

3 Optimizing the Saving Decision in a Growth Model

Chapter 3 Contents

| | |
|--|-----------|
| A. Topics and Tools | 2 |
| B. Introducing Dynamic Utility Functions | 3 |
| C. Discounting the Future in Discrete and Continuous Time | 4 |
| <i>The idea of discounting</i> | 4 |
| <i>Frequency of compounding and present value</i> | 6 |
| <i>Discounting money vs. discounting utility</i> | 8 |
| <i>Adding up values in continuous time using integrals</i> | 9 |
| <i>Discounting utility in continuous time</i> | 10 |
| D. Constrained Maximization: The Lagrangian | 11 |
| E. Using Indifference Curves to Understand Intertemporal Substitution | 13 |
| F. Understanding Romer's Chapter 2, Part A | 17 |
| <i>Household vs. individual utility</i> | 17 |
| <i>Choosing a functional form for the utility function</i> | 18 |
| <i>Consumption smoothing</i> | 19 |
| <i>Discounting with time-varying interest rates: $R(t)$ and $r(t)$</i> | 21 |
| <i>The positivity restriction on $\rho - n - (1 - \theta)g$</i> | 22 |
| <i>Deriving the Ramsey consumption-equilibrium equation</i> | 22 |
| <i>Understanding the Ramsey consumption-equilibrium equation</i> | 23 |
| <i>The steady-state balanced-growth path in the Ramsey model</i> | 25 |
| <i>Saddle-path convergence to the steady state</i> | 27 |
| <i>Ramsey equilibrium is Pareto optimal</i> | 28 |
| G. Understanding Romer's Chapter 2, Part B | 29 |
| <i>Consumer behavior in Diamond's overlapping-generations model</i> | 29 |
| <i>Steady-state equilibrium in the Diamond model</i> | 31 |
| <i>Welfare analysis in the Diamond model</i> | 31 |
| H. Government Spending in Growth Models | 33 |
| <i>The effects of government purchases</i> | 33 |
| I. Suggestions for Further Reading | 37 |
| <i>Original expositions of the models</i> | 37 |
| <i>Alternative presentations and mathematical methods</i> | 37 |
| I. Work Cited in Text | 37 |

A. Topics and Tools

One of our goals in approaching macroeconomic analysis is to make sure that our models are well-grounded in microeconomic behavior. The Solow model's assumption that people save a constant share of their income is exactly the kind of *ad hoc* assumption that we are trying to avoid. A reasonable theory of saving should allow people to decide how much of their income to save and consume. This choice should be influenced by such factors as the real interest rate, which is the market's incentive for people to save, and the relationship between their current income and their expected future income.

In microeconomics, we model saving and consumption choices using utility maximization. The Ramsey-Cass-Koopmans (hereafter abbreviated to just Ramsey) and Diamond growth models, which we study here in David Romer's Chapter 2, use the standard microeconomic theory of saving to make the saving rate endogenous. Because saving is a dynamic decision depending on past, present, and future income, we will need some new tools to analyze it. We use (at a somewhat superficial level) tools of dynamic optimal control theory to examine the household's optimal consumption/saving decision over time.

Most macroeconomic models being developed today begin from the Ramsey/Diamond framework of utility maximization, varying mainly in whether time is continuous (as in Ramsey) or discrete (as in Diamond) and whether households have infinite (Ramsey) or finite (Diamond) lifetimes.

Endogenous saving adds considerable complication to the dynamics of growth. The marginal net rate of return on capital (the equivalent of the real interest rate in this model) depends on the capital-labor ratio. As the capital-labor ratio changes during convergence toward the steady-state, the corresponding change in the return to capital and the real interest rate will cause changes in the saving rate. In order to track the dynamics of two variables as we move toward equilibrium, we will need a two-dimensional "phase plane" in which two variables simultaneously converge. Moreover, the nature of the equilibrium in this model is a "saddle point," which has interesting dynamic properties.

Chapter 2 is one of the most challenging chapters in the Romer text. Don't be discouraged if you don't understand everything immediately. Rely on a combination of the text, class lectures, and this coursebook chapter to help you achieve a working understanding of the model. As always, don't hesitate to ask for help!

B. Introducing Dynamic Utility Functions

Just as in microeconomics, we use utility functions to quantify people's preferences: what they like and what they dislike. The most common application of utility functions in microeconomics is to analyze choices between two different goods, say, asparagus and Brussels sprouts. In macroeconomics, we usually aggregate all goods together, so we do not worry much about choices among goods. Instead, we use utility functions to model preferences about generic "goods" consumed at different times and about preferences for leisure relative to these goods. The budget constraint in such models reflects the need to sacrifice leisure time in order to work and earn wages to purchase goods. (Or, recalling Jane in the simple one-person economy of Chapter 1, sacrificing leisure in order to produce the goods directly.)

In the growth models we shall study in this course, we take the labor/leisure decision as given—labor supply is perfectly inelastic—and focus only on the former decision: *when* to consume the goods that our lifetime worth of income allows. Should we consume a lot now and borrow against future earnings? Or should we be frugal now and save in order to consume more in the future? This is the essential decision on which we focus.

The utility functions that we use embody three basic preferences that we assume all individuals or households have:

- **People prefer more consumption to less, but at a decreasing rate.** In other words, they never become satiated with consumption goods, though the "marginal utility" of additional units of goods declines as they consume more at any point in time.
- **People prefer consumption sooner rather than later.** Consumption further in the future gives people less utility than consumption now or soon. One can attribute this property to people's innate impatience or, perhaps, to the "bird in the hand" phenomenon that something may happen to sidetrack future consumption but present consumption is certain. In our utility functions, the parameter ρ (Greek letter rho) will be used to measure impatience. People with a higher value of ρ have stronger preferences for current over future consumption.
- **People prefer a smooth consumption path over time rather than a lumpy one.** This follows from the assumption that marginal utility of consumption declines. It will always benefit households to shift consumption from high-consumption years (where marginal utility is low) to low-consumption years (where it is high). The result is a preference for a smooth path of consumption

over time. The parameter θ (Greek letter theta) will measure the strength of people's preference for smooth consumption. Those with a high θ want very smooth consumption and are not very willing to deviate from it; those with a low θ are more willing to substitute consumption across time.

The innate conflict between these last two preferences motivates the tradeoffs that we study here. The preference for current over future consumption would, if it were the only thing that mattered, cause people to consume their entire lifetime income right now. But this would lead to a very non-smooth consumption path, with extremely high consumption now and zero consumption in the future. Thus, the preference for smooth consumption prevents households from overdoing their preference to consume sooner.

Going the other direction, if consumption smoothing were the only thing that mattered, households would always consume at the same level at every moment in time. But given that they prefer consuming sooner rather than later, that is unlikely to be an optimal decision either.

In this chapter, we introduce a utility function called the constant-relative-risk-aversion (CRRA) function that embodies these three properties of preferences. The next section discusses in more detail how we incorporate the preference for sooner consumption into the utility function through discounting.

C. Discounting the Future in Discrete and Continuous Time

The idea of discounting

Introductory economics teaches you that comparing dollar values at different points in time requires *discounting*—expressing future and past payments in terms of comparable *present values*. For example, if the market interest rate at which you can borrow or lend is 10 percent, then you get the same consumption opportunity from receiving \$100 today as from receiving \$110 dollars one year from today.

Table 1 shows this by examining four cases in a 2×2 table. The top-left and bottom-right cells show what happens if the individual consumes the income when it is received; the top-right and bottom-left cells illustrate the individual's ability to perform intertemporal substitution through borrowing or saving at an interest rate of 10 percent.

The upper row shows your options if you receive \$100 now. If you wish to consume now, you simply spend the \$100. If you would rather spend the money next year, you lend the \$100 out at 10 percent interest. Next year you receive \$110 in principal and interest payments and spend it on \$110 worth of goods.

| | | |
|-------------------------|--|--|
| | Consume \$100 today | Consume \$110 next year |
| Receive \$100 today | Consume \$100 today when received | Lend \$100 today at 10%, receive and consume \$110 next year |
| Receive \$110 next year | Borrow \$100 today at 10% and consume; repay \$110 next year | Consume \$110 next year when received |

Table 1. Consumption opportunities

The lower row shows that you get the same consumption options from receiving \$110 next year. If you wish to consume next year, you simply spend the money when it is received. If you wish to consume today, you borrow \$100 and spend it today, then repay the principal and interest next year when you receive \$110. Thus, regardless of which of these payments is to be received, you have identical consumption options: consume \$100 today or \$110 next year. Thus, we say that these two payments have an identical *present value* of \$100.¹

We can generalize the concept of present value to allow payments to be made two or more years in the future. In doing so, we must take account of the *compounding* of interest—the fact that in future years you can earn interest not only on your principal but also on interest payments that have already been received. Suppose first that interest on loans is paid once per year and, again, that the interest rate is 10 percent per year. Each year that you lend, the value of your money increases by a factor of 1.10 or, more generally, by $1 + r$ where r is the interest rate. How large a payment made two years from today would give you consumption opportunities equivalent to those from a payment of \$100 today? If you received \$100 today, you could lend it out for the first year and receive \$110 back in one year ($\$100 \times 1.10 = \110). You could then lend out \$110 for the second year and receive \$121 back two years from today ($\$110 \times 1.10 = \121). Thus, \$121 two years from now has the same present value as \$100 today.

¹This example assumes that you can borrow and lend freely at a uniform interest rate. Calculation of present values is more complicated if consumers must pay a higher interest rate when they borrow than they receive when they lend, or if consumers are “liquidity constrained” and cannot borrow at all.

In terms of a mathematical formula, the future payment Q is related to its present value PV by $Q = PV \times (1 + r)^n$ if the payment is received n years in the future. Dividing both sides of this equation by the expression in parenthesis gives us the familiar discrete-time present-value formula:

$$PV = \frac{Q}{(1 + r)^n}. \quad (1)$$

We can use equation (1) to verify both of our examples above. In the one-year example, $\$100 = \$110/(1.10)^1$, so the present value of a \$110 payment received one year in the future is \$100 when the interest rate is 10 percent. For the two-year example, $\$100 = \$121/(1.10)^2$, so the present value of a \$121 payment two years in the future is \$100 when the interest rate is 10 percent.

Frequency of compounding and present value

Equation (1) is based on the assumption that interest is paid (or “compounded”) once per year. Would we get the same result if interest were paid each quarter or each month rather than once per year? No. The more frequently interest is compounded, the faster your money grows. This is exactly the same process as the compounding of growth rates discussed in the previous chapter.

Suppose that the annual interest rate is 10 percent, but that this is paid quarterly so that you receive $\frac{1}{4} \times 10$ percent = 2.5 percent each quarter. If you lend \$100 on January 1, then on April 1 you will have \$102.50, the \$100 principal and the first \$2.50 interest payment. Lending the entire \$102.50 for the second quarter will give you $\$102.50 \times 1.025 = \105.0625 on July 1. By the end of a year, you will have $\$100 \times (1.025)^4 \cong \110.38 , rather than the \$110 you would have if your interest was compounded annually.²

We can generalize this example into a formula as well. If the annual interest rate is r and interest is compounded k times per year, then the present value of a payment to be received in n years is

$$PV_k = \frac{Q}{(1 + \frac{1}{k}r)^{kn}}. \quad (2)$$

²In the United States, financial institutions are required to disclose “annual percentage rates” on loans, to make it easier for consumers to compare interest rates on loans with different compounding intervals. The APR on the quarterly-compounded loan in the example is 10.38%, the rate on an equivalent annually compounded loan.

Because the present-value formula depends on how often interest is compounded, we need to adopt a convention about which compounding interval to use. In discrete-time models, we usually assume that interest is paid once per period and express our interest rates in “percent per period.”³ This assumption means that we can use equation (1) to calculate present value.

An alternative assumption that is common in continuous-time models is that interest is *continuously compounded*. This amounts to a limiting case in which interest accrues at each instant, with each (infinitesimally small) payment of interest beginning to earn interest immediately. Mathematically, we can derive the continuous-compounding present-value formula by taking the limit of equation (2) as $k \rightarrow \infty$, *i.e.*, as the number of times interest is compounded per period gets very large.

Although this seems like it would complicate the mathematics, it can be shown that

$$\lim_{k \rightarrow \infty} PV_k = \lim_{k \rightarrow \infty} \frac{Q}{\left(1 + \frac{1}{k}r\right)^{kn}} = \frac{Q}{e^{rn}} = Qe^{-rn}, \quad (3)$$

where e is the exponential constant. Because the exponential function is much easier to work with in mathematical applications than the function in equation (1), many economic models, including almost all continuous-time models, use the formula in equation (3). Summarizing equation (3) in words, the present value of a payment is equal to the amount of the payment times e to the power of minus the interest rate (per period) times the number of periods in the future the payment is to be received.

We can see the similarity of equation (3) to the continuous-time growth formula given by equation (2) of Chapter 2 more easily if we solve equation (3) for Q to get $Q = PV e^{rn}$. This shows that for a given present value (amount invested) PV , the future value grows exponentially at continuous rate r over time.

³In a theoretical model, we do not usually specify what the length of the period is. Since real-world interest rates are universally quoted in “percent per year,” it may be most comfortable to think of a period as being a year. However, many of the discrete-time models we develop may be more realistic if the time period is shorter or longer. When applying the models to a time period other than a year, it is important to remember that interest rates (and also inflation rates and growth rates) must be expressed in terms of “percent per period” rather than the more-familiar “percent per year.”

Discounting money vs. discounting utility

The discussion above is framed entirely in terms of discounting a monetary payment. This monetary payment is worth less at a future date than it is today because you can earn interest on money that you receive today if you choose not to spend it immediately. Economists also use a formula that looks similar to equation (3) to discount future *utility*, arguing that utility received in the future is worth less than utility received now. What is the basis for using a formula like this to discount utility?

The discounting of utility cannot be justified in the same way as the discounting of payments because one cannot borrow or lend utility in a market. Suppose that for some reason you are extremely happy today, but you would rather “save” some of this happiness for tomorrow. There is no market in which you can lend today’s happiness to save it for tomorrow.⁴ Thus, the discounting of future utility relative to present utility cannot be based on a market argument similar to that used for discounting future money payments.

Rather, the basis for discounting utility is the observation that most people, if given a choice, seem to prefer to enjoy something now rather than in the future *if all else is equal*. Suppose, for example, that someone offers you an all-expenses-paid Hawaiian vacation, to be taken whenever you wish. You cannot sell this vacation to anyone else, nor can you “redeem” it for cash, so there is no way to earn interest on the vacation by choosing to take it earlier or later. Our observation above about human behavior claims that most people would prefer to take the vacation this year rather than, say, ten years from now.⁵

In order to capture this assumed preference for present over future utility, we discount future utility at a constant rate to its “present-value equivalent” whenever the agents in our model must compare utility at different points in time.⁶ The “rate of time discount” (which takes the place occupied by the rate of interest in monetary present-value calculations) is often represented by the Greek letter ρ . In discrete-time models,

⁴You may, of course, be able to lend *money* today by forgoing today’s purchases, which will give you money to make more purchases tomorrow. If purchases give you utility, then you can exchange current utility for future utility by this indirect means. Our intertemporal equilibrium consumption and saving decision relies on this kind of substitution. But this is not the same as being able to lend or borrow actual utility.

⁵Recall the “all other things equal” assumption. This assumption rules out “I’m too busy this year but I’ll have lots of free time ten years from now” and other similar cases.

⁶The assumption of a constant rate of discount makes the analysis easy, but is not necessarily realistic. For example, Laibson (1997) proposes “hyperbolic discounting,” in which households discount all future time more heavily relative to the present than they do points in the future relative to each other. However, Barro (1999) shows that hyperbolic discounting has only minor effects on the simple growth models that we discuss here.

we usually use a formula similar to equation (1). For example, if we want to represent the lifetime utility (U) of an individual who lives for two periods and gets utility $u(C_t)$ from consumption in period t , we might write

$$U = u(C_1) + \frac{1}{1+\rho} u(C_2). \quad (4)$$

Romer's equation (2.43) is an example of how equation (4) can be applied using a specific form for the $u(C_t)$ function. If we have more than two periods, or even an infinite number of periods, we can generalize equation (4) as⁷

$$U = \sum_{t=0}^{\infty} \frac{1}{(1+\rho)^t} u(C_t). \quad (5)$$

Adding up values in continuous time using integrals

In continuous time, we use an equation that differs from equation (5) in two ways. First, we use an exponential discounting expression similar to the one shown in equation (3). Second, because time is continuous we cannot simply sum up utility values corresponding to all the points in time—there are infinitely many such points. Instead, we must use the concept of an *integral*, which is drawn from basic calculus, to add up utility over time.

To see how integrals correspond to summations, think about adding up the amount of water flowing down a river during a day. There is a rate of flow at every moment of time, call it $w(t)$, measured in gallons per hour. But how are we to add up the infinite number of momentary flows that could potentially be observed at the infinite number of moments in the day?

One way would be simply to measure the flow (expressed in gallons per hour) at the beginning of the day $w(0)$ and multiply it times the number of hours in the day (24). This would be an accurate measure only if the rate of flow at the beginning of the day was exactly the average rate over the entire day. A better approximation could probably be achieved by taking two measurements, one at the beginning of the day $w(0)$ and one in the middle $w(12)$, multiplying each by the number of hours in the half day (12) and summing: $[w(0) \times 12] + [w(12) \times 12]$. Alternatively, we could measure every hour, multiplying each measurement by one (the number of hours in an hour), and adding up, or we could measure every minute, multiply each reading by 1/60 (the

⁷Equation (5) is the summation of an infinite number of terms. Depending on the path over time of $u(C_t)$, the value of this summation may be infinite or finite. We will only deal with problems in which the sum is finite. This requires that the $(1 + \rho)^t$ term in the denominator get large faster than the $u(C_t)$ term in the numerator.

number of hours in a minute) and add them up. Mathematically, measuring k times per hour would give us

$$W = \sum_{i=1}^{24k} w\left(\frac{i}{k}\right) \times \frac{1}{k}. \quad (6)$$

The most accurate of all would be the (impractical) limiting case where we would measure w continuously and add up the infinite number of such readings. This is what an integral does.⁸ We define the integral by

$$\int_{t=0}^{24} w(t) dt \equiv \lim_{k \rightarrow \infty} \sum_{i=1}^{24k} w\left(\frac{i}{k}\right) \frac{1}{k}. \quad (7)$$

The “limits of integration” at the bottom and top of the integral sign in equation (7) specify the values of the variable t over which the summation is to occur, $w(t)$ is the expression to be summed, and the dt term on the end indicates that it is that variable t that varies from 0 to 24.

Although taking an infinite number of readings is obviously impossible in practice, integrals such as equation (7) can often be evaluated in theoretical models if we can represent $w(t)$ by a mathematical function. For suitable functions, we can find a representation for the integral expression by finding the function $W(t)$ whose derivative is $w(t)$ and calculating $W(24) - W(0)$. Integration is the inverse operation of differentiation, so $W(t)$ is computed by finding the “anti-derivative” of $w(t)$. Any introductory calculus book can give you more details about integrals. However, we shall rarely be concerned with actually evaluating integrals, so we do not pursue these details here.

Discounting utility in continuous time

We can use the concept of the integral to add up the discounted values of momentary utility over a continuous interval. Suppose that utility at every moment depends on consumption at that moment according to the function $u[C(t)]$. If the rate of time preference is ρ , then the value of utility at t discounted back to the present ($t = 0$) is $e^{-\rho t} u[C(t)]$. Adding up this discounted utility for each moment from the present into the infinite future yields

$$U = \int_{t=0}^{\infty} e^{-\rho t} u[C(t)] dt. \quad (8)$$

⁸ Those who have studied the fundamentals of calculus will recognize the successive approximations above as the Riemann sums that are used in the formal definition of the integral.

Equation (8) combines the infinite-horizon summation in equation (5) with the continuous-time discounting formula of equation (3). Except for Romer's adjustment for the size of household (which is discussed below), it is equivalent to Romer's equation (2.2), with which he begins the analysis of the Ramsey-Cass-Koopmans model.

D. Constrained Maximization: The Lagrangian

As discussed in the previous chapter, setting the first derivative to zero can usually be used to determine the value(s) at which a function achieves a maximum or minimum value. However, there are many problems in economics where individuals are limited in the values of the variables they can choose in order to maximize utility or profit. Households and firms must often choose among the values that satisfy some economic constraint, such as the budget constraint that limits choices in utility maximization. Instead of looking for a general maximum, which can be done with the simple first-derivative rule, we must look for the maximum among only those values of the variables that fulfill the constraint.

The method of *Lagrange multipliers* is used to find the maximum or minimum of a function subject to a constraint. Courses in microeconomics (such as Reed's Econ 313) often spend considerable time solving Lagrange-multiplier problems. We shall introduce the concept briefly to make understanding Romer's Chapter 2 easier, but we will devote little time to actual problem solving.

The general objective of a constrained maximization problem is to choose the values of some variables, say, x_1 and x_2 , in a way that maximizes a given function $g(x_1, x_2)$ subject to the constraint that $a(x_1, x_2) = c$. For example, $g(x_1, x_2)$ could be a utility function with x_1 and x_2 being the levels of consumption of two goods, while $a(x_1, x_2)$ is the cost of consuming x_1 and x_2 and c is the consumer's income.

The theorem that underlies the method of Lagrange multipliers asserts that a maximum or minimum of $g(x_1, x_2)$ subject to the constraint that $a(x_1, x_2) = c$ occurs at the same values of x_1 and x_2 at which there is an *unconstrained* maximum or minimum value of the *Lagrangian* expression $L(x_1, x_2, \lambda) \equiv g(x_1, x_2) + \lambda[c - a(x_1, x_2)]$, where λ is called a Lagrange multiplier. Maximization of the Lagrangian is performed by the usual method of unconstrained maximization: setting the partial derivatives equal to zero.⁹

⁹As with unconstrained problems, either a maximum or a minimum can occur where the partial derivatives are zero. For the remainder of this section we will focus on maximization problems, since that is the nature of the problems in Chapter 2. In general, to determine whether a given point is a maximum or minimum one must examine second-order conditions. We will not

For the Lagrangian, which is a function of three variables, we maximize with respect to x_1 and x_2 and also with respect to λ , giving us three partial derivatives to set to zero. This leads to a system of three equations that we can attempt to solve for x_1 , x_2 , and λ . (These equations are called “*first-order conditions*” for a maximum.) The values of x_1 and x_2 that we obtain from this solution are the ones that maximize the function subject to the constraint. The value of λ is interpreted as the “*shadow price*” of the constraint. In the constrained utility-maximization problem discussed above, λ is the marginal utility of additional income—the improvement in the objective (utility) function that would be obtained from a one-unit relaxation of the (budget) constraint.

One of the partial derivatives that we set equal to zero is the partial derivative with respect to λ . A closer look at this derivative shows the logic underlying the method of Lagrange multipliers: $\partial L/\partial\lambda = c - a(x_1, x_2)$. Setting $c - a(x_1, x_2) = 0$ is equivalent to enforcing the budget constraint $a(x_1, x_2) = c$. Since $\partial L/\partial\lambda = 0$ is one of the three first-order conditions that we solve to get the values of x_1 and x_2 , we are assured that these values lie on the budget constraint.

A straightforward example of a Lagrangian is Romer’s equation (2.50), which is the consumer’s maximization problem in the Diamond model. The two first-order conditions shown in equations (2.51) and (2.52) result from setting equal to zero the partial derivatives of the Lagrangian with respect to the two choice variables, C_{1t} and C_{2t+1} . The third first-order condition, from the partial derivative with respect to the Lagrange multiplier λ , is not explicitly shown. It replicates the constraint (2.46).

Romer’s equation (2.17), which is the consumer-choice problem for the Ramsey model, is a more complicated application of a Lagrangian. The objective function being maximized is an integral representing the discounted value of utility. The constraint is the complicated expression in brackets, which says that the present value of lifetime income equals the present value of lifetime consumption. The Lagrangian is maximized with respect to λ and with respect to all the (infinite set of) values of $c(t)$.

While the method of Lagrange multipliers is very useful in economic analysis, we will spend no more time on it here. Interested students should consult an advanced microeconomics text or the relevant chapters of a book on mathematical economics such as Chapter 12 of Alpha Chiang, *Fundamental Methods of Mathematical Economics* 3d ed. (New York: McGraw-Hill, 1984).

discuss the second-order conditions of Lagrangian problems; assumptions about the parameters of our models assure that the second-order conditions for a maximum are fulfilled for the problems in Chapter 2.

E. Using Indifference Curves to Understand Intertemporal Substitution

We are most comfortable using indifference curves to analyze consumption choices, and this tool can easily be used to explain the intertemporal substitution model and consumption smoothing. Of course, relying on indifference curves allows us to examine only two dimensions at a time, so we can apply this method only to a two-period model.

Suppose that utility is given by

$$U = u(C_1) + e^{-\rho}u(C_2).$$

Even though the two-period model requires that we work in discrete time, we shall use continuously compounded discounting to retain more symmetry with the Ramsey framework. With continuous compounding of interest, the individual's budget constraint is

$$W_1 + e^{-r}W_2 \equiv Y = C_1 + e^{-r}C_2,$$

where W_1 and W_2 are wage earnings in periods one and two. We are taking labor supply to be perfectly inelastic and assuming that the household is always able to work its desired number of hours, so wage earnings are fixed. Thus, from the consumer's standpoint, Y and r (the real interest rate, reflecting the rate at which present goods can be exchanged for future goods) are given, ρ is a parameter of the utility function, and C_1 and C_2 are the individual's constrained choices.

Consider first the consumer's budget constraint. We will plot C_1 on the horizontal axis and C_2 on the vertical axis, so we begin by solving the budget constraint for C_2 to get

$$C_2 = \frac{Y}{e^{-r}} - \frac{C_1}{e^{-r}} = e^r(Y - C_1).$$

This is a straight line intersecting the vertical axis at $e^r Y$ (and the horizontal axis at Y) and having a slope equal to $-e^r$.

Next consider the consumer's indifference curves. These will not be linear; their exact form depends on the functional form of the function $u(\cdot)$. The properties assumed

for utility functions assure us that $u' > 0$ and $u'' < 0$ —marginal utility is positive (non-satiation) and diminishing. An indifference curve corresponding to a given level of utility U_0 is defined as the set of (C_1, C_2) for which

$$U = U_0 = u(C_1) + e^{-\rho}u(C_2).$$

We are interested in the slope of the indifference curve, which is the change in C_2 that leaves utility unchanged ($dU = 0$) following a unit change in C_1 :

$$\left. \frac{dC_2}{dC_1} \right|_{dU=0}.$$

Since (without assuming a specific form for the utility function) we cannot solve for C_2 here as an explicit function of C_1 and U_0 , we can only obtain the slope of the indifference curve by “implicit differentiation.” We begin by taking the “total differential” of the utility function:

$$dU = u'(C_1)dC_1 + e^{-\rho}u'(C_2)dC_2.$$

This equation says that the change in utility (dU) equals the change in consumption in period one (dC_1) times the marginal utility of a unit of C_1 (given by $u'(C_1)$) plus the change in consumption in period two (dC_2) times the marginal utility of a unit of C_2 (which is $u'(C_2)$). Utility is not changing as we move along an indifference curve, so we set $dU = 0$ and solve for dC_2/dC_1 :

$$\left. \frac{dC_2}{dC_1} \right|_{dU=0} = -\frac{u'(C_1)}{e^{-\rho}u'(C_2)} = -e^{\rho} \frac{u'(C_1)}{u'(C_2)}.$$

As we move from lower right to upper left in the positive quadrant, C_1 gets smaller and C_2 gets larger. The negative second derivative of the utility function assures us that $u'(C_1)$ gets larger as C_1 gets smaller and $u'(C_2)$ gets smaller as C_2 gets larger, so $u'(C_1)/u'(C_2)$ increases as we move from lower right to upper left: the indifference curves get steeper and are convex in the usual way.

We know that (barring a corner solution, which is improbable here) the consumer maximizes utility by consuming at a point where an indifference curve is tangent to

her budget constraint. This tangency occurs at a point where the slope of the indifference curve equals the slope of the budget constraint. Recall that the slope of the budget constraint is $-e^r$, so the mathematical equilibrium condition is

$$-e^\rho \frac{u'(C_1)}{u'(C_2)} = -e^r,$$

or

$$\frac{u'(C_1)}{u'(C_2)} = e^{(r-\rho)}.$$

This is the basic first-order condition for utility maximization in the two-period model. The relationship between C_1 and C_2 on the left-hand side clearly depends on the relative size of r and ρ on the right. We can establish more intuition about the consumer's choices by considering the three cases $r = \rho$, $r > \rho$, and $r < \rho$.

Consider first the case where the rate of return on capital (real interest rate) r is exactly equal to the consumer's marginal rate of time preference ρ . In this case $-e^\rho = -e^r$, so the tangency condition becomes $u'(C_1)/u'(C_2) = 1$, or $u'(C_1) = u'(C_2)$. Because $u'' < 0$, $u'(C_1) = u'(C_2)$ if and only if $C_1 = C_2$, which means in geometric terms that the tangency between the indifference curve and the budget line must lie on the 45° line through the origin. In economic terms, with the interest rate equals the rate of time preference, an individual will choose equal consumption in both periods: she will *smooth* consumption. Figure 1 shows this consumer equilibrium situation at point a .

Intuitively, the interest rate is the reward to saving offered by the market and the rate of time preference is the reward demanded by the consumer to justify postponing consumption. If the two are exactly in balance, then the consumer chooses future consumption equal to present consumption.

If, further, the consumer's income in the two periods is equal, she will choose zero saving. If she has higher income in one period than the other, she will save in the higher-income period and dissave in the lower-income period to smooth her consumption perfectly.




Figure 1. Consumer equilibrium when $r = \rho$

We have established that a perfect consumption-smoothing outcome maximizes utility when $r = \rho$. What happens when $r > \rho$ or $r < \rho$? Once again, we can answer these questions easily with the indifference-curve diagram.

If $r > \rho$, then $e^r > e^\rho$ and the budget constraint is steeper than the indifference curve when $C_1 = C_2$. (In Figure 1, twist the linear budget constraint to make it steeper.) This means that the tangency must occur at a point on the budget constraint above and to the left of point a in Figure 1, where $C_1 < C_2$. In terms of the mathematical equilibrium condition, $e^{(r-\rho)} > 1$, so $u'(C_1)/u'(C_2)$ must be greater than one and $u'(C_1) > u'(C_2)$. With marginal utility decreasing in consumption, $u'(C_1) > u'(C_2)$ if and only if $C_2 > C_1$.

Intuitively, when $r > \rho$, the consumer chooses more consumption in the future than in the present. She wants a consumption path that rises over time because the market reward to saving (r) exceeds her innate desire to consume early (ρ). If her income is the same in both periods, she will choose positive saving in period one. (However, it is important to note that if her income in period two is considerably higher than period one she may dissave in period one, although her consumption in period two will still be higher than in period one.)

The exact opposite happens if $r < \rho$. The budget constraint is flatter than the indifference curves at $C_1 = C_2$, so the tangency occurs below and to the right of Figure 1's

point a , where $C_1 > C_2$. Mathematically, $e^{(r-\rho)} < 1$, so at equilibrium $u'(C_1) > u'(C_2)$, which implies $C_2 < C_1$. In terms of intuition, when $r < \rho$, the market reward to postponing consumption falls short of the consumers innate desire for current consumption, thus she chooses higher consumption now and lower consumption in the future. If her income is equal in both periods, she will dissave (borrow) in the current period to finance high current consumption at the expense of lower future consumption when she must pay off her debt in the future. (Once again, if her income were enough lower in the future period, she might actually choose positive saving in period one, but she will not save enough to smooth her consumption perfectly.)

F. Understanding Romer's Chapter 2, Part A

As noted above, Chapter 2 is one of the most mathematically difficult in Romer's text. This section and the one that follows are intended to facilitate your understanding of the mathematically challenging sections.

Household vs. individual utility

The basic setup of Romer's equation (2.2) was discussed above in the context of continuous-time discounting. However, one aspect of the equation was ignored there: the presence of the $L(t)/H$ term. Writing the utility function in the way that Romer does implies that $u[C(t)]$ is to be interpreted as the utility gained at time t by one individual family member, but that decisions are made in a way that maximizes *total household* utility. $L(t)$ is the number of people in the economy and H is the (constant) number of households. Thus, $L(t)/H$ is the number of members in each household at time t , so multiplying by this factor translates individual utility into household or family utility.¹⁰

As Romer points out in his footnote 1, the problem can be easily reformulated with individual utility being maximized, but with the discount rate ρ being interpreted differently. The only effect of the total-family formulation on the model's conclusions that the form of the dynamic stability condition $\rho - n - (1 - \theta)g > 0$ is slightly different if the alternative formulation is chosen.

¹⁰ Obviously, fractional humans do not exist, so in real life $L(t)/H$ must be a non-negative integer. We will ignore this problem here (and throughout the course when it arises).

Choosing a functional form for the utility function

Problems such as this one cannot be solved for utility-maximizing consumption paths without choosing a particular functional form for the instantaneous utility function $u(\bullet)$. Romer's equation (2.3) gives the functional form he chooses, the *constant-relative-risk-aversion* (CRRA) utility function, which is common in growth analysis.¹¹ It is probably far from obvious to you why he chose this particular function, so let's think a little bit about some criteria one might use to choose a form for the utility function: admissibility, convenience, and flexibility.

First of all, the functional form must be *admissible*, meaning that it must satisfy the conventional properties of a utility function. We usually assume that the marginal utility of consumption $u'[C(t)] = du / dC(t)$ is positive for all values of $C(t)$, but that marginal utility is decreasing, which means that $u''[C(t)] = d^2u / dC(t)^2$ is negative. In model-building we often choose linear functions when possible, but the latter condition rules out a linear utility function, because the second derivative of a linear function is always zero (*i.e.*, if utility were linear then marginal utility would be constant, not decreasing). A quadratic function might be considered—utility could be represented by the upward-sloping part of a downward-opening parabola. But a quadratic utility function would only work over a limited region, because every downward-opening parabola eventually reaches a maximum at some level of consumption and for levels beyond that the marginal utility of consumption is negative, violating one of our assumptions. Utility can be approximated locally, but not globally, by a quadratic utility function.

Since linear and quadratic utility functions cannot provide a globally suitable functional form, a natural alternative to consider is a power function similar to the Cobb-Douglas production function, where utility equals a constant times consumption raised to some power. The constant-relative-risk-aversion function that Romer chooses is of this type and is easily shown to be admissible.

A second criterion for choosing a functional form is *convenience* or simplicity. Although the CRRA function does not appear to be very simple at first glance, it turns out (as you will see in a few pages) that its first and second derivatives are very simple, so the solution is of a particularly simple form when it is used for the utility-maximization problem.

¹¹ This function is sometimes called the “constant elasticity of intertemporal substitution” utility function, which may be a more appropriate title for our risk-free intertemporal application. When the function is used to analyze risky decisions, the relative rate of risk aversion is the constant θ ; when we use it to analyze intertemporal behavior, the elasticity of intertemporal substitution is $1/\theta$, which is also constant. So either name is justified.

A third criterion for choosing a function is *flexibility*. The CRRA function is quite flexible in that by varying the parameter θ it can represent a wide spectrum of consumption behavior: indifference curves can be sharply bending (large θ) or straight lines ($\theta = 0$). As Romer notes, θ measures the household's resistance against substituting consumption in one period for consumption in another. This is an important behavioral parameter in macroeconomic modeling. To see why, we digress for a moment on to examine in more detail the concept of consumption smoothing, which was introduced above.

Consumption smoothing

Suppose that we ignore issues of discounting for a moment (set the interest rate and the rate of time preference equal to zero) and consider the maximization of utility for a consumer who lives two periods. If the marginal utility of consumption is positive but decreasing, then the utility function is concave and looks similar to the one in Figure 2. Suppose that the consumer has a fixed amount of income Q to allocate between consumption in period one and in period two. We assume that the interest rate is zero, so the budget constraint is simply $C_1 + C_2 = Q$. One choice (which turns out to be the optimal choice) would be to consume the same amount in each period—to “smooth” consumption perfectly. This would imply consuming $\frac{1}{2}Q$ in each period and getting lifetime utility equal to $2u(\frac{1}{2}Q)$ —twice the utility of $\frac{1}{2}Q$. (Remember that we are ignoring the issue of discounting future utility so that total lifetime utility is just the unweighted sum of utility in the first and second periods.)

To see that consumption smoothing is the optimal plan, consider the alternative feasible plan of consuming $\frac{1}{2}Q + x$ in one period and $\frac{1}{2}Q - x$ in the other, where x is positive and less than $\frac{1}{2}Q$. This gives lifetime utility of $u(\frac{1}{2}Q + x) + u(\frac{1}{2}Q - x)$. However, notice from Figure 2 that because of the concavity of the utility function, the additional utility gained in the high-consumption period $u(\frac{1}{2}Q + x) - u(\frac{1}{2}Q)$ is smaller than the utility lost in the low-consumption period $u(\frac{1}{2}Q) - u(\frac{1}{2}Q - x)$. Because of this, the average utility per period under consumption smoothing, $u(\frac{1}{2}Q)$, exceeds the average utility from the uneven consumption path, $\frac{1}{2}[u(\frac{1}{2}Q + x) + u(\frac{1}{2}Q - x)]$. This implies that the utility of the smooth consumption plan is greater than that of the uneven plan, so when there is no interest or time preference ($r = \rho = 0$), then a utility-maximizing consumer with a concave utility function (diminishing marginal utility of consumption) will choose a smooth consumption path.

Now consider how the amount of curvature in the utility function affects this result. If the utility function is nearly linear (not very sharply curved), then the loss in utility from an uneven consumption plan is very small. If the utility function is sharply curved, then the loss is very large.

The parameter θ in the CRRA utility function controls the amount of curvature in the function. If θ is close to zero, then the function is almost linear and consumers are

quite willing to accept uneven consumption patterns. As the value of θ gets larger, the amount of curvature in the utility function increases and consumers' willingness to accept anything other than smooth consumption declines.¹²




Figure 2. Concave utility

As we shall see, introducing discounting of future utility and the earning of interest makes the issue of consumption smoothing a little more complicated, but the role of the θ parameter remains essentially the same. A small θ implies a high willingness to alter consumption patterns away from smoothness in response to such disturbances as changes in interest rates, while a large θ means that consumers are determined to pursue a smooth and regular consumption path in spite of these disturbances.

¹² The indifference curves between period-one and period-two consumption mirror this difference in curvature. With θ near zero, the indifference curves approach straight lines, making consumption in the two periods near-perfect substitutes. When θ is large, the indifference curves approach L-shaped and consumption in one and two are complements.

Discounting with time-varying interest rates: $R(t)$ and $r(t)$

In writing the simple formula for present value we usually assume that the interest rate is constant over time. In discrete time, this allows us to write the present-value formula as equation (1): $PV = Q / (1 + r)^n$, where Q is a payment to be received n years in the future. What would happen if r varies over time?

Consider a discrete-time example with annual compounding of interest. Q is to be received two years from now and the interest rate is 4% this year and will be 6% next year. One dollar lent at interest today would be worth $\$1 \times 1.04 = \1.04 after one year. Lending $\$1.04$ for the second year would increase its value to $\$1.04 \times 1.06 = \$1.1024 = \$1.00 \times 1.04 \times 1.06$. Thus, the present value of $\$1.1024$ two years from today is $\$1.1024 / [(1.04)(1.06)] = \1.00 . In the general case of a varying interest rate, the denominator of the present-value formula is the product of all the one-year $(1 + r(t))$ terms for all years between now and when the payment is received:

$$Q = PV \times \prod_{t=1}^n [1 + r(t)], \text{ or}$$
$$PV = \frac{Q}{\prod_{t=1}^n [1 + r(t)]}.$$

The large Π notation is similar to the familiar summation notation that uses Σ , except that the elements are multiplied together rather than added together.

How does this translate into continuous time? As discussed above, the continuous-compounding discount factor for payments n periods in the future (corresponding to $1 / (1 + r)^n$) is e^{-rn} if the interest rate is constant at r . The exponent of this discount factor is the interest rate multiplied by the number of periods, which would also be the result of summing the (constant) interest rate over n periods in much the same way that raising $1 + r$ to the power n takes the product of $(1 + r)$ over n periods.

What if the interest rate varies between now (time 0) and time n ? Then we must sum the varying values of the interest rate $r(t)$ between 0 and n . Because we are working in continuous time, we cannot just add up the interest rates corresponding to a finite set of points in time. Instead we must use an integral over the time interval 0 to n to sum up all the interest rates. Romer defines the term $R(n)$ to be the integral (sum)

$$R(n) = \int_0^n r(t) dt.$$

The appropriate discount factor for n periods in the future is then $e^{-R(n)}$. Note that if $r(t)$ is constant at r over the time interval 0 to n , then $R(n) = rn$ and the usual formula applies.

The positivity restriction on $\rho - n - (1 - \theta)g$

A final issue relevant to the utility function is the condition that $\rho - n - (1 - \theta)g$ must be positive. I don't know any intuition that would lead you to this condition prior to performing the dynamic analysis, so do not feel like you have missed something if the intuitive rationale is not obvious. This condition turns out to be necessary to assure the dynamic stability of the equilibrium of the growth model. Look at Romer's equation (2.13) and notice that when the utility function is expressed in terms of efficiency units of labor, the discount factor turns out to be exactly this expression. If $\beta \equiv \rho - n - (1 - \theta)g > 0$, then future utility in terms of consumption per efficiency unit of labor is discounted positively (*i.e.*, valued less than current utility).

We can think intuitively about why such a condition is necessary for a stable model. The income of each household grows in the steady state due to both population growth and technological progress (n and g). When we assume that $\rho > n + (1 - \theta)g$, we are assuming that there is a strong enough preference for current over future utility to outweigh the effects of population growth and growth in per-capita income (weighted by $1 - \theta$). If ρ were very small, then households would discount the future only slightly relative to the present. Since positive, ongoing growth will cause future levels of income and consumption to be much, much greater than current levels, small (proportional) changes in consumption in the infinitely distant future could have greater importance to household utility than large present changes. It is to assure that households care enough about changes in current consumption to provide a stable equilibrium that we require a sufficient degree of time preference to offset the explosive effects of growth.¹³ Households for which this condition did not hold would choose extremely high rates of saving that would lead the economy away from a stable, steady-state equilibrium.

Deriving the Ramsey consumption-equilibrium equation

The derivation of consumption equilibrium in this model is challenging. Unless you enjoy mathematical applications, you may skim the details of the math on Romer's pages 55 through 57 up to equation (2.21), which is the consumption-equilibrium equation he has been seeking. If you want to understand his approach, the paragraphs below can help.

Romer's intuitive approach to consumption equilibrium relies on the Lagrangian method discussed earlier in this chapter. The only formal problem with applying the Lagrangian here is that there are infinitely many $c(t)$ values to choose. Formally, differentiating the Lagrangian with respect to all of these values would lead to a system

¹³There is a mathematical side to this problem as well. If $\beta < 0$, then the integral in the last line of Romer's equation (2.14) does not exist because the integrand is getting larger and larger as $t \rightarrow \infty$. In this sense, the model "explodes" if $\beta < 0$.

of infinitely many equations in infinitely many unknowns. But the idea is the same and in this case we end up with the same answer.

The integrals in Romer’s Lagrangian expression (2.17) are summations over all values of t . When we differentiate with respect to one particular $c(t)$ we are holding all of the other c values constant. Thus, the only element of the summation (integral) that changes when $c(t)$ changes is the element corresponding to time t . So the partial derivative of the integral with respect to any individual $c(t)$ is just the derivative of the integrand with respect to $c(t)$. This partial differentiation leads to Romer’s equation (2.18).

A more formal way of finding the utility-maximization condition is the present-value Hamiltonian, a dynamic version of the Lagrangian. We shall use the Hamiltonian to analyze investment models later in the course. A main conceptual difference is that the Lagrange multiplier in a dynamic problem can vary over time, so it is replaced by a “costate variable” that is a function of time. For a solution to the Ramsey model that uses this method, see section 2.1 of Barro and Sala-i-Martin (2004).

Understanding the Ramsey consumption-equilibrium equation

As Romer notes in the discussion following equation (2.21), the outcome of this mathematical analysis is that growth rate of consumption per worker at time t is the ***Euler equation*** (equation 2.22 in Romer)

$$\frac{\dot{C}(t)}{C(t)} = \frac{r(t) - \rho}{\theta}. \tag{9}$$

Since θ is assumed to be positive (in order to give the utility function the appropriate shape), the sign of the growth rate of per-capita consumption on the left-hand side is the same as the sign of the numerator of the right-hand side, which is the difference between the current interest rate and the rate of time preference.

The intuition of this result is familiar; we studied it in the two-period model above. To reiterate, the interest rate measures the amount of additional future consumption the individual *can* obtain by sacrificing one unit of current consumption—the market reward for postponing consumption. Each marginal unit of current consumption that is forgone yields e^r units of consumption one period later (assuming continuous compounding). On the other side, the rate of time preference ρ measures the individual’s *unwillingness*, other things being equal, to postpone consumption. An individual with equal consumption in the two periods is indifferent to exchanging one unit of current consumption for e^ρ units (again, with continuous compounding) a period later.

If $r > \rho$ at time t , then the market reward for postponing consumption (the interest rate) exceeds the amount required to motivate an individual to move away from perfectly smooth consumption and forgo some current consumption, exchanging it for

future consumption through (additional) saving. Thus people want their future consumption to be higher than their current consumption when $r > \rho$, and as a result they choose a path on which consumption is rising at time t , which is represented mathematically by a positive consumption growth rate: $\dot{C}(t)/C(t) > 0$.

If $r < \rho$ at time t , then the interest reward offered by the market is insufficient for individuals to want to keep consumption smooth. In this case, people want more consumption now at the expense of the future and consumption will be declining at t , so $\dot{C}(t)/C(t) < 0$. The intermediate case in which the interest rate equals the rate of time preference is one in which individuals desire a constant level of per-capita consumption over time (exact consumption smoothing) and $\dot{C}(t)/C(t) = 0$. This is the case corresponding to Figure 1 in our earlier analysis of the two-period model.

The argument above explains why the relative magnitude of r and ρ determines whether desired consumption is rising or falling at each moment (the *sign* of the growth rate of consumption). We must still consider what determines *how much* any given deviation of the interest rate from the rate of time preference will cause consumers to alter their consumption paths away from smoothness. The sensitivity of consumption growth to the difference between the interest rate and the rate of time preference in equation (9) is $1/\theta$. We can now relate our above discussion of the θ parameter to the behavior of desired consumption.

If θ is near zero, then the instantaneous utility function $u(\cdot)$ in Figure 2 is nearly linear (as are the indifference curves in Figure 1). In this case, (as we discussed above) individuals have only a small preference for smooth consumption and are quite willing to change their consumption patterns in response to market conditions. Thus, when θ is small, (and $1/\theta$ is correspondingly large) consumption growth will react strongly to any given differential between the interest rate and the rate of time preference. Conversely, when θ is large (and $1/\theta$ is small), the utility function in Figure 2 and the indifference curves in Figure 1 bend sharply and they want to stick to their smooth consumption paths even when there is a strong market incentive to change. The expression $1/\theta$ is called the *elasticity of intertemporal substitution*.

Households are making decisions at time t about their future paths of consumption. Based on the future path of the interest rate $r(t)$, Romer's equation (2.21) tells us whether households want their consumption paths to be rising, falling, or flat at each moment. In other words, it tells us the slope of the time path of $\ln C(t)$ at every t . (Recall that $d(\ln C(t))/dt = \dot{C}(t)/C(t)$.) It does not, however, tell us the *level* of the consumption path, so we cannot yet determine exactly how much the household will consume at the instant t . There are many parallel paths of $\ln C$ (having the same slope at every time value t) that satisfy equation (2.20). How do we know which one the household will choose?

The missing ingredient that we have not yet brought into the analysis is the budget constraint (Romer's equation (2.7)), from which he derives the *no-Ponzi-game condition* of equation (2.11).¹⁴ Among these parallel consumption paths satisfying (2.21), some have very high levels of consumption so that in order to follow them the household would have to go ever deeper into debt as time passes. Others have such low consumption that the household would accumulate unspent assets consistently over time. Only one of these parallel consumption paths exactly exhausts lifetime income so that the present value of the household's wealth goes to zero as time goes to infinity as required by the no-Ponzi-game condition. That unique path is the consumption path that the household will choose and the point on that path corresponding to instant t determines $C(t)$.

We began this chapter with the goal of generalizing the Solow model's restrictive assumption about saving behavior. However, all of our discussion so far has focused on consumption rather than saving. What does this analysis imply about saving? At moment t , the household receives income $y(t)$ per efficiency unit of labor, which depends on the amount of capital available in the economy according to the intensive production function. We have just analyzed the determination of $c(t)$, the level of consumption per efficiency unit of labor. Saving (per efficiency unit) is just the difference between income and consumption at time t : $y(t) - c(t)$.

The steady-state balanced-growth path in the Ramsey model

As in the Solow model, we look for a steady-state value of the capital/effective labor ratio k . However, in the Ramsey model the dynamic analysis seems much more complicated, involving two variables (c and k) rather than just one (k). What is it about the Ramsey model that requires the more difficult analysis?

The equation of motion for k in both models is essentially the same as in the Solow model:

$$\dot{k}(t) = f(k(t)) - c(t) - (n + g)k(t), \quad (10)$$

which is Romer's equation (2.26). Notice that this equation involves two state variables: k and c . If we could find an explicit equation for $c(t)$ that we could substitute into the equation, then we could conduct the analysis using k alone. In the Solow model, we assume that $c(t) = (1 - s)f(k(t))$, so such a substitution can be made and equation (10) reduces to the familiar $\dot{k}(t) = sf(k(t)) - (n + g)k(t)$. However, in the Ramsey model we have no simple equation for $c(t)$; it is determined through the more complex process of first determining the growth rate at each point in time by equation (9) then

¹⁴ This condition is often called the *transversality condition*.

reconciling the growth path with the lifetime budget constraint. Because we cannot substitute c out to express the model solely in terms of k , we must analyze the dynamic behavior of c and k jointly.

To consider the joint movement of k and c , we need an equation of motion for c . This is provided by Romer's equation (2.25), which is derived from (2.21) by substituting the marginal product of capital for the interest rate. This equation is reproduced below:

$$\frac{\dot{c}(t)}{c(t)} = \frac{f'(k(t)) - \rho - \theta g}{\theta}. \quad (11)$$

In order to analyze the movements of the two variables together, we use a two-dimensional *phase diagram* such as Figures 2.1 through 2.3 in Romer. The little vertical and horizontal arrows in the phase diagram show, for any initial point (k_0, c_0) , the directions that the dynamic equations of motion (10) and (11) imply that k and c would move if the economy were at that point. For example, an upward arrow indicates that the economy's equations of motion would cause c (the variable on the vertical axis) to increase from that point, so over time the economy would tend to move to other points lying above (k_0, c_0) . The left and right arrows show the directions that the equations of motion would tend to move k (the variable on the horizontal axis) from any point in the plane.

The first step in constructing a phase diagram is to establish for each variable the set of points at which it is neither increasing nor decreasing. This means graphing the sets of (k, c) values at which $\dot{c} = 0$ and $\dot{k} = 0$. In the Ramsey model, equation (11), describing the dynamic behavior of consumption, is particularly simple because when it is set to zero, only $k(t)$ (and not $c(t)$) appears in the equation. Thus, there is a single unique value of k — call it k^* — for which $\dot{c} = 0$. That value is given by $f'(k^*) - \rho - \theta g = 0$ or $f'(k^*) = \rho + \theta g$. Regardless of the value of c , $\dot{c} = 0$ if $k = k^*$, so the locus of points at which $\dot{c} = 0$ is a vertical line at $k = k^*$, as shown in Romer's Figure 2.1.

When $k < k^*$, then $f'(k) > \rho + \theta g$. We know this because the marginal product of capital increases when the capital/effective labor ratio declines, so k dropping below k^* makes $f'(k)$ larger and thus, from equation (11), makes $\dot{c} > 0$. Similarly, if $k > k^*$, then the marginal product of capital will be lower than at k^* , so $f'(k) < \rho + \theta g$ and $\dot{c} < 0$. The vertical arrows in Romer's Figure 2.1 show the directions of motion of c at points off of the $\dot{c} = 0$ line. Table 2 summarizes the three possible conditions for the sign of \dot{c} in terms of algebra and in the phase diagram.

Table 2. Conditions for change in c

| Value of \dot{c} | Mathematical condition | Position on graph | Direction of arrow |
|--------------------|---------------------------|-----------------------|--------------------|
| $\dot{c} < 0$ | $f'(k) < \rho + \theta g$ | Right of $k=k^*$ line | ↓ |
| $\dot{c} = 0$ | $f'(k) = \rho + \theta g$ | On $k=k^*$ line | No vertical mvt. |
| $\dot{c} > 0$ | $f'(k) > \rho + \theta g$ | Left of $k=k^*$ line | ↑ |

The dynamic behavior of k is slightly more complicated than c because $\dot{k}(t)$ depends on the values of both $k(t)$ and $c(t)$, as shown by equation (10). The hump-shaped curve in Romer's Figure 2.2 shows the values of k and c for which $\dot{k} = 0$. As he describes in the text, points above this curve are values of c and k at which k is falling (hence the leftward arrow) and points below the curve are ones at which k rises. This curve is a graph of the levels of c that correspond to the different possible steady states for k . Constructing it is exactly analogous to the golden-rule experiment in the Solow model, where we considered the effect of different possible steady-state values of k on steady-state per-capita consumption. The maximum of the curve in Figure 2.2 corresponds to the golden-rule level of the capital/effective-labor ratio. The movement of k is summarized in Table 3.

Table 3. Conditions for change in k

| Value of \dot{k} | Mathematical condition | Position on graph | Direction of arrow |
|--------------------|------------------------|---------------------------|---------------------------|
| $\dot{k} < 0$ | $f(k) - c < (n + g)k$ | Above $\dot{k} = 0$ curve | ← |
| $\dot{k} = 0$ | $f(k) - c = (n + g)k$ | On $\dot{k} = 0$ curve | No \leftrightarrow mvt. |
| $\dot{k} > 0$ | $f(k) - c > (n + g)k$ | Below $\dot{k} = 0$ curve | → |

Putting the two curves (or, more precisely, the line and the curve) together gives Romer's Figure 2.3, which describes the dynamics of the system. The point at which the two curves intersect shows the steady-state equilibrium values of c and k . If the economy is at this point, neither k nor c will change. The arrows in each of the quadrants show how (or whether) the system will converge toward the steady state from any set of initial values of c and k .

Saddle-path convergence to the steady state

The convergence of the Ramsey model is far from obvious judging from the arrows in Figure 2.3. If the economy were to start from the upper-left or lower-right quadrant, it would move directly away from the steady-state point (such as from point A in

Romer's Figure 2.4). Even if the economy begins in the lower-left or upper-right quadrants, convergence to point E is not guaranteed. Notice the four single arrows placed on the $\dot{c} = 0$ and $\dot{k} = 0$ curves. All four of these arrows point into the unstable quadrants rather than into the potentially stable quadrants. This indicates that if the economy touches (or even gets close to) these curves on its way to E , it will veer off into the unstable region and diverge, as shown in Figure 2.4 by the paths starting from C and D .

The steady-state equilibrium in the Ramsey model is an example of a *saddle-point* equilibrium. There is a unique curve called the *saddle path* running from the interior of the lower-left quadrant and the interior of the upper-right quadrant to point E . If the economy begins on the saddle path, it will converge smoothly to the steady state at E . If it starts anywhere else, it diverges.¹⁵

The knife-edge nature of convergence to a saddle-point equilibrium may make you think that convergence is unlikely. It means that the economy must be in *exactly* the right place in order to converge. In our case, it means that given the value of k that we inherit from our past saving and investment decisions, the value of c must be exactly the right one to put us on the saddle path—a penny of consumption more or less than this amount leads to instability. Can we count on this?

Fortunately, the answer is yes. Because c is a “control” variable rather than a “state” variable, its value at time t is free to adjust upward or downward as necessary; it is not bound by its past history as is k . Furthermore, the value of $c(t)$ that puts the economy on the stable saddle path is *precisely* the value that puts the household on the optimal consumption path that balances its lifetime budget constraint. Thus, our utility-maximizing consumers will automatically choose the level of consumption per person that puts the economy on the saddle path to the steady state, and thus the steady state is stable.

Ramsey equilibrium is Pareto optimal

Those who have studied microeconomics in detail are familiar with the property that perfect competition in the absence of externalities leads to a Pareto optimal equilibrium. There is no way to reallocate resources that will make someone better off without making someone else worse off.

¹⁵The name “saddle” point for this equilibrium reflects this property. Think about releasing a marble or similar object from any point on a saddle (and ignore the effects of the marble's own momentum). From most points, the marble will slide to one side or the other and off the saddle. There is, however, one path running across the very center of the saddle on which, if you could balance it exactly right, the marble would slide right down into the center of the saddle and come to rest there. That stable path is the saddle path and the point of rest in the center of the saddle is the saddle-point equilibrium.

In the case of the Ramsey model, we fulfill all of the condition required for Pareto optimality, so it is not surprising that Ramsey equilibrium is Pareto optimal. Demonstrating this intuitively is straightforward because all households and individuals in the Ramsey model are identical. Each is maximizing its utility, so without any positive or negative externalities from one household to others, they make socially optimal decisions.

When we want to prove this formally (which we will not), we evaluate the “social planner problem,” contemplating how a benevolent social planner would maximize the utility of the typical household subject to the society-level constraints. In the case of the Ramsey model, this turns out to be exactly the same mathematical problem that the individual household faces, so the household’s solution also maximizes social welfare.

Note that the socially optimal capital/effective-labor ratio k^* is less than the Golden Rule level, which occurs at the peak of the $\dot{k} = 0$ curve. Why is that? Recall that the Golden Rule criterion maximizes the value of y in the steady state. Suppose that our economy began at the Golden Rule level of k . Households could, if they wanted, increase their consumption slightly. This would increase their utility now, but lower it over the steady-state future. Because households are assumed to have a positive preference for consumption now over the future, they would prefer this alternative. Our households with positive ρ (more precisely, $\rho > n + (1-\theta)g$) are not patient enough to value only steady-state income. They will always choose a consumption path that has too much consumption now to end up at the Golden Rule level of capital in the steady-state future.

Is this a bad thing? Well, not if you are a household with the preferences that we have assumed ($\rho > n + (1-\theta)g$). We have just argued that such households are *better off* with k less than the Golden Rule maximum.

Should households have positive time preference? (Should $\rho > 0$?) That is a philosophical question, not an economic one, but some have argued that society would be “better off” if everyone valued the future as much as the present.

G. Understanding Romer’s Chapter 2, Part B

Consumer behavior in Diamond’s overlapping-generations model

The Ramsey model has several important drawbacks. For example, the assumption of infinite lifetimes is clearly counterfactual (given the present state of medical science) and may lead to misleading conclusions if it leads agents in the model to be

unrealistically forward-looking.¹⁶ The infinite-lifetime model also makes it impossible to model life-cycle or generational effects in which agents save for retirement or leave bequests for their children. To analyze questions for which these effects are important, economists often prefer to use a modeling paradigm called *overlapping generations*, in which agents live a finite number of periods (usually two) and experience a working period and a retirement period. At every moment there is at least one generation working and at least one generation that is not working. Among the interesting issues that can be addressed with such a model are the interaction between retired and working generations and households' behavior in saving for their own retirement. The overlapping-generations model is a natural framework for analyzing such policy issues as Social Security reform.

The two major differences between the Diamond model and the Ramsey model are the infinite-lifetime vs. the overlapping-generations assumption and the use of discrete vs. continuous time. In most other ways, the assumptions of the two models are similar or identical, as are their answers to many questions. For example, the utility function in the Diamond model is given by Romer's equation (2.43). The instantaneous utility function has an identical CRRA form to the one we used in the Ramsey model. The discounting process is similar with two exceptions: (1) time is discrete so the discrete-time discounting formula is used, and (2) agents live only two discrete periods so a two-period sum replaces the infinite integral.

The equation of motion for consumption is Romer's equation (2.49), which is reproduced below:

$$\frac{C_{2,t+1}}{C_{1,t}} = \left[\frac{1 + r_{t+1}}{1 + \rho} \right]^{\frac{1}{\theta}}. \quad (12)$$

Note the similarities between equation (12) and the continuous-time version, equation (9). In both cases, the growth of the individual's consumption from one period to the next depends on the relationship between r and ρ . If the interest rate exceeds the rate of time preference, then individuals will desire a rising pattern of consumption over time. If the interest rate is less than the rate of time preference, then consumption will fall over time. In the borderline case of $r = \rho$, the desired consumption path will be constant over time. As in the continuous-time model, $1/\theta$ measures the sensitivity of consumption patterns to differences between the interest rate and the rate of time preferences. Thus, consumption behavior is essentially the same in the two models.

¹⁶ The infinite-lifetime model is sometimes justified by taking a "dynastic" view of the household. This view incorporates a particular assumption about bequests: that the current generation values the utility of future generations exactly as if they were extensions of the current generation, with future utility discounted at the same rate over all future years.

In the infinite-horizon model, it is not possible to achieve a *closed* solution for current consumption from equation (9) and the budget constraint, although these equations do lead to an *implicit* solution. We were able to characterize the steady state and its properties by using a phase diagram, but we were not able to find an expression for C_t in terms of the other variables of the model. Because the consumer in the overlapping-generations model lives only two finite periods, we can perform such a solution in the Diamond model. Solving Romer's equation (2.49) together with the budget constraint (2.46) yields (2.55), which gives current consumption by the younger generation as a function of income, the interest rate, and the parameters of the utility function.

Steady-state equilibrium in the Diamond model

Because we can solve for a closed-form consumption expression, deriving the steady state in the overlapping-generations model is more direct than in the Ramsey model. With the series of substitutions described on pages 79 to 81, we get (2.60), which gives an implicit relationship between k_{t+1} and k_t that does not involve c . We have simplified the model to allow us to work with only one variable (k), so we just need to characterize the condition that $\Delta k = 0$, or $k_t = k_{t+1}$, in order to find the steady state. As Romer points out, this condition can have multiple solutions that have interesting properties if the utility and production functions fail to satisfy the Inada conditions or otherwise differ from the simple log-utility and Cobb-Douglas cases.

In particular, the possibility raised by panel (d) of Romer's Figure 2.12 is one in which macroeconomists have been very interested. There are three different possible steady-state equilibria in this model. The one to which the economy will converge choose is arbitrary; it depends on initial conditions.

The model of panel (d) raises the possibility of *poverty traps* in economic growth and development. It is possible for an economy to get "stuck" at a low- k , low- y equilibrium at the same time that other economies with similar parameters are at high- k equilibria. Growth economists are always on the lookout for potential explanations of large gaps in equilibrium per-capita income (because we observe these in the world) and models with multiple equilibria could provide such an explanation.

Welfare analysis in the Diamond model

In the Ramsey model, we are able to do some simple welfare analysis: comparing everyone's lifetime well-being under alternative possible states. We are able to do that because all Ramsey households are identical—each household exists at every point in time, each grows at the same rate, and each has the same utility function. Thus, we can determine whether a change in economic conditions leaves the representative household better off or worse off and immediately generalize that outcome to all other households in the economy.

We cannot do welfare analysis quite so neatly in the overlapping-generations framework. The agents in the Diamond model are similar to one another in many ways. They all have the same utility functions and live two periods. However, they do not all live in the *same* two periods. Thus, a change in the growth path can cause one generation's welfare to be improved but another generation to be worse off, even though everyone has the same utility function and maximizes her own utility. Thus, sometimes we cannot evaluate welfare unambiguously—some changes will be good for some generations and bad for others.

The only changes that we can evaluate with confidence are those that make everyone better (or worse) off. This is the *Pareto criterion* for optimality: an equilibrium is Pareto efficient if there is no way to make one individual better off without making someone else worse off. Romer shows on pages 87 through 90 that the equilibrium of the Diamond model may not be Pareto efficient. He gives an example of a situation in which *each* generation can be made better off than at the current equilibrium.

This interesting example is especially timely for debates about the funding of the Social Security system in the United States. The source of the dynamic inefficiency that can arise in the Diamond model is that each generation has only one way of providing for its retirement consumption—saving in the form of capital. Thus, there are two motivations for household saving: (1) enjoying a rising living standard through the positive return to investment (as in the Ramsey model) and (2) simply providing for any consumption at all in retirement (which is not relevant to Ramsey's infinitely lived households). However, only the former motivation applies to society as a whole since society never retires.

This additional motive for private saving makes it possible that the private propensity to save could exceed the social desirability of saving. In order to have enough consumption to thrive in retirement, people may need to save a lot when they are young and accumulate a large amount of capital. This large accumulation could push the marginal product of capital very low—in a limiting case, to zero. While it is obviously not socially desirable for agents with positive time preference to accumulate useless capital (that has a marginal product of zero), households that have no other option for transferring wealth from working to retirement periods might save even with zero or negative rates of return on capital. (A durable good that has no productive use but wears out over time would have a negative rate of return. Such saving in the form of stocks of grain is common among rural households in countries where they have no access to financial markets for borrowing or lending.)

As Romer points out, dynamic inefficiency of this kind can be mitigated through a government policy that redistributes money from the young (workers) to the old (retirees), giving the retirees an additional source of income that does not require saving and capital accumulation. This, of course, is exactly what the current Social Security system does in the United States. Most economists have favored shifting toward a

“fully funded” system in which the current transfers from young to old would be replaced by (possibly institutionally mandated) saving/investment by the young toward their own retirement income. This would be a shift toward a system in which young households would have to accumulate capital rather than receiving retirement transfers from the labor income of the next generation. Compared with today’s system, the fully funded scheme increases the possibility that dynamic inefficiency could arise, though it seems improbable that this is a real threat for the U.S. economy.

H. Government Spending in Growth Models

Until the late 1970s, macroeconomists usually analyzed the effects of fiscal policy—aggregate expenditures of government and how those expenditures are financed—in a static framework such as the *IS/LM* model, looking at one period at a time and relying on simple rules of thumb such as the assumption of a constant saving rate. How will the presence of government spending and taxes affect consumer behavior in our optimal growth models? Romer takes up this question in the latter sections of Part A and Part B of Chapter 2 (Sections 2.7 and 2.12). We shall return to related issues at the end of the book when we discuss fiscal policy.

The effects of government purchases

What happens when the government buys goods and services? The traditional Keynesian approach to fiscal policy says that this rise in aggregate demand leads to an increase in the amount of goods and services produced. We shall study this approach to macroeconomics in more detail later in the course.¹⁷

In the Solow, Ramsey, and Diamond growth models, however, we model *natural output*, so production is determined by the amount of labor and capital resources available in the economy and the economy’s technological capability (represented by the production function and the technology parameter A), not by aggregate demand. An increase in government spending does not directly change the available amounts of resources or the economy’s technology, so goods and services that are purchased by the government must come at the expense of private expenditures on consumption

¹⁷ This is the standard *IS/LM* analysis in which a rightward shift in the *IS* curve causes an increase in aggregate demand and, perhaps in the short run, in output.

and/or investment.¹⁸ This tendency for rises in government spending to cause offsetting declines in private spending is called *crowding out*. This phenomenon is reflected in Romer's equation (2.41): for given levels of output and consumption, an increase in government spending lowers investment.

From this description, it looks as though the effects of government spending in the Ramsey model are entirely negative. Consumption yields utility directly and investment provides future utility through greater productive capacity, so increases in government spending that reduce one or both of these apparently must reduce utility. This is, of course, unrealistic. Most of government spending is on public goods such as national defense, police protection, highways, and education (Is education really a public good?) that either yield utility or make the economy more productive. To capture this positive effect of government spending we would need to include government-provided goods and services in consumers' utility functions and/or to add government goods and services (or accumulated government capital) to the production function.

Romer chooses not to do this in Chapter 2. The absence of a way of measuring how government spending is useful prevents us from being able to use this model to assess the welfare impacts of changes in the size of government. We can however ask questions about how the size of government affects other variables in the model, including private consumption and private capital accumulation in the steady state.¹⁹

Romer's model of government spending is simple to analyze because changes in G do not affect the Euler equations for \dot{C} or \dot{c} . However, the curve representing no change in the capital/effective-labor ratio ($\dot{k} = 0$) shifts downward, as shown by negative sign on $G(t)$ in equation (2.41). As Romer shows in equation (2.42), an increase in government spending affects consumption by reducing the amount of "disposable" income available. This decline in the level of consumption with no change in the slope

¹⁸ In the long run, government expenditures and taxes may affect real output indirectly in several ways. Changes in tax policy may affect the incentives to save and invest, leading to changes in private capital accumulation. The government may also invest in infrastructure or in research and development, which might lead over time to greater productivity.

¹⁹ We are assuming that changes in the level of government spending affect consumers *only* through their budget constraints. Romer alludes on page 73 to the fact that the outcome would be the same if "utility equals the sum of utility from private consumption and utility from government-provided goods." This means that we could have a utility function where utility depends on $u(C) + v(G)$. As long as utility is additive in this way, the marginal utility of C , which is what matters for the consumption/saving decision, is not affected by the level of G , so the analysis is still simple. If the level of public goods affects the *marginal* utility of private consumption, the analysis becomes more difficult (though still tractable), so it is less suitable for a textbook explanation.

of its time path is reflected in Romer's Figure 2.8 by the downward shift in the saddle path with an unchanged $\dot{c} = 0$ line.

Because consumption depends on both current and future disposable income, permanent increases in government spending will have different effects than temporary ones. The permanent case is quite straightforward. The once-and-for-all downward shift in the $\dot{k} = 0$ locus moves the steady-state equilibrium from E to E' in Figure 2.8. We know from our earlier study of the dynamics of saddle-point equilibria that consumption must jump vertically to the new saddle path, then converge along the saddle path to the new steady state. In this case, the point on the saddle path directly below E is the new steady state E' , so the economy jumps immediately to the new steady state with lower consumption and an unchanged capital/labor ratio. In the Ramsey model with the marginal utility of private consumption independent of government spending, permanent increases in government spending crowd out consumption dollar for dollar.

The analysis of temporary, unexpected increases in government spending is more difficult and interesting. The exercise that Romer describes is as follows: Before time t_0 , the economy is in a steady state with low government spending at G_L . At time t_0 , everyone discovers that government spending is going to be at the higher level G_H until time t_1 , when it will return to G_L .

Those accustomed to thinking of policy effects in static terms might predict that the temporary increase would move the economy temporarily from E to E' in Figure 2.8, then back to E when the increase in government spending was reversed. This is what would occur under a different assumption about information: if the increase in G at t_0 was thought to be permanent at the time it happened but then turned out to be temporary, so that consumers were surprised again at t_1 when G goes back down.

However, the lifetime nature of the consumption decision implies that this simple move to E' and back cannot be correct if people correctly perceive that the change is temporary. If the change in government spending is only temporary, it has a smaller effect on lifetime disposable income (the right-hand side of equation (2.41)) than if the change is permanent. Thus, households would attempt to substitute intertemporally to smooth consumption, reducing consumption less at the current time than they would have to if the change were permanent. This means that c declines part way, but not all the way to E' , as shown in the top panel of Romer's Figure 2.9. The magnitude of the reduction in c depends on how long the increase in G is expected to last. The longer the increase lasts, the greater is the decline in lifetime disposable income and the greater is the decline in consumption. E' is the limiting case: an increase in G of infinite duration. A very short-lived change in G would cause a very small decline in c , as shown (sort of) in the bottom panel of Romer's Figure 2.9.

The point below E to which the economy initially moves is not on a stable saddle path for either the high- G steady-state equilibrium or the low- G one. This might seem

a little bit unsettling, since it implies a dynamic path that, if it continued forever, would send the economy to $\dot{k} = 0$. But this is precisely the point: because the change is temporary, the economy will *not* continue on this unstable path forever, only until government spending returns to its lower level. Because consumers know that their disposable incomes will rise in the future, they can consume more than the level (E) that they could sustain if their disposable income was going to be permanently lower.

The dynamics of the economy from the point below E are governed initially by its position relative to the temporary, high- G equilibrium E' . The direction of motion at this point is straight to the left, since the point is on the $\dot{c} = 0$ locus and above the $\dot{k} = 0$ locus, which means that $\dot{c} = 0$ and $\dot{k} < 0$. As soon as the economy begins moving to the left, it leaves the $\dot{c} = 0$ locus and moves into the region where $\dot{c} > 0$, so it begins to turn upward and move in a northwesterly direction.

As noted above, if the economy were to stay on this path forever, it would eventually head into oblivion with k falling to zero. However, at time t_1 government spending falls back to its original level, which shifts the $\dot{k} = 0$ locus and the saddle path back to the upper position. At t_1 , the economy *must be exactly on the saddle path* leading back to E . How can we be sure that this will occur? We know that it must because that is the only way that consumers will exactly exhaust their lifetime budget constraint. The amount of the vertical decline in c at time t_0 *must be exactly* the amount that puts the economy on an unstable (northwesterly) path that intersects the stable (northeasterly) saddle path to E at exactly time t_1 . Thus, the path followed by c and k resembles a triangle: an initial vertical drop followed by movement up and to the left then up and to the right, as shown in Figure 2.9.

As Romer points out, a testable implication of this model can be derived by noting that interest rates should track the marginal product of capital, moving in the opposite direction of the capital/effective-labor ratio according to the pattern of panel (b) in Figure 2.9. Wars seem like a naturally occurring experiment with temporary increases in government spending, thus they have been the basis of several empirical tests such as Barro (1987). One might raise several objections to using war periods as tests of this hypothesis, however. For one, agents do not know exactly how long a war is going to last, so a more sophisticated model with uncertainty about t_1 might be more appropriate. For another, economies undergo significant structural change during wars, which might affect consumers' incomes in other ways. Some governments have applied price controls or other forms of non-market resource allocation during wars that could distort consumers' and producers' decisions. Finally, wars are usually periods of intense patriotism, which might cause consumers' preferences about work and consumption to be different than in peaceful periods. To summarize these papers, tests using wartime data have been quite supportive of the theory for the United Kingdom, but less so for the United States.

I. Suggestions for Further Reading

Original expositions of the models

Ramsey, Frank P., "A Mathematical Theory of Saving," *Economic Journal* 38(4), December 1928, 543–559.

Cass, David, "Optimum Growth in an Aggregative Model of Capital Accumulation," *Review of Economic Studies* 32(3), July 1965, 233–240.

Koopmans, Tjalling C., "On the Concept of Optimal Economic Growth," in *The Economic Approach to Development Planning* (Amsterdam: Elsevier, 1965).

Diamond, Peter A., "National Debt in a Neoclassical Growth Model," *American Economic Review* 55(5), December 1965, 1126–1150.

Alternative presentations and mathematical methods

Barro, Robert J., and Xavier Sala-i-Martin, *Economic Growth*, 2nd ed., (Cambridge, Mass.: MIT Press, 2004), Chapter 2. (An alternative presentation at a slightly higher level than Romer.)

Chiang, Alpha C., *Elements of Dynamic Optimization* (New York: McGraw-Hill, 1992). (A fairly sophisticated introduction to the dynamic techniques used in this chapter.)

I. Work Cited in Text

Barro, Robert J. 1987. "Government Spending, Interest Rates, Prices, and Budget Deficits in the United Kingdom, 1701-1918." *Journal of Monetary Economics* 20 (2):221-47.

Barro, Robert J. 1999. "Ramsey Meets Laibson in the Neoclassical Growth Model." *Quarterly Journal of Economics* 114 (4):1125-52.

Barro, Robert J., and Xavier Sala-i-Martin. 2004. *Economic Growth*. 2nd ed. Cambridge, Mass.: MIT Press.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (2):443-477.