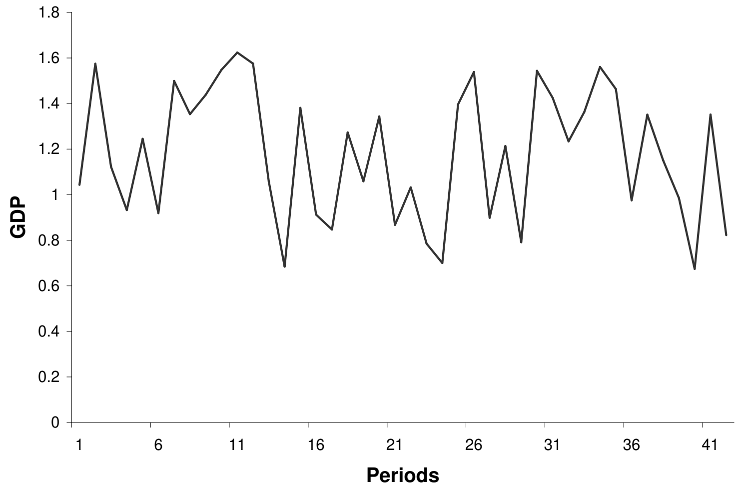

Figure S.6: Simulated GDP

with technologies, and a market system that holds everything together. Real business cycle theory takes the simplest models of this sort as a point of departure to explain business cycles. Model testing is most often done with the "calibration" method. This means that first the model parameters are determined by making them consistent with empirical facts other than the business cycle facts that are supposed to be explained. The parameterized model is then simulated, and the outcomes are compared with real world data.

Exercise 9.8
Plosser's model does not contain a government, and even if there were one, there would be no need to stabilize the economy. There are no market frictions in the model; the outcomes are competitive equilibria. By the First Welfare Theorem we know that equilibria are efficient, so there is nothing a government could do to improve economic outcomes. It is possible to extend the model to allow for a government, and we could add frictions to the model to make intervention beneficial, without changing the general framework very much. Also, any government is certainly able to produce additional shocks in the economy. Still, real business cycle theory works fine without a government, both as a source of disturbance and as a possible stabilizer.

Exercise 10.1

1. This is the myth of small business job creation again. The SBA has every reason to tout the influence of small small businesses, but, as DHS point out, the dominant job market role if played by large, old firms and plants.

2. This rather entertaining quote has several immediate and glaring errors, but it does contain an argument quite in vogue at the moment. There is a common idea that

jobs are a scarce resource, and that the pool of jobs is shrinking under pressure from greedy company owners, slave labor factories abroad and so on. In reality, as we've seen in this chapter, the pool of jobs is churning all the time. Ten percent of all jobs are typically destroyed in a year, and ten percent are created. In the face of this turmoil, one or two high profile plant closings is simply not important.

Exercise 10.2
The term g_{est} is defined as:

$$g_{est} = \frac{X_{es,t} - X_{es,t-1}}{0.5(X_{es,t} + X_{es,t-1})}.$$

For a new plant $X_{es,t-1} = 0$ and for a dying plant $X_{es,t} = 0$. Thus for a new plant:

$$g_{est} = \frac{X_{es,t} - 0}{0.5(X_{es,t} + 0)} = 2.$$

And for a dying plant:

$$g_{est} = \frac{0 - X_{es,t-1}}{0.5(0 + X_{es,t-1})} = -2.$$

Exercise 10.3
The only thing tricky about this problem is remembering how to deal with absolute values. If $a = b$, then $|a| = b$ if a is positive and $|a| = -b$ if a is negative. For shrinking plants, $\Delta X_{es,t}$ is negative, so for shrinking plants:

$$g_{est} = \frac{\Delta X_{es,t}}{Z_{est}} = -\frac{|\Delta X_{es,t}|}{Z_{est}}.$$

Now we work through the algebra required to get the answer to the first identity. We begin with the definition of c_{st}:

$$c_{st} = \frac{C_{st}}{Z_{st}} = \frac{1}{Z_{st}} \sum_{e \in S^+} \Delta X_{es,t} = \frac{1}{Z_{st}} \sum_{e \in S^+} \frac{Z_{est} \Delta X_{es,t}}{Z_{est}} = \frac{1}{Z_{st}} \sum_{e \in S^+} Z_{est} g_{est}.$$

Turning to the next identity, we begin with the definition of net_{st}:

$$\begin{aligned}
\text{net}_{st} &= \frac{C_{st} - D_{st}}{Z_{st}}, \\
&= \frac{1}{Z_{st}} \sum_{e \in S^+} \Delta X_{es,t} - \frac{1}{Z_{st}} \sum_{e \in S^-} |\Delta X_{es,t}|, \\
&= \frac{1}{Z_{st}} \sum_{e \in S^+} Z_{est} g_{est} - \frac{1}{Z_{st}} \sum_{e \in S^+} Z_{est}(-g_{est}), \\
&= \frac{1}{Z_{st}} \sum_{e \in S} Z_{est} g_{est}.
\end{aligned}$$

Exercise 10.4
This question really just boils down to plugging the definitions of R_t and NET_{st} into the definition of covariance. However, the algebra shouldn't detract from an interesting statistical regularity. Begin with the definition of covariance (supplied in the question):

$$\text{cov}(R_t, NET_t) < 0.$$
$$\frac{1}{N}\sum_{i=1}^{N}(R_i - \overline{R})(NET_i - \overline{NET}) < 0.$$
$$\frac{1}{N}\sum_{i=1}^{N}(C_i + D_i - \overline{C} - \overline{D})(C_i - D_i - \overline{C} + \overline{D}) < 0.$$
$$\frac{1}{N}\sum_{i=1}^{N}\left[(C_i - \overline{C}) + (D_i - \overline{D})\right]\left[(C_i - \overline{C}) - (D_i - \overline{D})\right] < 0.$$
$$\frac{1}{N}\sum_{i=1}^{N}(C_i - \overline{C})^2 - \frac{1}{N}\sum_{i=1}^{N}(D_i - \overline{D})^2 < 0.$$

Using the definition definition of variance supplied in the question, this last inequality can be written $\text{var}(C) - \text{var}(D) < 0$, so $\text{var}(C) < \text{var}(D)$. That was a lot of algebra, but it was all straightforward. Thus if periods of large net job loss coincide with periods of larger than normal job reallocation, it must be the case that job destruction has a higher variance than job creation.

Exercise 10.5
Here is the original chart, now augmented with the answers.

Year	$X_{1,t}$	$X_{2,t}$	$X_{3,t}$	c_t	d_t	net_{st}	UB	LB
1990	1000	0	500					
1991	800	100	800	0.250	0.125	0.125	600	200
1992	1200	200	700	0.263	0.053	0.210	600	400
1993	1000	400	600	0.098	0.146	-0.048	500	100
1994	800	800	500	0.195	0.146	0.049	700	100
1995	400	1200	600	0.233	0.186	0.047	900	100
1996	200	1400	600	0.091	0.091	0	400	0
1997	0	2000	500	0.255	0.128	0.127	900	300

Exercise 10.6
All of these statements referred to specific charts or graphs in DHS. This question was on the Spring 1997 midterm exam in Econ 203.

1. Most students were at least able to say that this hypothesis wasn't exactly true, even if they couldn't identify specifically why. Any two of the following facts were acceptable:
    (a) Even the industries in the highest important ratio quintile had an import penetration rate of about 13.1%, which is pretty low.
    (b) The relationship between import penetration quintile and net job growth and job destruction is not monotone.
    (c) For the highest import penetration quintile, net job growth averaged −2.8% annually.
2. Robots replacing workers is another favorite canard (thankfully less common recently) of the chattering classes. The reality is reflected in DHS Table 3.6 showing gross job flows by capital intensity decile. The most fascinating part of this table is the final entry, showing an average annual *net* employment growth rate of 0.7% for plants in the highest capital intensity decile. Plants in the lowest capital intensity decile shed about 10% of their jobs, net, each year. That is, over the 15-year sample period, they must have become nearly extinct. Thus high capital plants (plants with lots and lots of robots, one presumes) have been steadily adding excellent jobs of the past 20 years.
3. What we were looking for here was some version of Figure 2.2 in DHS, giving the distributions of plant-level job creation and destruction by employment growth g. They have a distinctive "double hump" shape with the first peak at about $g = 0.10$. However, we accepted more general statements about how most destruction occurs at plants which are shutting down and so on.
4. This question is drawn directly from Table 3.6, showing that highly specialized plants have high job creation and destruction rates, and a net growth rate of -2%. Because of their high job destruction rate, and the tendency of plants to close in recessions (the cyclical behavior of job destruction), highly specialized plants are indeed at risk of closing in recessions.
5. For this question we wanted students to tell us about job creation and destruction rates by wage quintile (Table 3.4 in DHS). Any two of the following facts were acceptable:
    (a) Job creation and destruction are falling by plant wage quintile.
    (b) Of all jobs destroyed each year, only about 18% are accounted for by the highest wage quintile, while about 26% are accounted for by the lowest wage quintile.
    (c) High wage jobs tend to be more durable (longer creation persistence).

Exercise 11.1
The aggregate production technology is $Y = 3L^{.7}K^{.3}$, and we have $L = 150$, $\delta = 0.1$, and $s = 0.2$. The law of motion for capital is given by:

$$K_t = (1 - \delta)K_{t-1} + sY_t.$$

Therefore the steady state level of capital $\bar{K}$ has to satisfy:

$$\bar{K} = (1 - \delta)\bar{K} + s(3L^{0.7}\ \bar{K}^{0.3}).$$

Plugging in the values for labor, depreciation, and the saving rate yields:

$$\bar{K} = 0.9\bar{K} + (0.6)(150)^{0.7}\ \bar{K}^{0.3},\ \text{or:}$$
$$0.1\bar{K} = (0.6)(150)^{0.7}\ \bar{K}^{0.3},\ \text{or:}$$
$$\bar{K} = \left[(10)(0.6)(150)^{0.7}\right]^{1/0.7}.$$

Evaluating this expression results in $\bar{K} \approx 1940$. Steady state output $\bar{Y}$ is given by:

$$\bar{Y} = 3L^{0.7}\ \bar{K}^{0.3},$$

which gives us the solution $\bar{Y} \approx 970$.

Exercise 11.2
In terms of the Solow model, the war temporarily reduced the capital stock in Kuwait. Given the lower capital stock, per capita incomes will be lower in the next years. In the long run, the economy reaches the steady state again, so the war does not affect per capita income any more. Similarly, the effect on the growth rate of per capita income is also temporary. In the short run, the growth rate will be higher, because the growth rate of per capita income is inversely related to the capital stock. In the long run, the growth rate of per capita income is determined by the rate of technological progress, so the war does not have an effect on the long-run growth rate. Recovery will be faster if foreigners are allowed to invest, because more investment implies that the economy returns faster to the steady state level of capital. The gains and losses of workers and capitalists depend on the reaction of wages and the return on capital to a higher capital stock due to foreign investment. Our formulas for wage and interest, equations (11.3) and (11.4), indicate that the wage is positively related to the capital stock, while the return to capital is negatively related to the capital stock. Since a prohibition of foreign investment lowers the capital stock, workers would lose, and capitalists would gain by a prohibition.

Exercise 12.1

1. True. Under an unfunded pension system payments to the old are made by taxing the young, *not* by investing in the bond market. Hence the volume of physical savings between periods of life is higher under a funded than an unfunded pension system.

2. Check the *Economic Report of the President* to get a good sense of n, and the back of the *Economist* magazine to get the latest value for r. Unless something very odd is happening, n is probably considerably lower than r.

3. From the *Economic Report of the President* we see that (among others) the U.S government spends more than 20% of its total outlays on interest payments on the Federal debt, social security and defense. We shall have quite a bit more to say about the Federal debt in Chapter 14 and Chapter 18.

Exercise 12.2
The household's budget constraint is:

$$C + I + G = Y^P + Y^G.$$

We are given that private output Y^P is fixed at Y and that government output Y^G is ϕG. Thus government spending G must satisfy:

$$\begin{aligned} G &= Y + \phi G - C - I, \text{ or:} \\ (1-\phi)G &= Y - C - I. \end{aligned}$$

Obviously, as G grows, C and I are going to have to shrink (although not one-for-one with G). The maximum allowed level for government spending occurs when consumption and investment are each zero, so $C = I = 0$. In that case:

$$G = \frac{Y}{1-\phi}.$$

The government can spend more than total private output since its spending is productive. As ϕ is closer to zero, the closer G must be to Y. As ϕ is closer to unity, the larger G may be relative to Y.

Exercise 12.3
To calculate the market-clearing interest rate, we have to find the interest rate that makes the household want to consume precisely its endowment stream *net of government taxes*. Since in this question consumption in each period t must just be $C_t = Y_t - G_t$, we find that:

$$\begin{aligned} 1 + r_0 &= \frac{1}{\beta}\frac{U'(Y-G)}{U'(Y-G)}, \text{ so:} \\ r_0 &= \rho, \text{ and:} \\ 1 + r_0^* &= \frac{1}{\beta}\frac{U'(Y-G_0)}{U'(Y-G_1)}. \end{aligned}$$

We cannot characterize r_0^* further without more information about Y, G and U, but we can say that, since $G_0 > G_1$, the marginal utility in the first period must be greater than the marginal utility in the second period, that is, $U'(Y - G_0) > U'(Y - G_1)$. Thus:

$$\frac{U'(Y-G_0)}{U'(Y-G_1)} > 1.$$

As a result, $r_0^* > r_0$. This fits well with the results of this chapter, which hold that temporary increases in government spending increase the real interest rate.

Exercise 12.4

The household's maximization problem becomes:

$$\max_{S_t} \left\{ 2\sqrt{(1-\tau)y - S_t} + 2\beta\sqrt{(1+r)S_t + (1+n)\tau y} \right\}.$$

The first-order condition with respect to S_t is:

$$-\frac{1}{\sqrt{(1-\tau)y - S_t}} + \frac{\beta\sqrt{1+r}}{\sqrt{(1+r)S_t + (1+n)\tau y}} = 0.$$

Solving for S_t produces:

$$S_t = \frac{\beta^2(1+r)}{1+\beta^2(1+r)}y - \frac{\beta^2(1+r) + \frac{1+n}{1+r}}{1+\beta^2(1+r)}\tau y.$$

Notice that private savings is (as usual) decreasing in τ. Also notice that the larger n is relative to r, the greater this effect. When $n > r$, contributions to the social security system supplant private savings at a greater rate than in a funded system. The reason is because, when $n > r$, the social security system is more attractive than private savings.

Exercise 12.5

1. Grace trades consumption today for consumption tomorrow via schooling S (since there is no bond market). Her maximization problem is:

   $$\max_{S} \left\{ \ln(1-S) + \beta \ln(AS) \right\}.$$

   Recall that $\ln(ab) = \ln(a) + \ln(b)$. Hence the first-order condition is:

   $$-\frac{1}{1-S} + \frac{\beta}{S} = 0.$$

   Solving for Grace's optimal schooling provides $S = \beta/(1+\beta)$. In this setup, $K_1 = S$.

2. Now Grace's problem becomes:

   $$\max_{S} \left\{ \ln(1-S-G) + \beta \ln[A(S+\phi G)] \right\}.$$

   The first-order condition for maximization is:

   $$-\frac{1}{1-S-G} + \frac{\beta}{S+\phi G} = 0.$$

   Thus Grace's optimal schooling choice becomes:

   $$S = \frac{\beta}{1+\beta} - \frac{\beta+\phi}{1+\beta}G.$$

Grace's schooling is certainly decreasing in G (thus investment is, to a certain extent, being crowded out). Grace's human capital is $K_1 = S + \phi G$, so substituting in provides:

$$K_1 = \frac{\beta}{1+\beta} + (\phi - 1)G\frac{\beta}{1+\beta}.$$

Notice that if $\phi < 1$, the government is less efficient at providing schooling than the private sector, and Grace's human capital decreases in G.

3. Now Grace's maximization problem becomes:

$$\max_S \{\ln(1-S) + \beta \ln[A(S+\phi G)]\},$$

since Grace does not have to pay a lump-sum tax in the first period. The first order condition is now:

$$-\frac{1}{1-S} + \frac{\beta}{S+\phi G} = 0.$$

Grace's optimal schooling choice is:

$$S = \frac{\beta}{1+\beta} - \phi G\frac{1}{1+\beta},$$

and her human capital becomes:

$$K_1 = \frac{\beta}{1+\beta}(1+\phi G).$$

Notice that Grace's schooling is still being crowded out, but that her human capital is increasing in G no matter what the value of ϕ, as long as $\phi > 0$.

Exercise 13.1

1. If the agent works, $c^i(\ell^i = 1) = 1 - \tau$, while if the agent does not work, $c^i(\ell^i = 0) = 0$.

2. If the agent works, $u^i(\ell^i = 1) = 1 - \tau - \gamma^i$ while if the agent does not work, $u^i(\ell^i = 0) = 0$. An agent will work if the utility of working is greater than the utility of not working, or if $1 - \tau - \gamma^i \geq 0$.

3. From our previous answer, it is easy to see that $\gamma^*(\tau) = 1 - \tau$.

4. We know that the fraction of agents with γ less than or equal to some number, say γ^*, is just γ^* if $0 \leq \gamma^* \leq 1$. Thus aggregate labor supply as a function of the tax rate is just $\ell(\tau) = \gamma^*(\tau) = 1 - \tau$. On each agent who works, the government collects revenue τ. Thus $\mathcal{T}(\tau) = \tau(1-\tau)$. This is sketched in Figure (13.1).

5. There is a Laffer curve in the tax system.

Exercise 13.2
Briefly, although such a result might be evidence for a Laffer curve, the regression does not control for changes in real income over time. There may not truly be a Laffer curve, but it would look like there was one if real incomes were high when taxes were low and low when taxes were high.

Exercise 13.3
The point of this simple problem was to clear up the difference between the *tax system* $\mathcal{H}(a;\psi)$ and the government's *revenue function* $\mathcal{T}(\psi)$. This problem should also give you some practice in thinking about exemptions.

1. The parameters of the tax system are the choices of the flat tax rate τ and the lump-sum tax S. The household chooses an effort level L in response. Thus $\psi = [\tau, S]$ and $a = L$ here.

2. The tax system $\mathcal{H}(a;\psi)$ maps household actions a and tax system parameters ψ into an amount of tax:

$$\mathcal{H}(L;[\tau, S]) = S + \tau(L - S).$$

   Recall that income directed towards the lump-sum tax S is exempt from the flat tax. We do not consider (yet) that L is itself a function of S and τ.

3. A household's tax bill is always the same as the tax system. In this case, if the household works an amount L it owes $S + \tau(L - S)$.

4. The household's income as a function of L is just L. Hence the household consumes $L - \mathcal{H}(L;[\tau, S])$ or:

$$C = L - [S + \tau(L - S)] = (1 - \tau)(L - S).$$

5. Substituting in to the household's utility function gives:

$$U(C, L) = 2\sqrt{(1 - \tau)(L - S)} - L.$$

   The first-order condition for maximization with respect to L is:

$$\frac{1 - \tau}{L - S} = 1.$$

   (Where did the 2 go?) Solving for L produces:

$$L(\tau, S) = S + 1 - \tau.$$

   This gives the household's optimal response to the tax system $\mathcal{H}$. In the chapter we called this $a_{\max}(\psi)$.

6. The government revenue function is the tax system with the household's action a optimized out. That is:

$$\mathcal{T}(psi) = \mathcal{H}\left[a_{\max}(\psi); \psi\right].$$

In this case, this produces:

$$\mathcal{T}([\tau, S]) = \tau(1-\tau) + S.$$

Notice that there is a Laffer curve (as expected) in the tax parameter τ.

For further practice: Assume that income spent on the lump-sum tax is no longer exempt from the flat tax. How do your answers change? You should be able to show that the Laffer curve in τ vanishes.

Exercise 13.4

1. If the household works ℓ, it raises gross income of ℓ and must pay a tax bill of $\tau\ell$. It consumes the residual, $(1-\tau)\ell$.
2. Substitute $c(\ell, \tau)$ into the household's utility function to find utility purely as a function of labor effort. The household's maximization problem becomes:

$$\max_{\ell}\left\{4\sqrt{(1-\tau)\ell} - \ell\right\}.$$

Taking the derivative with respect to ℓ gives the first order condition for maximization:

$$2\sqrt{\frac{1-\tau}{\ell}} - 1 = 0.$$

We can solve this to find the household's optimal choice of labor effort given taxes, $\ell(\tau)$:

$$\ell(\tau) = 4(1-\tau).$$

3. That government's tax revenue is:

$$\mathcal{T}(\tau) = \tau\ell(\tau) = 4\tau(1-\tau).$$

4. The government wishes to raise revenue of 3/4. We are looking for the tax rate τ that satisfies:

$$\begin{aligned} \mathcal{T}(\tau) &= 3/4, \text{ or:} \\ 4\tau(1-\tau) &= 3/4. \end{aligned}$$

Inspection reveals that there are two such tax rates: $\{1/4, 3/4\}$. Since the government is nice, it will choose the lower tax rate, at which the household consumes more.

Exercise 13.5
In the question, you were allowed to assume that $r = 0$ and that Tammy had an implicit discount factor of $\beta = 1$. These solutions are a little more general. To check your solutions, substitute $\beta = 1$ and $r = 0$.

1. In the first period of life, Tammy earns an income of $y = w\ell$ of which she must pay $w\ell\tau_1$ in taxes. Thus her income net of taxes is $w\ell(1-\tau_1)$. She splits this between consumption in the first period of life, c_1 and savings, b. Thus:

$$\begin{aligned} c_1 + b &\leq (1-\tau_1)w\ell, \text{ and:} \\ c_2 &\leq (1+r)b. \end{aligned}$$

   Notice that $b = c_2/(1+r)$ so we can collapse the two one-period budget constraints into a single present-value budget constraint. Thus:

$$c_1 + \frac{1}{1+r}c_2 \leq (1-\tau_1)w\ell.$$

2. Tammy's Lagrangian is:

$$\mathcal{L}(c_1, c_2, \ell) = \sqrt{c_1} + \beta\sqrt{c_2} - \ell + \lambda\left((1-\tau_1)w\ell - c_1 - \frac{1}{1+r}c_2\right).$$

   This has first-order conditions with respect to c_1, c_2, and ℓ of:

$$\begin{aligned} &\frac{1}{2}\frac{1}{\sqrt{c_1}} - \lambda = 0, \\ &\frac{\beta}{2}\frac{1}{\sqrt{c_2}} - \frac{1}{1+r}\lambda = 0, \text{ and:} \\ &-1 + \lambda w(1-\tau_1) = 0. \end{aligned}$$

   Manipulating each of these equations produces the system:

$$\begin{aligned} c_1 &= \frac{1}{4}\left(\frac{1}{\lambda}\right)^2, \\ c_2 &= \frac{1}{4}\left(\frac{(1+r)\beta}{\lambda}\right)^2, \text{ and:} \\ \frac{1}{\lambda} &= (1-\tau_1)w. \end{aligned}$$

   We can further manipulate these three equations, by substituting out the multiplier λ to find the optimal choices of consumption:

$$\begin{aligned} c_1 &= \frac{1}{4}(1-\tau_1)^2 w^2, \text{ and:} \\ c_2 &= \frac{1}{4}((1+r)\beta)^2(1-\tau_1)^2 w^2. \end{aligned}$$

We can find labor effort ℓ by substituting the optimal consumption decisions (calculated above) into the budget constraint. This will tell us how many hours Tammy must work in order to earn enough (after taxes) to afford to consume c_1, c_2. The budget constraint is:

$$\begin{aligned}
w(1-\tau_1)\ell &= c_1 + \frac{1}{1+r}c_2 \\
&= \frac{w^2(1-\tau_1)^2}{4}\left(1 + \frac{(1+r)^2\beta^2}{(1+r)}\right) \\
&= \frac{w^2(1-\tau_1)^2}{4}(1+(1+r)\beta^2), \text{ so:} \\
\ell &= \frac{w(1-\tau_1)}{4}(1+(1+r)\beta^2).
\end{aligned}$$

Notice that Tammy's effort is strictly decreasing in τ_1 and that at $\tau_1 = 1$, $\ell = 0$. In other words, if the government taxes Tammy to the limit, we expect her not to work at all. This will induce a Laffer curve.

Once we've figured out how much Tammy works, it's an easy matter to deduce how much revenue the government raises by taxing her. The government revenue function here is:

$$\mathcal{T}_1(\tau_1) = \tau_1 w \ell = \frac{w^2\tau_1(1-\tau_1)}{4}[1+(1+r)\beta^2].$$

In terms of τ_1, this is just the equation for a parabola:

$$\mathcal{H}_1(\tau_1) = \tau_1(1-\tau_1)(\text{constant term}).$$

Hence $\tau_1^* = 1/2$, and:

$$\mathcal{H}_1(\tau_1^*) = \left(\frac{1}{4}\right)(\text{constant term}) = \left(\frac{w^2}{16}\right)\left[1+(1+r)\beta^2\right].$$

Thus there is a strict limit on the amount of revenue that the government can squeeze out of Tammy. As the tax rate τ_1 increases, Tammy works less, although if $\tau_1 < 1/2$, the government collects more revenue.

3. There is indeed a Laffer curve in this problem. We should have expected it the instant we saw how Tammy's hours worked, ℓ, responded to the tax rate.

Exercise 13.6

Now Tammy is allowed to deduct savings held over for retirement. This is also known as being able to save in "pre-tax dollars." Almost all employers feature some kind of tax-sheltered savings plan.

1. Tammy's tax bill at the end of period 1 is $\tau_2(wn - b)$. Tammy faces a sequence of budget constraints:

$$\begin{aligned}
c_1 &= (1-\tau_2)(w\ell - b), \text{ and:} \\
c_2 &= (1+r)b.
\end{aligned}$$

Once again, we use the trick of $b = c_2/(1 + r)$, so that Tammy's present-value budget constraint becomes:

$$c_1 = (1 - \tau_2)\left(w\ell - \frac{c_2}{R}\right), \text{ so:}$$

$$c_1 + (1 - \tau_2)\frac{1}{1+r}c_2 = (1 - \tau_2)w\ell$$

Note that as $\tau_2 \to 1$, Tammy's ability to consume in the first period of life goes to zero, but her ability to consume in the second period of life is unchanged.

2. Tammy's Lagrangian is:

$$\mathcal{L}(c_1, c_2, \ell) = \sqrt{c_1} + \beta\sqrt{c_2} - \ell + \lambda\left((1 - \tau_2)w\ell - c_1 - \frac{(1 - \tau_2)}{1+r}c_2\right).$$

The first-order conditions with respect to c_1, c_2 and ℓ are:

$$\frac{1}{2}\frac{1}{\sqrt{c_1}} - \lambda = 0,$$
$$\frac{\beta}{2}\frac{1}{\sqrt{c_2}} - \frac{1 - \tau_2}{1+r}\lambda = 0, \text{ and:}$$
$$-1 + \lambda(1 - \tau_2)w = 0.$$

We can write c_1 and c_2 easily as a function of λ:

$$c_1 = \frac{1}{4\lambda^2}, \text{ and:}$$
$$c_2 = \frac{\beta^2}{4\lambda^2}\left(\frac{1+r}{1-\tau_2}\right)^2.$$

So it's an easy matter to substitute out for λ and calculate optimal consumption c_1, c_2. Thus:

$$c_1 = \frac{1}{4}w^2(1 - \tau_2)^2, \text{ and:}$$
$$c_2 = \frac{(\beta(1+r))^2}{4}w^2.$$

Notice that c_2 does not depend on τ_2. We can substitute in the optimal consumptions above into the budget constraint to determine how many hours Tammy has to work to be able to afford her optimal consumption plan:

$$(1 - \tau_2)w\ell = c_1 + (1 - \tau_2)\frac{1}{1+r}c_2$$
$$= \frac{w^2(1 - \tau_2)^2}{4} + (1 - \tau_2)\frac{\beta^2(1+r)w^2}{4}, \text{ so:}$$
$$\ell = \frac{w}{4}[(1 - \tau_2) + (1 + r)\beta^2].$$

Notice from the first line above that:

$$w\ell = \frac{c_1}{1-\tau_2} + b.$$

The trick here is to substitute back into the right budget constraint. The first couple of times I did this I substituted back into the budget constraint from Exercise (13.5) and got all sorts of strange answers. Notice that Tammy always consumes a certain amount of c_2, no matter what τ_2 is, so she always works a certain amount. However this may not overturn the Laffer curve since she is paying for c_2 with pre-tax dollars.

3. Once again, this is a bit tricky. Remember that Tammy's tax bill is $\tau_2(w\ell - b)$ and that $b = c_2/(1+r)$. Thus:

$$\begin{aligned}
\mathcal{T}_2(\tau_2) &= \tau_2(w\ell - b) \\
&= \tau_2\left(w\ell - \frac{c_2}{1+r}\right) \\
&= \tau_2\left(w\ell - \frac{\beta^2(1+r)}{4}w^2\right) \\
&= \tau_2\left[w\left(\frac{w}{4}(1-\tau_2) + \frac{w}{4}\beta^2(1+r)\right) - \beta^2(1+r)\frac{w^2}{4}\right] \\
&= \tau_2(1-\tau_2)\left(\frac{w^2}{4}\right).
\end{aligned}$$

This was a matter of remembering to substitute into the right revenue equation. Although Tammy always works at least enough to finance a certain amount of consumption while old, this amount of income is tax-deductible, so the government can't get at it.

As $\tau_2 \to 1$, government revenue goes to zero, as before. The maximizing tax rate, τ_2^*, is $\tau_2^* = 0.5$ and the maximum amount of revenue that government can raise is:

$$\mathcal{T}_2(\tau_2^*) = \frac{w^2}{16}.$$

Notice that $\mathcal{T}_2(\tau_2) < \mathcal{T}_1(\tau_1)$.

4. There is still a Laffer curve present. Unfortunately for the government, tax revenue is now lower.

5. Our answers are indeed different. Because Tammy is able to shelter some of her income from the government, total tax revenue will be lower.

Exercise 14.1

1. About \$5.2 trillion/\$7 trillion.

2. About 1.07 in 1945 (Barro p.362).

3. About 0.38 in 1981 (Barro p.341).

4. 3%.

Exercise 14.2
The consumer's problem is thus:

$$\max_{x_1,x_2} \{\ln(x_1) + x_2\}, \text{ subject to:}$$
$$(p_1 + t_1)x_1 + (p_2 + t_2)x_2 = M.$$

The two first-order conditions for this problem are:

$$\frac{1}{x_1} = \lambda(p_1 + t_1), \text{ and:}$$
$$1 = \lambda(p_2 + t_2).$$

These first order conditions plus the budget constraint can be used to solve for the three unknowns λ, x_1, and x_2 in terms of the givens in the problem, p_1, p_2, t_1, t_2, and M. Solving:

$$x_1 = \frac{p_2 + t_2}{p_1 + t_1}.$$
$$x_2 = \frac{M}{p_2 + t_2} - 1.$$

The government's revenue function can be calculated accordingly:

$$\mathcal{T}(t_1, t_2, p_1, p_2, M) = t_1x_1 + t_2x_2 = \frac{t_1(p_2 + t_2)}{p_1 + t_1} + \frac{t_2M}{p_2 + t_2} - t_2.$$

Substitute the above demand functions for x_1 and x_2 into the objective function to obtain the household's indirect utility:

$$V(p_1 + t_1, p_2 + t_2, M) = \ln\left(\frac{p_2 + t_2}{p_1 + t_1}\right) - \frac{M}{p_2 + t_2} + 1.$$

So at this point, we have found the government's revenue function which tells us how much the government can raise from taxes *given* that consumers respond optimally to the given tax rates. By deriving the consumer's indirect utility function, we know how consumers compare different tax rates and income levels in utility terms.

Potentially the government is faced with the need to raise a certain level of revenue, G. It can raise this revenue a number of different ways by taxing the two goods in different amounts with the constraint that in the end, it must have raised G in revenues.

A benevolent government could decide to choose the combination of taxes (t_1, t_2) such that consumer utility is maximized, subject to the constraint that it raises the necessary revenue G. The government's optimal-tax problem would then be:

$$\max_{t_1,t_2} V(p_1 + t_1, p_2 + t_2, M), \text{ subject to:}$$
$$\mathcal{T}(t_1, t_2, p_1, p_2, M) = G,$$

and some given M. Make sure you understand the intuition of this problem. Both the indirect utility function and the government revenue function account for the fact that households respond optimally to the given tax policy. Before you ever write down the optimal tax problem, you must know how consumers will respond to any possible tax policy given by (t_1, t_2). Implicit in the indirect utility function and government revenue function is the fact that consumers are responding optimally to their environment.

Exercise 14.3

The key to this problem is realizing that the household's budget set will be kinked at the point $\{y_1 - \mathcal{T}_1, y_2 - \mathcal{T}_2\}$. For points to the left of this kink, the household is saving, and the budget set is relatively flat. For points to the right of this kink, the household is borrowing, and the budget set is relatively steep. The government's optimal plan will be to levy very low taxes initially and then high taxes later, in essence borrowing on behalf of the household.

1. If the household neither borrows nor lends, it consumes:

$$c_1 = y_1 - \mathcal{T}_1, \text{ and:}$$
$$c_2 = y_2 - \mathcal{T}_2.$$

   This is the location of the kink in the budget constraint: to consume more in period $t = 1$ than $y_1 - \mathcal{T}_1$, it will have to borrow at the relatively high rate r' and the budget set will have a slop of $-(1+r')$. The government will be able to move the kink around, increasing or decreasing the number of points in the household's budget set.

2. For convenience, all of the answers to the next three questions are placed on the same set of axes (below). The solid line gives the answer to the first question.

3. The dotted line gives the answer to this question. Notice that the household has more points to choose from.

4. The dashed line gives the answer to this question. Notice that the household has fewer points to choose from.

5. The government chooses $\mathcal{T}_1 = 0$, in essence borrowing at the low interest rate r on behalf of the household.

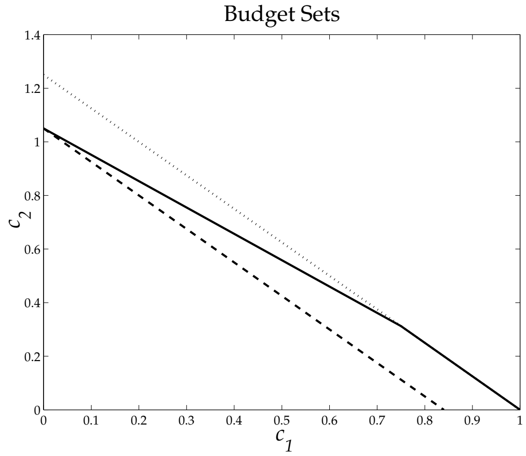

Exercise 14.4
In this question the government runs a deficit of unity in the first period (period $t = 0$), because expenditures exceed revenues by exactly unity. In all subsequent periods, government revenues just match direct expenditures in each period, but are not enough to repay the interest cost of the initial debt. As a result, the government will have to continually roll over its debt each period. Under the proposed plan, the government has not backed the initial borrowing with any future revenues, so it does not ever intend to repay its debt. From the government's flow budget constraint, assuming that $B^g_{-1} = 0$:

$$\begin{aligned} B^g_0 &= G_0 - \mathcal{T}_0 = 1. \\ B^g_1 &= 1 + r. \\ B^g_2 &= (1 + r)^2. \\ &\vdots\ \vdots \\ B^g_t &= (1 + r)^t. \end{aligned}$$

So the government debt level is exploding. Substituting in to the transversality condition, we get:

$$\lim_{t\to\infty} (1 + r)^{-t} B^g_t = \lim_{t\to\infty} (1 + r)^{-t} (1 + r)^t = 1.$$

Since the limit does not equal zero, we see that the government's debt plan does not meet the transversality condition.

Exercise 15.1
The Lagrangean is:

$$\mathcal{L} = (c^P_w)^\gamma (c^P_b)^{1-\gamma} + \lambda[m^P - c^P_w p_w + c^P_b p_b].$$

The first-order conditions are:

(FOC c_w^P) $\gamma(c_w^{P\star})^{\gamma-1}(c_b^{P\star})^{1-\gamma} + \lambda^\star[-p_w] = 0,$ and:

(FOC c_b^P) $\gamma(c_w^{P\star})^{\gamma}(1-\gamma)(c_b^{P\star})^{-\gamma} + \lambda^\star[-p_b] = 0.$

Combining these to get rid of $\lambda^\star$ yields:

$$\frac{\gamma(c_w^{P\star})^{\gamma-1}(c_b^{P\star})^{1-\gamma}}{p_w} = \frac{\gamma(c_w^{P\star})^{\gamma}(1-\gamma)(c_b^{P\star})^{-\gamma}}{p_b}, \text{ or:}$$

$$\frac{p_w}{p_b} = \frac{\gamma c_b^{P\star}}{(1-\gamma)c_w^{P\star}}.$$

Solving this for $c_b^{P\star}$ and plugging back into the budget equation gives us:

$$c_w^{P\star}p_w + \left[\frac{(1-\gamma)p_w c_w^{P\star}}{\gamma p_b}\right] p_b = m^P.$$

After some algebra, we get the first result:

$$c_w^{P\star} = \frac{\gamma m^P}{p_w}.$$

When we plug this back into the budget equation and solve for $b^{P\star}$, we get the other result:

$$c_b^{P\star} = \frac{(1-\gamma)m^P}{p_w}.$$

Exercise 15.2

1. See Figure S.7.

2. See Figure S.7.

3. From the graph, we know that $p_w/p_b = 3$. Suppose the relative price is less than 3. Then only Pat will make wine, and supply will be 4 jugs. Plugging 4 jugs into the demand function gives a relative price of $p_w/p_b = 7/2$, but $7/2 > 3$, which is a contradiction, so the equilibrium relative price can't be less than 3.

   By a similar argument, you can show that the equilibrium relative price can't be more than 3.

4. Pat makes (i) 4 jugs of wine and (ii) 0 jugs of beer. Chris makes (iii) 1 jug of wine and (iv) 1 jug of beer.

5. Pat has an absolute advantage in wine production, since $2 < 6$.

6. Pat has a comparative advantage in wine production, since $2/1 < 6/2$. Chris makes wine anyway, since the equilibrium price is so high.

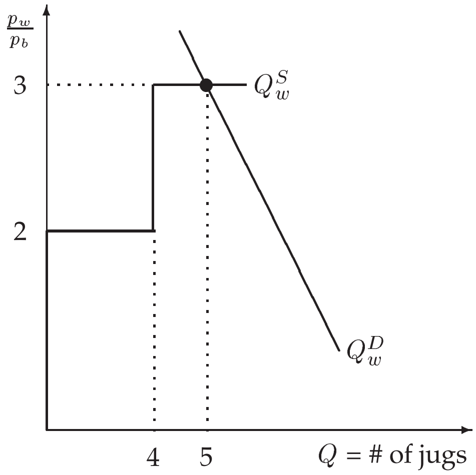

Figure S.7: The Supply of Wine by Pat and Chris Together

Exercise 17.1
We know that:

$$1 + r_2 = F\frac{1 - \theta(1 + r_1)}{1 - \theta}.$$

We can manipulate this to produce:

$$r_2 = F - 1 - \frac{\theta}{1 - \theta}Fr_1.$$

This is an interesting result. Essentially, r_2 is the net return on turnips held until period $t = 2$ (that is the $F - 1$ term) minus a risk premium term that is increasing in r_1.

Exercise 17.2

1. The bank's assets are the value of the loans outstanding net of loss reserve, in other words, the expected return on its loans. The bank's liabilities are the amount it owes its depositors. Considering that the bank must raise a unit amount of deposits to make a single loan, this means that the bank must pay $1 + r$ to make a loan. Thus the bank's expected profits are:

$$\alpha p_S x + (1 - \alpha)p_R x - (1 + r).$$

   That is, the bank gets x only if the borrower does not default.

2. Now we solve for the lowest value of x which generates non-negative expected prof-

its. Setting expected profits to zero produces:

$$\alpha p_S x + (1 - \alpha)p_R x = (1 + r), \text{ so:}$$
$$x^\star(r, \alpha) = \frac{1 + r}{\alpha p_S + (1 - \alpha)p_R}.$$

Since safe borrowers repay more frequently than risky ones because:

$$p_S > p_R,$$

the amount repaid, x, is decreasing as the mix of agents becomes safer, that is, as α increases. As expected, x is increasing in the interest rate r.

3. We assumed that agents are risk neutral. Thus if the project succeeds (with probability p_S), a safe agent consumes $\pi_S - x^\star(r, \alpha)$ and a risky agent (with probability p_R) consumes $\pi_R - x^\star(r, \alpha)$. If the project fails, agents consume nothing. Their expected utilities therefore are:

$$V_S(r) = p_S[\pi_S - x^\star(r, \alpha)], \text{ and:}$$
$$V_R(r) = p_R[\pi_R - x^\star(r, \alpha)].$$

Where $x^\star(r, \alpha)$ is the equilibrium value of x. Notice that since $p_S\pi_S = p_R\pi_R$ and $p_S > p_R$ that we can write V_S and V_R as:

$$V_R(r) = V_S(r) + (p_S - p_R)x^\star(r, \alpha).$$

Thus at any given interest rate $r > 0$, the expected utility of risky borrowers is greater than the expected utility of safe borrowers, $V_R(r) > V_S(r)$.

4. Since $V_R(r) > V_S(r)$, it is easy to see that if $V_S(r) > 0$ then $V_R(r)$ must also be greater than zero. Next we find $r^\star$ such that $V_S(r^\star) = 0$. Substituting:

$$0 = V_S(r^\star) = p_S[\pi_s - x^\star(r^\star, \alpha)] = p_S\pi_S - (1 + r^\star)\frac{p_S}{\alpha p_S + (1 - \alpha)p_R}, \text{ so:}$$
$$1 + r^\star = \pi_S + \alpha\pi_S(p_S - p_R).$$

At interest rates above $r^\star$ all safe agents stop borrowing to finance their projects. Realizing this, banks adjust their equilibrium payments to: $x^\star(r, \alpha = 0)$, so $(1 + r)/p_R$.

Exercise 17.3

If the revenue functions $\pi(x, \gamma)$ all shift up by some amount, then, for any given interest rate r, intermediaries can make loans to agents with higher audit costs. That is, $\gamma^\star(r)$ also shifts up as a result. This shifts the demand for capital up and out, but leaves the supply schedule untouched. As a result the equilibrium interest rate increases, as does the equilibrium quantity of capital saved by type-1 (worker) agents. As a result, type-1 agents work harder, accumulate more capital and more type-2 (entrepreneurial) agents' projects are funded so aggregate output goes up. Type-1 agents are made better off by the increase in the interest rate because their consumption goes up (although they are working harder too). Type-2 agents who had been credit rationed are made better off, but type-2 agents who previously had not been credit rationed are made worse off because the interest rate paid on their loans goes up.

Exercise 17.4

This question uses slightly different notation from that used in this chapter. Most bothersome is probably the fact that r here denotes the *gross* interest rate, which elsewhere is denoted $1+r$. This question is a reworking of the model of moral hazard from this chapter. This question was taken directly from the Spring 1998 Econ 203 final exam.

1. A rich Yalie can finance the tuition cost of Yale from her own wealth (that is, $a > 1$). If she gets the good job, she consumes $w + r(a - 1)$, if she does not, she consumes $r(a - 1)$. Hence her maximization problem is:

$$\max_{\pi} \left\{ \pi[w + r(a-1)] + (1-\pi)[r(a-1)] - \frac{w}{\alpha}\frac{\pi^2}{2} \right\}.$$

   The first-order condition with respect to π is:

$$w - \frac{w}{\alpha}\pi = 0.$$

   We can easily solve this to find that $\pi = w$.

2. Poor Yalies are required to repay an amount x only if they land the good job. Hence if they land the good job, they consume $w - x$, while if they go unemployed, they consume 0. Thus their optimization problem may be written as:

$$\max_{\pi} \left\{ \pi(w - x) + (1-\pi)\cdot 0 - \frac{w}{\alpha}\frac{\pi^2}{2} \right\}.$$

   The first-order condition with respect to effort π is now:

$$w - x - \frac{w}{\alpha}\pi = 0.$$

   We can solve this to find the optimal effort as a function of repayment amount:

$$\pi(x) = \alpha\left(1 - \frac{x}{w}\right).$$

   Notice that effort is decreasing in x.

3. Yale University must also pay r to raise the funds to loan to its students. If it is making these loans out of its endowment, then it is paying an opportunity cost of r. A student of wealth $a < 1$ needs a loan of size $1 - a$, which costs Yale an amount $r(a - 1)$. Thus Yale's profit on this loan is:

$$x\pi(x) + 0\cdot[1 - \pi(x)] - r(a-1).$$

   But we know $\pi(x)$ from the previous question, so:

$$x\alpha\left(1 - \frac{x}{w}\right) - r(a-1).$$

   This is the usual quadratic in x.

4. Yale's "fair lending policy" guarantees that *all* borrowers pay the same interest rate, regardless of wealth. Since we know $\pi(x)$ from above, and we are given $x(a)$, it is an easy matter to calculate $\pi(a)$:

$$\pi(a) = \alpha - \frac{r}{w}(1-a).$$

Notice that effort is decreasing in r and increasing in wealth a.

5. Here we are supposed to show that $\pi(a) \leq \pi^{\star}$ from above, where $\pi^{\star} = \alpha$. If $a < 1$ then $1 - a > 0$, and $r > 0$ and $w > 0$ by assumption. It's easy to see that this must be true.

6. Now we are supposed to show that Yale's profits are negative on loans and that poor borrowers cost it more than richer borrowers. The fair lending policy charge all borrowers the same interest rate. Further, this interest rate guarantees Yale zero profits assuming that they exert α effort. Poor borrowers will exert less than α effort, and so Yale will lose money. Return to Yale's profit function:

$$\pi(a)x(a) - r(1-a).$$

Substituting in for $\pi(a)$ and $x(a)$ we get:

$$\left[\alpha - \frac{r}{w}(1-a)\right]\left[\frac{r(1-a)}{\alpha}\right] - r(1-a).$$

We can manipulate this to produce:

$$-\frac{[r(1-a)]^2}{\alpha w}.$$

All other terms canceled out. This is certainly negative, and increasing in a. Thus Yale loses no money on "borrowers" of wealth $a = 1$, and loses the most money on borrowers of wealth $a = 0$.

Exercise 18.1
This question has been given on previous problem sets. In particular, we have amassed a few years' data on students currency holding habits.

1. According to Friedman and Schwartz, the stock of money fell 33% from 1929 to 1933. Household holdings of currency increased over the period.

2. Real income fell by 36% over the same period and prices decreased.

3. From the Barro textbook: Real interest rates have been negative in the years 1950-51, 1956-57 and 1973-79. Inflation was negative in 1949 and 1954.

4. From the Barro textbook: There is evidence in looking cross-sectionally at different countries that changes in money stocks are positively correlated with changes in prices, or inflation. Long run time-series evidence demonstrates a positive correlation between money growth and inflation as well.

5. From the Porter article on the location of U.S. currency: The stock of Federal Reserve notes outside of banks (vault cash) at the end of 1995 was about \$375 billion, or about \$1440 per American. Nobody had quite this much cash on them, although some students were carrying over \$100. I assume these students were well trained in self-defense. According to Porter, between \$200 and \$250 billion, that is, more than half, was abroad, primarily in the former Soviet Union and South America.

6. Generally people keep their money in low interest assets because they are liquid and provide transactions services. It's tough to buy lunch with shares of GM stock rather than Hyde Park bank checks.

7. Sargent states that inflation can seem to have momentum if people have persistent expectations that the government will continue to pursue inflationary fiscal and monetary policies.

8. Since currency is a debt of the government, whenever the government prints money, it is devaluing the value of its debt. This is a form of taxation and the value by which its debt is reduced is called seignorage. The government obtained \$23 billion in seignorage in 1991.

9. The quantity theory is the theory that the stock of money is directly related to the nominal value of output in the economy. It is usually written as the identity:

$$M = PY/V$$

   where M is the money stock, P is the price level, and Y is the real amount of output. It is an accounting identity in that the velocity of money, V, is defined residually as whatever it takes to make the above identity true.

10. A gold standard is a monetary system where the government promises to exchange dollars for a given amount of gold. If the world quantity of gold changes (for example, gold is discovered in the Illinois high country) then the quantity of money also changes. Our current monetary system is a fiat system, where money isn't backed by any other real asset. It is simply money by "fiat".

Exercise 18.2
Government austerity programs involve reducing government expenditures and increasing tax revenue. Both cause immediate and obvious dislocations. Governments typically reduce spending by firing lots of government workers, closing or privatizing loss-making government-owned industries and reducing subsidies on staples like food and shelter. Governments increase revenue by charging for previously-free services and pushing up the tax rates. From the point of view of a typical household, expenses are likely to go up

while income is likely to fall. Thus austerity programs can indeed cause immediate civil unrest.

On the other hand, we know that subsidies are a bad way to help the poor (since most of the benefit goes to middle-class and rich households), that state-owned businesses tend to be poorly run, depressing the marginal product of workers and tying up valuable capital and that bloated government bureaucracies are rarely beneficial.

Leave all this to one side: the fact is that no government willingly embarks on an austerity program. They only consider austerity when they are forced to choose between austerity and hyperinflation. Like Germany in 1921, an austerity program has to be seen as better than the alternative, hyperinflation. The central European countries in the early 1920s tried both hyperinflation and austerity, and found austerity to be the lesser of the two evils. That early experience has since been confirmed by a host of different countries. Austerity may indeed be painful, but it is necessary in the long run and better than hyperinflation.

Exercise 18.3

1. We know that the money supply must evolve to completely cover the constant per-capita deficit of d. So we know that:

$$\frac{M_t - M_{t-1}}{P_t} = D_t = dN_t. \tag{S.10}$$

We know from the Quantity Theory of Money given in the problem that:

$$P_t = \frac{M_t}{Y_t} = \frac{M_t}{N_t}. \tag{S.11}$$

Thus we can put equations (S.10) and (S.11) together to produce:

$$dN_t = \frac{M_t - M_{t-1}}{\frac{M_t}{N_t}} = N_t \frac{M_t - M_{t-1}}{M_t}, \text{ so:}$$

$$d = \frac{M_t - M_{t-1}}{M_t}, \text{ so:}$$

$$1 - \frac{M_{t-1}}{M_t} = d, \text{ and:}$$

$$\frac{M_t}{M_{t-1}} = \frac{1}{1-d}. \tag{S.12}$$

Thus $M_t = [1/(1-d)]M_{t-1}$. This gives us an expression by how much the total stock of money must evolve to raise enough seignorage revenue to allow the government to run a constant per-capita deficit of d each period.

2. To answer this question we will use the quantity-theoretic relation, equation (S.11) above and the effect of d on the evolution of money in equation (S.12) above to a find a value for d at which prices are stable, that is, at which $P_t = P_{t-1}$. Notice that:

$$\frac{P_t}{P_{t-1}} = \frac{M_t/N_t}{M_{t-1}/N_{t-1}} = \frac{N_{t-1}}{N_t}\frac{M_t}{M_{t-1}} = \frac{1}{1+n}\frac{M_t}{M_{t-1}} = \frac{1}{1+n}\frac{1}{1-d}. \tag{S.13}$$

If $P_t/P_{t-1} = 1$ then, continuing from (S.13):

$$\frac{1}{1+n}\frac{1}{1-d} = 1, \text{ so:}$$

$$d = \frac{n}{1+n}. \tag{S.14}$$

By (S.14) we see that the government can run a constant per capita deficit of $n/(1+n)$ by printing money and not cause any inflation, where n is the growth rate of the economy/population (they are the same thing in this example).

3. As $n \to 0$, the non-inflationary deficit also goes to zero. At $n = 1$ (the economy doubles in size every period) the non-inflationary per-capita deficit goes to 1/2. That is, the government can run a deficit of 50% of GDP by printing money and not cause inflation. At the supplied estimate of $n = 0.03$, the critical value of d is 0.03/1.03 or about 0.029 or 2.9% of GDP.

4. From equation (S.13) above, if $d = 0$ then:

$$\frac{P_t}{P_{t-1}} = \frac{1}{1+n} < 1, \text{ so:}$$
$$P_t = \frac{1}{1+n}P_{t-1}, \text{ and:}$$
$$P_t < P_{t-1}.$$

So there will be deflation over time—prices will fall at the rate n.

Exercise 18.4

Although we will accept a variety of answers, I will outline briefly what we were looking for. As with the central European countries in 1921-23, Kolyastan is politically unstable and in economic turmoil. Many of the same policies that worked in those countries should also work in Kolyastan. The government should move quickly to improve its tax collection system and radically decrease spending. This will probably mean closing down state-run factories and ending subsidies. The argument, often advanced, that such direct measures will hurt the citizens ignores the fact that the people are already paying for them through the inefficient means of the inflation tax. With its fiscal house in order, the government should reform the monetary sector by liberating the central bank, appointing a dour old man to be its head and undertaking a currency reform. For these changes to be credible, Kolyastan must somehow commit not to return to its bad old ways. It could do so by signing treaty agreements with the IMF, World Bank or some other dispassionate outside entity. Further, it should write the law creating the central bank in such a way that it is more or less independent from transitory political pressures. The bank ought to be prohibited from buying Kolyastani Treasury notes.

Exercise 19.1

1. True: The CPI calculates the change in the price of a market-basket of goods over fairly short time periods. If one element of that basket were to increase in price dra-

matically, even if they were compensated enough to buy the new market basket, consumers would choose one with less of the newly-expensive good (substituting away from it).

2. Inflation is bad because it leads consumers to undertake a privately useful but socially wasteful activity (economizing on cash balances). The Fed cannot effectively fight inflation with short-term actions, it must maintain a long-term low-inflation regime.

Exercise 19.2
The slope of the Phillips curve gives the relative price (technological tradeoff) between inflation and unemployment. If inflationary expectations are fixed, the government can achieve a higher utility if it does not have to accept more inflation for lower unemployment. In other words, if the Phillips curve is flatter. It is interesting to note that a perfectly flat Phillips curve would mean that unemployment was purely a choice of the government and did not affect inflation at all. If the government and the private sector engage in a Nash game, the Nash outcomes inflation rate is directly proportional to γ, so low values of γ mean lower Nash inflation.

Exercise 19.3
The point of this question was bested summed up by Goethe in Faust. His Mephistopheles at one point describes himself: “That Power I serve / Which wills forever Evil / And does forever good.” Or as Nick Lowe put: “You’ve got to be cruel to be kind.” The higher ϕ is the higher the inflation rate, but unemployment is only marginally lower (depending on expectations).

1. The government’s maximization problem is:

$$\max_{\pi} \left\{ -\phi(u^* + \gamma\pi^e - \gamma\pi)^2 - \pi^2 \right\}.$$

We can solve this to find:

$$\pi_0(\phi) = \frac{\phi\gamma}{1+\phi\gamma^2}u^* + \frac{\phi\gamma^2}{1+\phi\gamma^2}\pi^e.$$

Thus the optimal inflation choice is increasing in ϕ.

2. The corresponding unemployment rate is:

$$u_0(\phi) = \frac{1}{1+\phi\gamma^2}u^* + \frac{\gamma}{1+\phi\gamma^2}\pi^e.$$

3. Now we assume that government continues to take expectations as fixed, but that the private sector adjusts its expectations so that they are perfectly met. Recall that, given expectations π^e, the government’s optimal inflationary response is:

$$\pi = \frac{\phi\gamma}{1+\phi\gamma^2}u^* + \frac{\phi\gamma^2}{1+\phi\gamma^2}\pi^e.$$

Now define π_1 as:

$$\pi_1 = \frac{\phi\gamma}{1+\phi\gamma^2}u^* + \frac{\phi\gamma^2}{1+\phi\gamma^2}\pi_1.$$

We can solve this for π_1 to find:

$$\pi_1 = \phi\gamma u^*.$$

The associated unemployment rate is $u_1 = u^*$, since $\pi^e = \pi$ in this case.

4. Given that agents form expectations rationally, eventually π^e will converge to π. If the government is playing Ramsey (because it has a commitment device), then $\pi = 0$ and $u = u^*$ no matter what ϕ is. If the government is playing Nash, then unemployment is still at the natural rate, $u = u^*$, but inflation is $\pi = \phi\gamma u^*$. Thus the lower the value of ϕ, the lower the Nash inflation rate. The point of this question is that if $\phi = 0$, the Nash and Ramsey inflation rates coincide. Having $\phi = 0$ is an effective device with which to commit to low inflation.

Exercise 19.4
Think of the dynamics in this question as sliding along the government's best response curve, as depicted in Figure 19.2. Expectations will creep up, always lagging behind actual inflation, until the gap between the two vanishes and the private sector expects the Nash inflation, and the government (of course) delivers it.

1. In period t, given inflationary expectations π_t^e, the government solves:

$$\max_{\pi_t} \left\{-\{u^\star + \gamma(\pi_t^e - \pi_t)\}^2 - \pi_t^2\right\}.$$

The government's optimal choice is:

$$\pi_t^\star(\pi_t^e) = \frac{\gamma}{1+\gamma^2}u^\star + \frac{\gamma^2}{1+\gamma^2}\pi^e.$$

2. Since expectations are just last period's inflation rate, and since we know that the government inflation policy rule is given by $\pi_t^\star$ above, the dynamics of the system are given by the pair of equations:

$$\begin{aligned}\pi_t &= A(u^\star + \gamma\pi_t^e), \text{ for all } t = 0,1,\dots,\infty, \text{ and:}\\ \pi_t^e &= \pi_{t-1}, \text{ for all } t = 1,2,\dots,\infty.\end{aligned}$$

Recall that initial inflationary expectations are $\pi_0^e = 0$. We can substitute out the expectations term to produce a single law of motion in inflation:

$$\pi_t = Au^\star + \gamma A\pi_{t-1}, \text{ for all } t = 1,2,\dots,\infty.$$

For notational convenience we have defined $A = \gamma/(1+\gamma^2)$.

3. Since expectations start at zero, the first period's inflation rate is:

$$\pi_0 = Au^\star.$$

Where is as defined above. Thus in the first few periods inflation evolves as:

$\pi_0 = Au^\star.$

$\pi_1 = Au^\star + \gamma A\pi_0 = Au^\star + \gamma A^2 u^\star = Au^\star(1 + \gamma A).$

$\pi_2 = Au^\star + \gamma A\pi_1 = Au^\star + (1 + \gamma A)\gamma A^2 u^\star = Au^\star(1 + \gamma A + (\gamma A)^2).$

$\pi_3 = Au^\star + \gamma A\pi_2 = Au^\star + \gamma Au^*(1 + \gamma A + (\gamma A)^2) = Au^\star(1 + \gamma A + (\gamma A)^2 + (\gamma A)^3).$

The pattern ought to be pretty clear. In general, inflation in period t will be:

$$\pi_t = Au^\star \sum_{i=0}^{t} (\gamma A)^i.$$

So as time moves forward, we have:

$$\lim_{t\to\infty} \pi_t = Au^\star \sum_{i=0}^{\infty} (\gamma A)^i.$$

We can solve the summation using the geometric series to get:

$$\lim_{t\to\infty} \pi_t = \frac{A}{1 - \gamma A} u^\star.$$

Recall that we defined A to be:

$$A = \frac{\gamma}{1 + \gamma^2}.$$

So we can further simplify to get:

$$\lim_{t\to\infty} \pi_t = \gamma u^\star.$$

This is just the Nash inflation rate. Expectations are also converging to this level, so at the limit, unemployment will also be at the Nash level of the natural rate $u^\star$.

Given that inflationary expectations were initially low, the government was able to surprise the private sector and push unemployment below its natural level. Over time the private adapted its expectations and as expected inflation rose, so did unemployment. Thus the time paths of inflation and unemployment are both rising over time, until they achieve the Nash level.

4. The steady-state levels of inflation and unemployment are not sensitive to the initial expected inflation. If the private sector were instead anticipating very high inflation levels at the beginning of the trajectory, the government would consistently produce surprisingly low inflation levels (but still above the Nash level) and the unemployment rate would be above its natural rate. Over time both inflation and unemployment would fall to their Nash levels.

5. The government's optimal choice of inflation in period t, π_t, now becomes:

$$\pi_t = A(u^\star + \gamma\pi_{t-1} + \gamma\varepsilon_t), \text{ for all } t = 1, 2, \ldots, \infty.$$

Since the shock term is mean zero, over time we would expected the inflation rate to settle down in expectation to the same level as before, although each period the shock will push the inflation rate above or below the Nash level. In Figure (c19:fa3) we plot the mean and actual trajectories for inflation and unemployment.

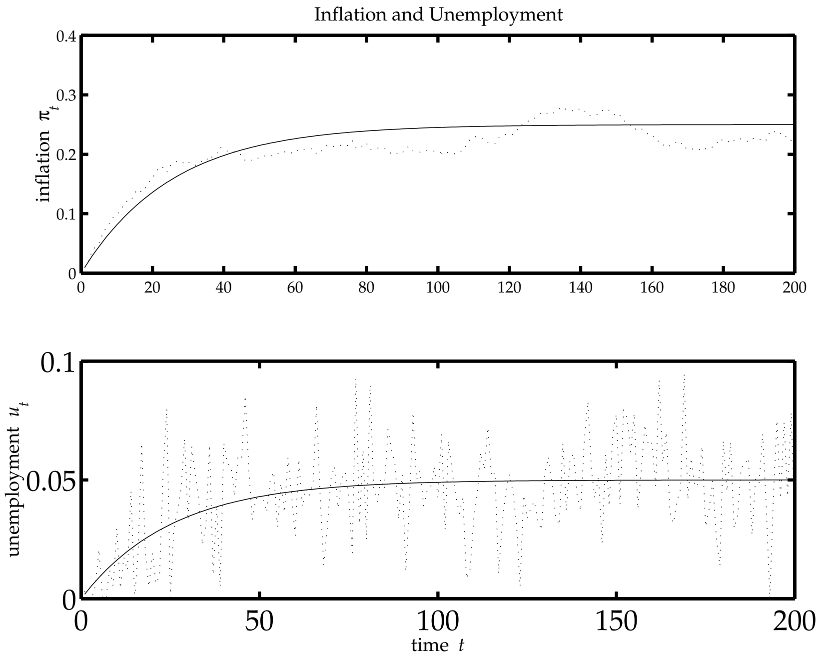

Figure S.8: The dotted line gives the actual time paths for inflation and unemployment with adaptive expectations when there is a mean-zero i.i.d. Normal shock to the Phillips curve, while the solid lines give the same thing with the shock turned off.

Exercise 19.5
As in the previous question, the dynamics of expectations and inflation are given by the system:

$$\pi_t = A(u^\star + \gamma\pi_t^e), \text{ for all } t = 0, 1, \ldots, \infty, \text{ and:}$$
$$\pi_t^e = \delta\pi_{t-1}, \text{ for all } t = 1, 2, \ldots, \infty.$$

Recall that initial inflationary expectations are defined to be $\pi_0^e = 0$. Again, the term A is defined to be:

$$A = \frac{\gamma}{1 + \gamma^2}.$$

We can substitute out the expectations term above to determine the law of motion for inflation:

$$\pi_t = Au^\star + \delta A\gamma\pi_{t-1}, \text{ for all } t = 1, 2, \dots, \infty.$$

Eventually this will converge to a steady-state level of inflation, at which $\pi_{t+1} = \pi_t = \pi_1$. Substituting in:

$$\pi_1 = Au^\star + \delta\gamma A\pi_1.$$

Solving for π_1 produces:

$$\pi_1 = \frac{A}{1-\delta\gamma A}u^\star.$$

The associated inflation rate, u_1, is:

$$u_1 = \frac{1-\gamma A}{1-\delta\gamma A}u^\star.$$

Notice that if $\delta = 1$ this is just the normal Nash outcome. As δ moves closer to zero, so that the private sector puts more and more weight on the government's (utterly mendacious) announcement, inflation and unemployment both fall.