


part VI

More on the Microeconomics Behind Macroeconomics



CHAPTER

SIXTEEN

Consumption

Consumption is the sole end and purpose of all production.

— Adam Smith

How do households decide how much of their income to consume today and how much to save for the future? This is a microeconomic question because it addresses the behavior of individual decisionmakers. Yet its answer has macroeconomic consequences. As we have seen in previous chapters, households' consumption decisions affect the way the economy as a whole behaves both in the long run and in the short run.

The consumption decision is crucial for long-run analysis because of its role in economic growth. The Solow growth model of Chapters 7 and 8 shows that the saving rate is a key determinant of the steady-state capital stock and thus of the level of economic well-being. The saving rate measures how much of its income the present generation is putting aside for its own future and for future generations.

The consumption decision is crucial for short-run analysis because of its role in determining aggregate demand. Consumption is two-thirds of GDP, so fluctuations in consumption are a key element of booms and recessions. The *IS-LM* model of Chapters 10 and 11 shows that changes in consumers' spending plans can be a source of shocks to the economy, and that the marginal propensity to consume is a determinant of the fiscal-policy multipliers.

In previous chapters we explained consumption with a function that relates consumption to disposable income: $C = C(Y - T)$. This approximation allowed us to develop simple models for long-run and short-run analysis, but it is too simple to provide a complete explanation of consumer behavior. In this chapter we examine the consumption function in greater detail and develop a more thorough explanation of what determines aggregate consumption.

Since macroeconomics began as a field of study, many economists have written about the theory of consumer behavior and suggested alternative ways of interpreting the data on consumption and income. This chapter presents the views of six prominent economists to show the diverse approaches to explaining consumption.

16-1 John Maynard Keynes and the Consumption Function

We begin our study of consumption with John Maynard Keynes's *General Theory*, which was published in 1936. Keynes made the consumption function central to his theory of economic fluctuations, and it has played a key role in macroeconomic analysis ever since. Let's consider what Keynes thought about the consumption function, and then see what puzzles arose when his ideas were confronted with the data.

Keynes's Conjectures

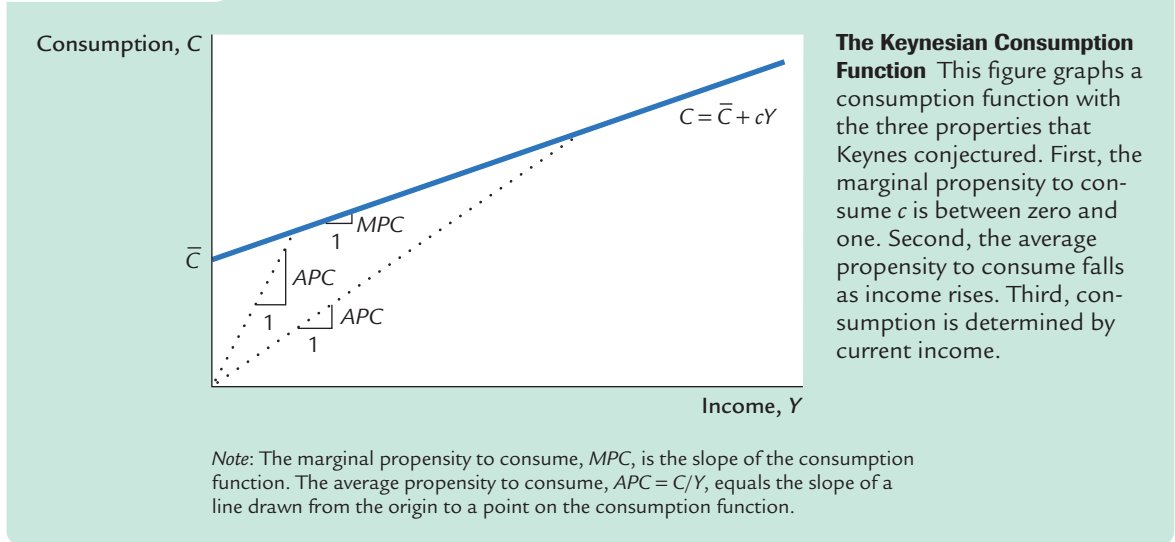
Today, economists who study consumption rely on sophisticated techniques of data analysis. With the help of computers, they analyze aggregate data on the behavior of the overall economy from the national income accounts and detailed data on the behavior of individual households from surveys. Because Keynes wrote in the 1930s, however, he had neither the advantage of these data nor the computers necessary to analyze such large data sets. Instead of relying on statistical analysis, Keynes made conjectures about the consumption function based on introspection and casual observation.

First and most important, Keynes conjectured that the **marginal propensity to consume**—the amount consumed out of an additional dollar of income—is between zero and one. He wrote that the “fundamental psychological law, upon which we are entitled to depend with great confidence, . . . is that men are disposed, as a rule and on the average, to increase their consumption as their income increases, but not by as much as the increase in their income.” That is, when a person earns an extra dollar, he typically spends some of it and saves some of it. As we saw in Chapter 10 when we developed the Keynesian cross, the marginal propensity to consume was crucial to Keynes's policy recommendations for how to reduce widespread unemployment. The power of fiscal policy to influence the economy—as expressed by the fiscal-policy multipliers—arises from the feedback between income and consumption.

Second, Keynes posited that the ratio of consumption to income, called the **average propensity to consume**, falls as income rises. He believed that saving was a luxury, so he expected the rich to save a higher proportion of their income than the poor. Although not essential for Keynes's own analysis, the postulate that the average propensity to consume falls as income rises became a central part of early Keynesian economics.

Third, Keynes thought that income is the primary determinant of consumption and that the interest rate does not have an important role. This conjecture stood in stark contrast to the beliefs of the classical economists who preceded him. The classical economists held that a higher interest rate encourages saving and discourages consumption. Keynes admitted that the interest rate could influence consumption as a matter of theory. Yet he wrote that “the main conclusion suggested

figure 16-1



by experience, I think, is that the short-period influence of the rate of interest on individual spending out of a given income is secondary and relatively unimportant.”

On the basis of these three conjectures, the Keynesian consumption function is often written as

$$C = \bar{C} + cY, \quad \bar{C} > 0, \quad 0 < c < 1,$$

where C is consumption, Y is disposable income, \bar{C} is a constant, and c is the marginal propensity to consume. This consumption function, shown in Figure 16-1, is graphed as a straight line.

Notice that this consumption function exhibits the three properties that Keynes posited. It satisfies Keynes’s first property because the marginal propensity to consume c is between zero and one, so that higher income leads to higher consumption and also to higher saving. This consumption function satisfies Keynes’s second property because the average propensity to consume APC is

$$APC = C/Y = \bar{C}/Y + c.$$

As Y rises, \bar{C}/Y falls, and so the average propensity to consume C/Y falls. And finally, this consumption function satisfies Keynes’s third property because the interest rate is not included in this equation as a determinant of consumption.

The Early Empirical Successes

Soon after Keynes proposed the consumption function, economists began collecting and examining data to test his conjectures. The earliest studies indicated that the Keynesian consumption function is a good approximation of how consumers behave.

In some of these studies, researchers surveyed households and collected data on consumption and income. They found that households with higher income consumed more, which confirms that the marginal propensity to consume is greater than zero. They also found that households with higher income saved more, which confirms that the marginal propensity to consume is less than one. In addition, these researchers found that higher-income households saved a larger fraction of their income, which confirms that the average propensity to consume falls as income rises. Thus, these data verified Keynes's conjectures about the marginal and average propensities to consume.

In other studies, researchers examined aggregate data on consumption and income for the period between the two world wars. These data also supported the Keynesian consumption function. In years when income was unusually low, such as during the depths of the Great Depression, both consumption and saving were low, indicating that the marginal propensity to consume is between zero and one. In addition, during those years of low income, the ratio of consumption to income was high, confirming Keynes's second conjecture. Finally, because the correlation between income and consumption was so strong, no other variable appeared to be important for explaining consumption. Thus, the data also confirmed Keynes's third conjecture that income is the primary determinant of how much people choose to consume.

Secular Stagnation, Simon Kuznets, and the Consumption Puzzle

Although the Keynesian consumption function met with early successes, two anomalies soon arose. Both concern Keynes's conjecture that the average propensity to consume falls as income rises.

The first anomaly became apparent after some economists made a dire—and, it turned out, erroneous—prediction during World War II. On the basis of the Keynesian consumption function, these economists reasoned that as incomes in the economy grew over time, households would consume a smaller and smaller fraction of their incomes. They feared that there might not be enough profitable investment projects to absorb all this saving. If so, the low consumption would lead to an inadequate demand for goods and services, resulting in a depression once the wartime demand from the government ceased. In other words, on the basis of the Keynesian consumption function, these economists predicted that the economy would experience what they called *secular stagnation*—a long depression of indefinite duration—unless fiscal policy was used to expand aggregate demand.

Fortunately for the economy, but unfortunately for the Keynesian consumption function, the end of World War II did not throw the country into another depression. Although incomes were much higher after the war than before, these higher incomes did not lead to large increases in the rate of saving. Keynes's conjecture that the average propensity to consume would fall as income rose appeared not to hold.

The second anomaly arose when economist Simon Kuznets constructed new aggregate data on consumption and income dating back to 1869. Kuznets assembled


these data in the 1940s and would later receive the Nobel Prize for this work. He discovered that the ratio of consumption to income was remarkably stable from decade to decade, despite large increases in income over the period he studied. Again, Keynes's conjecture that the average propensity to consume would fall as income rose appeared not to hold.

The failure of the secular-stagnation hypothesis and the findings of Kuznets both indicated that the average propensity to consume is fairly constant over long periods of time. This fact presented a puzzle that motivated much of the subsequent work on consumption. Economists wanted to know why some studies confirmed Keynes's conjectures and others refuted them. That is, why did Keynes's conjectures hold up well in the studies of household data and in the studies of short time-series, but fail when long time-series were examined?

Figure 16-2 illustrates the puzzle. The evidence suggested that there were two consumption functions. For the household data or for the short time-series, the Keynesian consumption function appeared to work well. Yet for the long time-series, the consumption function appeared to have a constant average propensity to consume. In Figure 16-2, these two relationships between consumption and income are called the short-run and long-run consumption functions. Economists needed to explain how these two consumption functions could be consistent with each other.

In the 1950s, Franco Modigliani and Milton Friedman each proposed explanations of these seemingly contradictory findings. Both economists later won

figure 16-2



The Consumption Puzzle

Studies of household data and short time-series found a relationship between consumption and income similar to the one Keynes conjectured. In the figure, this relationship is called the short-run consumption function. But studies of long time-series found that the average propensity to consume did not vary systematically with income. This relationship is called the long-run consumption function. Notice that the short-run consumption function has a falling average propensity to consume, whereas the long-run consumption function has a constant average propensity to consume.

Nobel Prizes, in part because of their work on consumption. But before we see how Modigliani and Friedman tried to solve the consumption puzzle, we must discuss Irving Fisher's contribution to consumption theory. Both Modigliani's life-cycle hypothesis and Friedman's permanent-income hypothesis rely on the theory of consumer behavior proposed much earlier by Irving Fisher.

16-2 Irving Fisher and Intertemporal Choice

The consumption function introduced by Keynes relates current consumption to current income. This relationship, however, is incomplete at best. When people decide how much to consume and how much to save, they consider both the present and the future. The more consumption they enjoy today, the less they will be able to enjoy tomorrow. In making this tradeoff, households must look ahead to the income they expect to receive in the future and to the consumption of goods and services they hope to be able to afford.

The economist Irving Fisher developed the model with which economists analyze how rational, forward-looking consumers make intertemporal choices—that is, choices involving different periods of time. Fisher's model illuminates the constraints consumers face, the preferences they have, and how these constraints and preferences together determine their choices about consumption and saving.

The Intertemporal Budget Constraint

Most people would prefer to increase the quantity or quality of the goods and services they consume—to wear nicer clothes, eat at better restaurants, or see more movies. The reason people consume less than they desire is that their consumption is constrained by their income. In other words, consumers face a limit on how much they can spend, called a *budget constraint*. When they are deciding how much to consume today versus how much to save for the future, they face an **intertemporal budget constraint**, which measures the total resources available for consumption today and in the future. Our first step in developing Fisher's model is to examine this constraint in some detail.

To keep things simple, we examine the decision facing a consumer who lives for two periods. Period one represents the consumer's youth, and period two represents the consumer's old age. The consumer earns income Y_1 and consumes C_1 in period one, and earns income Y_2 and consumes C_2 in period two. (All variables are real—that is, adjusted for inflation.) Because the consumer has the opportunity to borrow and save, consumption in any single period can be either greater or less than income in that period.

Consider how the consumer's income in the two periods constrains consumption in the two periods. In the first period, saving equals income minus consumption. That is,

$$S = Y_1 - C_1,$$

where S is saving. In the second period, consumption equals the accumulated saving, including the interest earned on that saving, plus second-period income. That is,

$$C_2 = (1 + r)S + Y_2,$$

where r is the real interest rate. For example, if the interest rate is 5 percent, then for every \$1 of saving in period one, the consumer enjoys an extra \$1.05 of consumption in period two. Because there is no third period, the consumer does not save in the second period.

Note that the variable S can represent either saving or borrowing and that these equations hold in both cases. If first-period consumption is less than first-period income, the consumer is saving, and S is greater than zero. If first-period consumption exceeds first-period income, the consumer is borrowing, and S is less than zero. For simplicity, we assume that the interest rate for borrowing is the same as the interest rate for saving.

To derive the consumer's budget constraint, combine the two preceding equations. Substitute the first equation for S into the second equation to obtain

$$C_2 = (1 + r)(Y_1 - C_1) + Y_2.$$

To make the equation easier to interpret, we must rearrange terms. To place all the consumption terms together, bring $(1 + r)C_1$ from the right-hand side to the left-hand side of the equation to obtain

$$(1 + r)C_1 + C_2 = (1 + r)Y_1 + Y_2.$$

Now divide both sides by $1 + r$ to obtain

$$C_1 + \frac{C_2}{1 + r} = Y_1 + \frac{Y_2}{1 + r}.$$

This equation relates consumption in the two periods to income in the two periods. It is the standard way of expressing the consumer's intertemporal budget constraint.

The consumer's budget constraint is easily interpreted. If the interest rate is zero, the budget constraint shows that total consumption in the two periods equals total income in the two periods. In the usual case in which the interest rate is greater than zero, future consumption and future income are discounted by a factor $1 + r$. This **discounting** arises from the interest earned on savings. In essence, because the consumer earns interest on current income that is saved, future income is worth less than current income. Similarly, because future consumption is paid for out of savings that have earned interest, future consumption costs less than current consumption. The factor $1/(1 + r)$ is the price of second-period consumption measured in terms of first-period consumption: it is the amount of first-period consumption that the consumer must forgo to obtain 1 unit of second-period consumption.

Figure 16-3 graphs the consumer's budget constraint. Three points are marked on this figure. At point A, the consumer consumes exactly his income in each

FYI Present Value, or Why a \$1,000,000 Prize Is Worth Only \$623,000

The use of discounting in the consumer's budget constraint illustrates an important fact of economic life: a dollar in the future is less valuable than a dollar today. This is true because a dollar today can be deposited in an interest-bearing bank account and produce more than one dollar in the future. If the interest rate is 5 percent, for instance, then a dollar today can be turned to \$1.05 dollars next year, \$1.1025 in two years, \$1.1576 in three years, . . . , or \$2.65 in 20 years.

Economists use a concept called *present value* to compare dollar amounts from different times. The present value of any amount in the future is the amount that would be needed today, given available interest rates, to produce that future amount. Thus, if you are going to be paid X dollars in T years and the interest rate is r , then the present value of that payment is


$$\text{Present Value} = X / (1 + r)^T.$$

In light of this definition, we can see a new interpretation of the consumer's budget constraint in our two-period consumption problem. The intertemporal budget constraint states that the present value of consumption must equal the present value of income.

The concept of present value has many applications. Suppose, for instance, that you won a million-dollar lottery. Such prizes are usually paid out over time—say, \$50,000 a year for 20 years. What is the present value of such a delayed prize? By applying the preceding formula for each of the 20 payments and adding up the result, we learn that the million-dollar prize, discounted at an interest rate of 5 percent, has a present value of only \$623,000. (If the prize were paid out as a dollar a year for a million years, the present value would be a mere \$20!) Sometimes a million dollars isn't all it's cracked up to be.

period ($C_1 = Y_1$ and $C_2 = Y_2$), so there is neither saving nor borrowing between the two periods. At point B, the consumer consumes nothing in the first period ($C_1 = 0$) and saves all income, so second-period consumption C_2 is $(1 + r)Y_1 + Y_2$. At point C, the consumer plans to consume nothing in the second period ($C_2 = 0$) and borrows as much as possible against second-period income, so first-period

figure 16-3



The Consumer's Budget Constraint This figure shows the combinations of first-period and second-period consumption the consumer can choose. If he chooses points between A and B, he consumes less than his income in the first period and saves the rest for the second period. If he chooses points between A and C, he consumes more than his income in the first period and borrows to make up the difference.

consumption C_1 is $Y_1 + Y_2/(1 + r)$. Of course, these are only three of the many combinations of first- and second-period consumption that the consumer can afford: all the points on the line from B to C are available to the consumer.

Consumer Preferences

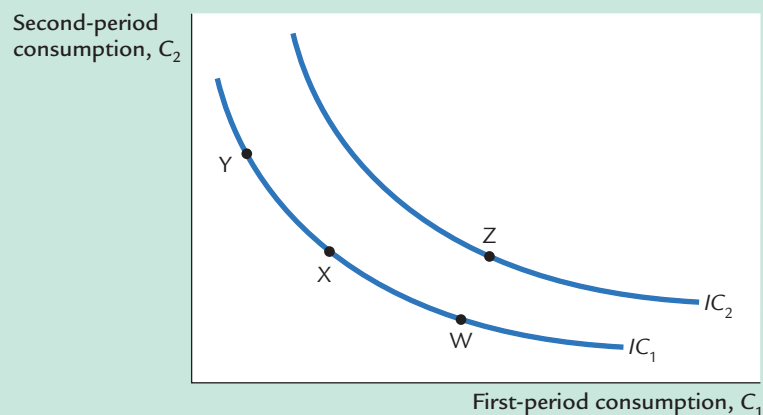
The consumer's preferences regarding consumption in the two periods can be represented by **indifference curves**. An indifference curve shows the combinations of first-period and second-period consumption that make the consumer equally happy.

Figure 16-4 shows two of the consumer's many indifference curves. The consumer is indifferent among combinations W, X, and Y, because they are all on the same curve. Not surprisingly, if the consumer's first-period consumption is reduced, say from point W to point X, second-period consumption must increase to keep him equally happy. If first-period consumption is reduced again, from point X to point Y, the amount of extra second-period consumption he requires for compensation is greater.

The slope at any point on the indifference curve shows how much second-period consumption the consumer requires in order to be compensated for a 1-unit reduction in first-period consumption. This slope is the **marginal rate of substitution** between first-period consumption and second-period consumption. It tells us the rate at which the consumer is willing to substitute second-period consumption for first-period consumption.

Notice that the indifference curves in Figure 16-4 are not straight lines and, as a result, the marginal rate of substitution depends on the levels of consumption in the two periods. When first-period consumption is high and second-period consumption is low, as at point W, the marginal rate of substitution is low: the consumer requires only a little extra second-period consumption to give up

figure 16-4



The Consumer's Preferences

Indifference curves represent the consumer's preferences over first-period and second-period consumption. An indifference curve gives the combinations of consumption in the two periods that make the consumer equally happy. This figure shows two of many indifference curves. Higher indifference curves such as IC_2 are preferred to lower curves such as IC_1 . The consumer is equally happy at points W, X, and Y, but prefers point Z to points W, X, or Y.

1 unit of first-period consumption. When first-period consumption is low and second-period consumption is high, as at point Y, the marginal rate of substitution is high: the consumer requires much additional second-period consumption to give up 1 unit of first-period consumption.

The consumer is equally happy at all points on a given indifference curve, but he prefers some indifference curves to others. Because he prefers more consumption to less, he prefers higher indifference curves to lower ones. In Figure 16-4, the consumer prefers the points on curve IC_2 to the points on curve IC_1 .

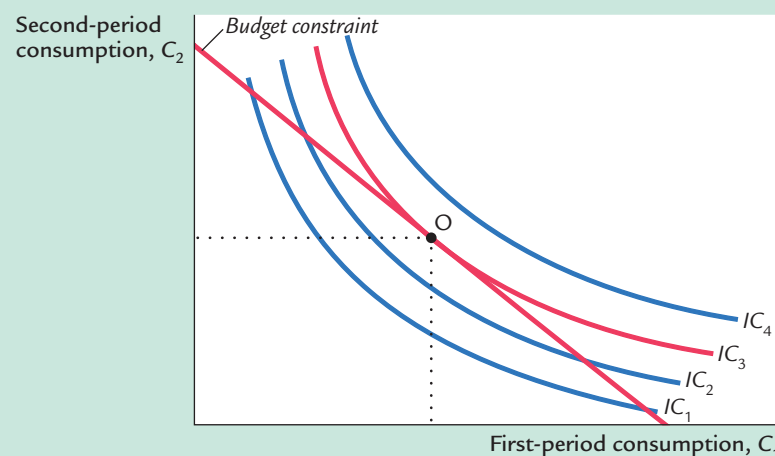
The set of indifference curves gives a complete ranking of the consumer's preferences. It tells us that the consumer prefers point Z to point W, but that may be obvious because point Z has more consumption in both periods. Yet compare point Z and point Y: point Z has more consumption in period one and less in period two. Which is preferred, Z or Y? Because Z is on a higher indifference curve than Y, we know that the consumer prefers point Z to point Y. Hence, we can use the set of indifference curves to rank any combinations of first-period and second-period consumption.

Optimization

Having discussed the consumer's budget constraint and preferences, we can consider the decision about how much to consume. The consumer would like to end up with the best possible combination of consumption in the two periods—that is, on the highest possible indifference curve. But the budget constraint requires that the consumer also end up on or below the budget line, because the budget line measures the total resources available to him.

Figure 16-5 shows that many indifference curves cross the budget line. The highest indifference curve that the consumer can obtain without violating the budget constraint is the indifference curve that just barely touches the budget

figure 16-5



The Consumer's Optimum

The consumer achieves his highest level of satisfaction by choosing the point on the budget constraint that is on the highest indifference curve. At the optimum, the indifference curve is tangent to the budget constraint.

line, which is curve IC_3 in the figure. The point at which the curve and line touch—point O for “optimum”—is the best combination of consumption in the two periods that the consumer can afford.

Notice that, at the optimum, the slope of the indifference curve equals the slope of the budget line. The indifference curve is *tangent* to the budget line. The slope of the indifference curve is the marginal rate of substitution MRS , and the slope of the budget line is 1 plus the real interest rate. We conclude that at point O,

$$MRS = 1 + r.$$

The consumer chooses consumption in the two periods so that the marginal rate of substitution equals 1 plus the real interest rate.

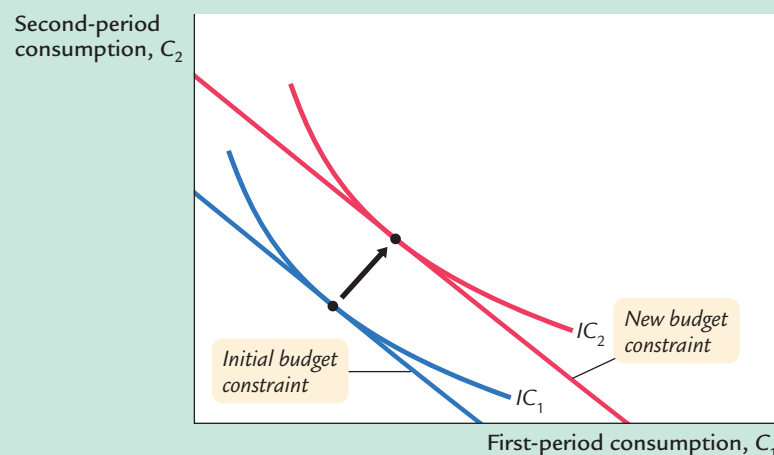
How Changes in Income Affect Consumption

Now that we have seen how the consumer makes the consumption decision, let's examine how consumption responds to an increase in income. An increase in either Y_1 or Y_2 shifts the budget constraint outward, as in Figure 16-6. The higher budget constraint allows the consumer to choose a better combination of first- and second-period consumption—that is, the consumer can now reach a higher indifference curve.

In Figure 16-6, the consumer responds to the shift in his budget constraint by choosing more consumption in both periods. Although it is not implied by the logic of the model alone, this situation is the most usual. If a consumer wants more of a good when his or her income rises, economists call it a **normal good**. The indifference curves in Figure 16-6 are drawn under the assumption that consumption in period one and consumption in period two are both normal goods.

The key conclusion from Figure 16-6 is that regardless of whether the increase in income occurs in the first period or the second period, the consumer spreads it over consumption in both periods. This behavior is sometimes called

figure 16-6



An Increase in Income An increase in either first-period income or second-period income shifts the budget constraint outward. If consumption in period one and consumption in period two are both normal goods, this increase in income raises consumption in both periods.

consumption smoothing. Because the consumer can borrow and lend between periods, the timing of the income is irrelevant to how much is consumed today (except, of course, that future income is discounted by the interest rate). The lesson of this analysis is that consumption depends on the present value of current and future income—that is, on

$$\text{Present Value of Income} = Y_1 + \frac{Y_2}{1 + r}.$$

Notice that this conclusion is quite different from that reached by Keynes. *Keynes posited that a person's current consumption depends largely on his current income. Fisher's model says, instead, that consumption is based on the resources the consumer expects over his lifetime.*

How Changes in the Real Interest Rate Affect Consumption

Let's now use Fisher's model to consider how a change in the real interest rate alters the consumer's choices. There are two cases to consider: the case in which the consumer is initially saving and the case in which he is initially borrowing. Here we discuss the saving case; Problem 1 at the end of the chapter asks you to analyze the borrowing case.

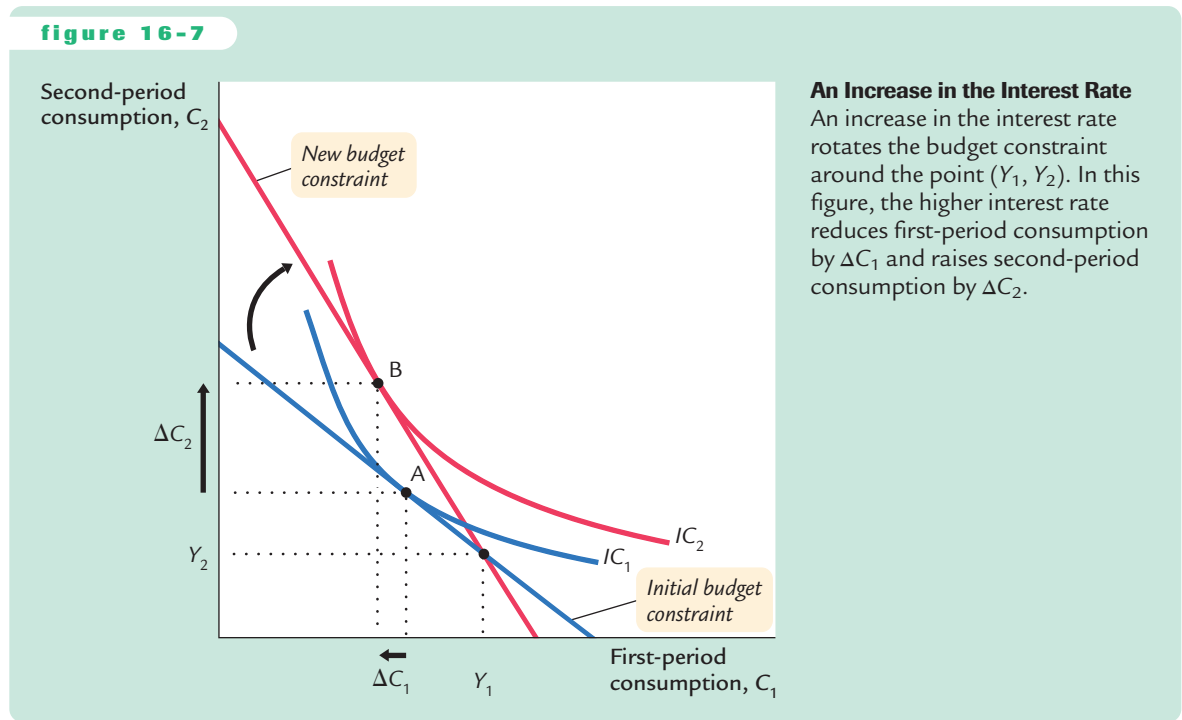
Figure 16-7 shows that an increase in the real interest rate rotates the consumer's budget line around the point (Y_1, Y_2) and, thereby, alters the amount of consumption he chooses in both periods. Here, the consumer moves from point A to point B. You can see that for the indifference curves drawn in this figure first-period consumption falls and second-period consumption rises.

Economists decompose the impact of an increase in the real interest rate on consumption into two effects: an **income effect** and a **substitution effect**. Textbooks in microeconomics discuss these effects in detail. We summarize them briefly here.

The income effect is the change in consumption that results from the movement to a higher indifference curve. Because the consumer is a saver rather than a borrower (as indicated by the fact that first-period consumption is less than first-period income), the increase in the interest rate makes him better off (as reflected by the movement to a higher indifference curve). If consumption in period one and consumption in period two are both normal goods, the consumer will want to spread this improvement in his welfare over both periods. This income effect tends to make the consumer want more consumption in both periods.

The substitution effect is the change in consumption that results from the change in the relative price of consumption in the two periods. In particular, consumption in period two becomes less expensive relative to consumption in period one when the interest rate rises. That is, because the real interest rate earned on saving is higher, the consumer must now give up less first-period consumption to obtain an extra unit of second-period consumption. This substitution effect tends to make the consumer choose more consumption in period two and less consumption in period one.

The consumer's choice depends on both the income effect and the substitution effect. Both effects act to increase the amount of second-period consumption;



hence, we can confidently conclude that an increase in the real interest rate raises second-period consumption. But the two effects have opposite impacts on first-period consumption, so the increase in the interest rate could either raise or lower it. Hence, depending on the relative size of the income and substitution effects, an increase in the interest rate could either stimulate or depress saving.

Constraints on Borrowing


Fisher’s model assumes that the consumer can borrow as well as save. The ability to borrow allows current consumption to exceed current income. In essence, when the consumer borrows, he consumes some of his future income today. Yet for many people such borrowing is impossible. For example, a student wishing to enjoy spring break in Florida would probably be unable to finance this vacation with a bank loan. Let’s examine how Fisher’s analysis changes if the consumer cannot borrow.

The inability to borrow prevents current consumption from exceeding current income. A constraint on borrowing can therefore be expressed as

$$C_1 \leq Y_1.$$

This inequality states that consumption in period one must be less than or equal to income in period one. This additional constraint on the consumer is called a **borrowing constraint** or, sometimes, a *liquidity constraint*.

Figure 16-8 shows how this borrowing constraint restricts the consumer's set of choices. The consumer's choice must satisfy both the intertemporal budget constraint and the borrowing constraint. The shaded area represents the combinations of first-period consumption and second-period consumption that satisfy both constraints.



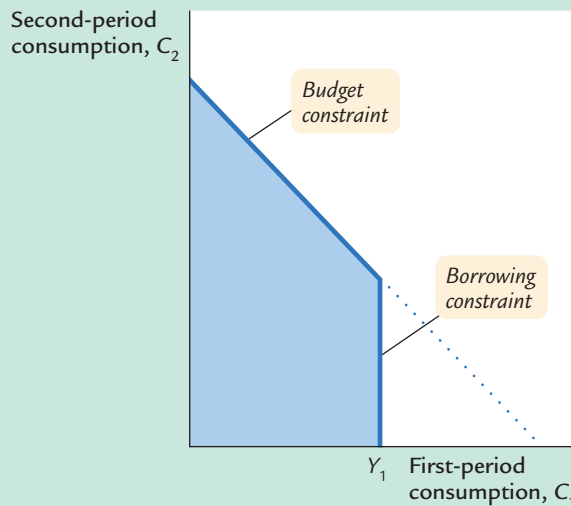
Drawing by Handelsman, © 1985 The New Yorker Magazine, Inc.

“What I’d like, basically, is a temporary line of credit just to tide me over the rest of my life.”

Figure 16-9 shows how this borrowing constraint affects the consumption decision. There are two possibilities. In panel (a), the consumer wishes to consume less in period one than he earns. The borrowing constraint is not binding and, therefore, does not affect consumption. In panel (b), the consumer would like to choose point D, where he consumes more in period one than he earns, but the borrowing constraint prevents this outcome. The best the consumer can do is to consume his first-period income, represented by point E.


The analysis of borrowing constraints leads us to conclude that there are two consumption functions. For some consumers, the borrowing constraint is not binding, and consumption in both periods depends on the present value of lifetime income, $Y_1 + [Y_2/(1 + r)]$. For other consumers, the borrowing constraint binds, and the consumption function is $C_1 = Y_1$ and $C_2 = Y_2$. Hence, for those consumers who would like to borrow but cannot, consumption depends only on current income.

figure 16-8



A Borrowing Constraint If the consumer cannot borrow, he faces the additional constraint that first-period consumption cannot exceed first-period income. The shaded area represents the combination of first-period and second-period consumption the consumer can choose.

figure 16-9



The Consumer's Optimum With a Borrowing Constraint When the consumer faces a borrowing constraint, there are two possible situations. In panel (a), the consumer chooses first-period consumption to be less than first-period income, so the borrowing constraint is not binding and does not affect consumption in either period. In panel (b), the borrowing constraint is binding. The consumer would like to borrow and choose point D. But because borrowing is not allowed, the best available choice is point E. When the borrowing constraint is binding, first-period consumption equals first-period income.

CASE STUDY

The High Japanese Saving Rate

Japan has one of the world's highest saving rates, and this fact is important for understanding both the long-run and short-run performance of its economy. On the one hand, many economists believe that the high Japanese saving rate is a key to the rapid growth Japan experienced in the decades after World War II. Indeed, the Solow growth model developed in Chapters 7 and 8 shows that the saving rate is a primary determinant of a country's steady-state level of income. On other other hand, some economists have argued that the high Japanese saving rate contributed to Japan's slump during the 1990s. High saving means low consumer spending, which according to the *IS-LM* model of Chapters 10 and 11 translates into low aggregate demand and reduced income.

Why do the Japanese consume a much smaller fraction of their income than Americans? One reason is that it is harder for households to borrow in

Japan. As Fisher's model shows, a household facing a binding borrowing constraint consumes less than it would without the borrowing constraint. Hence, societies in which borrowing constraints are common will tend to have higher rates of saving.

One reason that households often wish to borrow is to buy a home. In the United States, a person can usually buy a home with a down payment of 10 percent. A homebuyer in Japan cannot borrow nearly this much: down payments of 40 percent are common. Moreover, housing prices are very high in Japan, primarily because land prices are high. A Japanese family must save a great deal if it is ever to afford its own home.

Although constraints on borrowing are part of the explanation of high Japanese saving, there are many other differences between Japan and the United States that contribute to the difference in the saving rates. The Japanese tax system encourages saving by taxing capital income very lightly. In addition, cultural differences may lead to differences in consumer preferences regarding present and future consumption. One prominent Japanese economist writes, "The Japanese are simply *different*. They are more risk averse and more patient. If this is true, the long-run implication is that Japan will absorb all the wealth in the world. I refuse to comment on this explanation."¹

16-3 Franco Modigliani and the Life-Cycle Hypothesis

In a series of papers written in the 1950s, Franco Modigliani and his collaborators Albert Ando and Richard Brumberg used Fisher's model of consumer behavior to study the consumption function. One of their goals was to solve the consumption puzzle—that is, to explain the apparently conflicting pieces of evidence that came to light when Keynes's consumption function was brought to the data. According to Fisher's model, consumption depends on a person's lifetime income. Modigliani emphasized that income varies systematically over people's lives and that saving allows consumers to move income from those times in life when income is high to those times when it is low. This interpretation of consumer behavior formed the basis for his **life-cycle hypothesis**.²

¹ Fumio Hayashi, "Why Is Japan's Saving Rate So Apparently High?" *NBER Macroeconomics Annual* (1986): 147–210.

² For references to the large body of work on the life-cycle hypothesis, a good place to start is the lecture Modigliani gave when he won the Nobel Prize. Franco Modigliani, "Life Cycle, Individual Thrift, and the Wealth of Nations," *American Economic Review* 76 (June 1986): 297–313.

The Hypothesis

One important reason that income varies over a person's life is retirement. Most people plan to stop working at about age 65, and they expect their incomes to fall when they retire. Yet they do not want a large drop in their standard of living, as measured by their consumption. To maintain consumption after retirement, people must save during their working years. Let's see what this motive for saving implies for the consumption function.

Consider a consumer who expects to live another T years, has wealth of W , and expects to earn income Y until she retires R years from now. What level of consumption will the consumer choose if she wishes to maintain a smooth level of consumption over her life?

The consumer's lifetime resources are composed of initial wealth W and lifetime earnings of $R \times Y$. (For simplicity, we are assuming an interest rate of zero; if the interest rate were greater than zero, we would need to take account of interest earned on savings as well.) The consumer can divide up her lifetime resources among her T remaining years of life. We assume that she wishes to achieve the smoothest possible path of consumption over her lifetime. Therefore, she divides this total of $W + RY$ equally among the T years and each year consumes

$$C = (W + RY)/T.$$

We can write this person's consumption function as

$$C = (1/T)W + (R/T)Y.$$

For example, if the consumer expects to live for 50 more years and work for 30 of them, then $T = 50$ and $R = 30$, so her consumption function is

$$C = 0.02W + 0.6Y.$$

This equation says that consumption depends on both income and wealth. An extra \$1 of income per year raises consumption by \$0.60 per year, and an extra \$1 of wealth raises consumption by \$0.02 per year.

If every individual in the economy plans consumption like this, then the aggregate consumption function is much the same as the individual one. In particular, aggregate consumption depends on both wealth and income. That is, the economy's consumption function is


$$C = \alpha W + \beta Y,$$

where the parameter α is the marginal propensity to consume out of wealth, and the parameter β is the marginal propensity to consume out of income.

Implications

Figure 16-10 graphs the relationship between consumption and income predicted by the life-cycle model. For any given level of wealth W , the model yields a conventional consumption function similar to the one shown in Figure 16-1.

figure 16-10



Notice, however, that the intercept of the consumption function, which shows what would happen to consumption if income ever fell to zero, is not a fixed value, as it is in Figure 16-1. Instead, the intercept here is αW and, thus, depends on the level of wealth.

This life-cycle model of consumer behavior can solve the consumption puzzle. According to the life-cycle consumption function, the average propensity to consume is


$$C/Y = \alpha(W/Y) + \beta.$$

Because wealth does not vary proportionately with income from person to person or from year to year, we should find that high income corresponds to a low average propensity to consume when looking at data across individuals or over short periods of time. But, over long periods of time, wealth and income grow together, resulting in a constant ratio W/Y and thus a constant average propensity to consume.

To make the same point somewhat differently, consider how the consumption function changes over time. As Figure 16-10 shows, for any given level of wealth, the life-cycle consumption function looks like the one Keynes suggested. But this function holds only in the short run when wealth is constant. In the long run, as wealth increases, the consumption function shifts upward, as in Figure 16-11. This upward shift prevents the average propensity to consume from falling as income increases. In this way, Modigliani resolved the consumption puzzle posed by Simon Kuznets's data.

The life-cycle model makes many other predictions as well. Most important, it predicts that saving varies over a person's lifetime. If a person begins adulthood with no wealth, she will accumulate wealth during her working


figure 16-11



How Changes in Wealth Shift the Consumption Function If consumption depends on wealth, then an increase in wealth shifts the consumption function upward. Thus, the short-run consumption function (which holds wealth constant) will not continue to hold in the long run (as wealth rises over time).

years and then run down her wealth during her retirement years. Figure 16-12 illustrates the consumer's income, consumption, and wealth over her adult life. According to the life-cycle hypothesis, because people want to smooth consumption over their lives, the young who are working save, while the old who are retired dissave.

figure 16-12



Consumption, Income, and Wealth Over the Life Cycle If the consumer smooths consumption over her life (as indicated by the horizontal consumption line), she will save and accumulate wealth during her working years and then dissave and run down her wealth during retirement.

CASE STUDY

The Consumption and Saving of the Elderly

Many economists have studied the consumption and saving of the elderly. Their findings present a problem for the life-cycle model. It appears that the elderly do not dissave as much as the model predicts. In other words, the elderly do not run down their wealth as quickly as one would expect if they were trying to smooth their consumption over their remaining years of life.

There are two chief explanations for why the elderly do not dissave to the extent that the model predicts. Each suggests a direction for further research on consumption.

The first explanation is that the elderly are concerned about unpredictable expenses. Additional saving that arises from uncertainty is called **precautionary saving**. One reason for precautionary saving by the elderly is the possibility of living longer than expected and thus having to provide for a longer than average span of retirement. Another reason is the possibility of illness and large medical bills. The elderly may respond to this uncertainty by saving more in order to be better prepared for these contingencies.

The precautionary-saving explanation is not completely persuasive, because the elderly can largely insure against these risks. To protect against uncertainty regarding life span, they can buy *annuities* from insurance companies. For a fixed fee, annuities offer a stream of income that lasts as long as the recipient lives. Uncertainty about medical expenses should be largely eliminated by Medicare, the government's health insurance plan for the elderly, and by private insurance plans.

The second explanation for the failure of the elderly to dissave is that they may want to leave bequests to their children. Economists have proposed various theories of the parent-child relationship and the bequest motive. In Chapter 15 we discussed some of these theories and their implications for consumption and fiscal policy.

Overall, research on the elderly suggests that the simplest life-cycle model cannot fully explain consumer behavior. There is no doubt that providing for retirement is an important motive for saving, but other motives, such as precautionary saving and bequests, appear important as well.³

16-4 Milton Friedman and the Permanent-Income Hypothesis

In a book published in 1957, Milton Friedman proposed the **permanent-income hypothesis** to explain consumer behavior. Friedman's permanent-income hypothesis complements Modigliani's life-cycle hypothesis: both use

³ To read more about the consumption and saving of the elderly, see Albert Ando and Arthur Kennickell, "How Much (or Little) Life Cycle Saving Is There in Micro Data?" in Rudiger Dornbusch, Stanley Fischer, and John Bossons, eds., *Macroeconomics and Finance: Essays in Honor of Franco Modigliani* (Cambridge, Mass.: MIT Press, 1986); and Michael Hurd, "Research on the Elderly: Economic Status, Retirement, and Consumption and Saving," *Journal of Economic Literature* 28 (June 1990): 565-589.

Irving Fisher's theory of the consumer to argue that consumption should not depend on current income alone. But unlike the life-cycle hypothesis, which emphasizes that income follows a regular pattern over a person's lifetime, the permanent-income hypothesis emphasizes that people experience random and temporary changes in their incomes from year to year.⁴

The Hypothesis

Friedman suggested that we view current income Y as the sum of two components, **permanent income** Y^P and **transitory income** Y^T . That is,

$$Y = Y^P + Y^T.$$

Permanent income is the part of income that people expect to persist into the future. Transitory income is the part of income that people do not expect to persist. Put differently, permanent income is average income, and transitory income is the random deviation from that average.

To see how we might separate income into these two parts, consider these examples:

- ▶ Maria, who has a law degree, earned more this year than John, who is a high-school dropout. Maria's higher income resulted from higher permanent income, because her education will continue to provide her a higher salary.
- ▶ Sue, a Florida orange grower, earned less than usual this year because a freeze destroyed her crop. Bill, a California orange grower, earned more than usual because the freeze in Florida drove up the price of oranges. Bill's higher income resulted from higher transitory income, because he is no more likely than Sue to have good weather next year.

These examples show that different forms of income have different degrees of persistence. A good education provides a permanently higher income, whereas good weather provides only transitorily higher income. Although one can imagine intermediate cases, it is useful to keep things simple by supposing that there are only two kinds of income: permanent and transitory.

Friedman reasoned that consumption should depend primarily on permanent income, because consumers use saving and borrowing to smooth consumption in response to transitory changes in income. For example, if a person received a permanent raise of \$10,000 per year, his consumption would rise by about as much. Yet if a person won \$10,000 in a lottery, he would not consume it all in one year. Instead, he would spread the extra consumption over the rest of his life. Assuming an interest rate of zero and a remaining life span of 50 years,

⁴ Milton Friedman, *A Theory of the Consumption Function* (Princeton, N.J.: Princeton University Press, 1957).

consumption would rise by only \$200 per year in response to the \$10,000 prize. Thus, consumers spend their permanent income, but they save rather than spend most of their transitory income.

Friedman concluded that we should view the consumption function as approximately

$$C = \alpha Y^P,$$

where α is a constant that measures the fraction of permanent income consumed. The permanent-income hypothesis, as expressed by this equation, states that consumption is proportional to permanent income.

Implications

The permanent-income hypothesis solves the consumption puzzle by suggesting that the standard Keynesian consumption function uses the wrong variable. According to the permanent-income hypothesis, consumption depends on permanent income; yet many studies of the consumption function try to relate consumption to current income. Friedman argued that this *errors-in-variables problem* explains the seemingly contradictory findings.

Let's see what Friedman's hypothesis implies for the average propensity to consume. Divide both sides of his consumption function by Y to obtain

$$APC = C/Y = \alpha Y^P/Y.$$

According to the permanent-income hypothesis, the average propensity to consume depends on the ratio of permanent income to current income. When current income temporarily rises above permanent income, the average propensity to consume temporarily falls; when current income temporarily falls below permanent income, the average propensity to consume temporarily rises.

Now consider the studies of household data. Friedman reasoned that these data reflect a combination of permanent and transitory income. Households with high permanent income have proportionately higher consumption. If all variation in current income came from the permanent component, the average propensity to consume would be the same in all households. But some of the variation in income comes from the transitory component, and households with high transitory income do not have higher consumption. Therefore, researchers find that high-income households have, on average, lower average propensities to consume.

Similarly, consider the studies of time-series data. Friedman reasoned that year-to-year fluctuations in income are dominated by transitory income. Therefore, years of high income should be years of low average propensities to consume. But over long periods of time—say, from decade to decade—the variation in income comes from the permanent component. Hence, in long time-series, one should observe a constant average propensity to consume, as in fact Kuznets found.

CASE STUDY

The 1964 Tax Cut and the 1968 Tax Surcharge

The permanent-income hypothesis can help us to interpret how the economy responds to changes in fiscal policy. According to the *IS–LM* model of Chapters 10 and 11, tax cuts stimulate consumption and raise aggregate demand, and tax increases depress consumption and reduce aggregate demand. The permanent-income hypothesis, however, predicts that consumption responds only to changes in permanent income. Therefore, transitory changes in taxes will have only a negligible effect on consumption and aggregate demand. If a change in taxes is to have a large effect on aggregate demand, it must be permanent.

Two changes in fiscal policy—the tax cut of 1964 and the tax surcharge of 1968—illustrate this principle. The tax cut of 1964 was popular. It was announced to be a major and permanent reduction in tax rates. As we discussed in Chapter 10, this policy change had the intended effect of stimulating the economy.

The tax surcharge of 1968 arose in a very different political climate. It became law because the economic advisers of President Lyndon Johnson believed that the increase in government spending from the Vietnam War had excessively stimulated aggregate demand. To offset this effect, they recommended a tax increase. But Johnson, aware that the war was already unpopular, feared the political repercussions of higher taxes. He finally agreed to a temporary tax surcharge—in essence, a one-year increase in taxes. The tax surcharge did not have the desired effect of reducing aggregate demand. Unemployment continued to fall, and inflation continued to rise.

The lesson to be learned from these episodes is that a full analysis of tax policy must go beyond the simple Keynesian consumption function; it must take into account the distinction between permanent and transitory income. If consumers expect a tax change to be temporary, it will have a smaller impact on consumption and aggregate demand.

16-5 Robert Hall and the Random-Walk Hypothesis

The permanent-income hypothesis is based on Fisher's model of intertemporal choice. It builds on the idea that forward-looking consumers base their consumption decisions not only on their current income but also on the income they expect to receive in the future. Thus, the permanent-income hypothesis highlights that consumption depends on people's expectations.

Recent research on consumption has combined this view of the consumer with the assumption of rational expectations. The rational-expectations assumption states that people use all available information to make optimal forecasts

about the future. As we saw in Chapter 13, this assumption can have profound implications for the costs of stopping inflation. It can also have profound implications for the study of consumer behavior.

The Hypothesis

The economist Robert Hall was the first to derive the implications of rational expectations for consumption. He showed that if the permanent-income hypothesis is correct, and if consumers have rational expectations, then changes in consumption over time should be unpredictable. When changes in a variable are unpredictable, the variable is said to follow a **random walk**. According to Hall, the combination of the permanent-income hypothesis and rational expectations implies that consumption follows a random walk.

Hall reasoned as follows. According to the permanent-income hypothesis, consumers face fluctuating income and try their best to smooth their consumption over time. At any moment, consumers choose consumption based on their current expectations of their lifetime incomes. Over time, they change their consumption because they receive news that causes them to revise their expectations. For example, a person getting an unexpected promotion increases consumption, whereas a person getting an unexpected demotion decreases consumption. In other words, changes in consumption reflect “surprises” about lifetime income. If consumers are optimally using all available information, then they should be surprised only by events that were entirely unpredictable. Therefore, changes in their consumption should be unpredictable as well.⁵

Implications

The rational-expectations approach to consumption has implications not only for forecasting but also for the analysis of economic policies. *If consumers obey the permanent-income hypothesis and have rational expectations, then only unexpected policy changes influence consumption. These policy changes take effect when they change expectations.* For example, suppose that today Congress passes a tax increase to be effective next year. In this case, consumers receive the news about their lifetime incomes when Congress passes the law (or even earlier if the law’s passage was predictable). The arrival of this news causes consumers to revise their expectations and reduce their consumption. The following year, when the tax hike goes into effect, consumption is unchanged because no news has arrived.

Hence, if consumers have rational expectations, policymakers influence the economy not only through their actions but also through the public’s expectation of their actions. Expectations, however, cannot be observed directly. Therefore, it is often hard to know how and when changes in fiscal policy alter aggregate demand.

⁵ Robert E. Hall, “Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence,” *Journal of Political Economy* 86 (April 1978): 971–987.

CASE STUDY

Do Predictable Changes in Income Lead to Predictable Changes in Consumption?

Of the many facts about consumer behavior, one is impossible to dispute: income and consumption fluctuate together over the business cycle. When the economy goes into a recession, both income and consumption fall, and when the economy booms, both income and consumption rise rapidly.

By itself, this fact doesn't say much about the rational-expectations version of the permanent-income hypothesis. Most short-run fluctuations are unpredictable. Thus, when the economy goes into a recession, the typical consumer is receiving bad news about his lifetime income, so consumption naturally falls. And when the economy booms, the typical consumer is receiving good news, so consumption rises. This behavior does not necessarily violate the random-walk theory that changes in consumption are impossible to forecast.

Yet suppose we could identify some *predictable* changes in income. According to the random-walk theory, these changes in income should not cause consumers to revise their spending plans. If consumers expected income to rise or fall, they should have adjusted their consumption already in response to that information. Thus, predictable changes in income should not lead to predictable changes in consumption.

Data on consumption and income, however, appear not to satisfy this implication of the random-walk theory. When income is expected to fall by \$1, consumption will on average fall at the same time by about \$0.50. In other words, predictable changes in income lead to predictable changes in consumption that are roughly half as large.

Why is this so? One possible explanation of this behavior is that some consumers may fail to have rational expectations. Instead, they may base their expectations of future income excessively on current income. Thus, when income rises or falls (even predictably), they act as if they received news about their lifetime resources and change their consumption accordingly. Another possible explanation is that some consumers are borrowing-constrained and, therefore, base their consumption on current income alone. Regardless of which explanation is correct, Keynes's original consumption function starts to look more attractive. That is, current income has a larger role in determining consumer spending than the random-walk hypothesis suggests.⁶

⁶ John Y. Campbell and N. Gregory Mankiw, "Consumption, Income, and Interest Rates: Reinterpreting the Time-Series Evidence," *NBER Macroeconomics Annual* (1989): 185–216; Jonathan Parker, "The Response of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review* 89 (September 1999): 959–973; and Nicholas S. Souleles, "The Response of Household Consumption to Income Tax Refunds," *American Economic Review* 89 (September 1999): 947–958.

16.6 David Laibson and the Pull of Instant Gratification

Keynes called the consumption function a “fundamental psychological law.” Yet, as we have seen, psychology has played little role in the subsequent study of consumption. Most economists assume that consumers are rational maximizers of utility who are always evaluating their opportunities and plans in order to obtain the highest lifetime satisfaction. This model of human behavior was the basis for all the work on consumption theory from Irving Fisher to Robert Hall.

More recently, economists have started to return to psychology. They have suggested that consumption decisions are not made by the ultrarational *homo economicus* but by real human beings whose behavior can be far from rational. The most prominent economist infusing psychology into the study of consumption is Harvard professor David Laibson.

Laibson notes that many consumers judge themselves to be imperfect decisionmakers. In one survey of the American public, 76 percent said they were not saving enough for retirement. In another survey of the baby-boom generation, respondents were asked the percentage of income that they save and the percentage that they thought they should save. The saving shortfall averaged 11 percentage points.

According to Laibson, the insufficiency of saving is related to another phenomenon: the pull of instant gratification. Consider the following two questions:

Question 1: Would you prefer (A) a candy today or (B) two candies tomorrow.

Question 2: Would you prefer (A) a candy in 100 days or (B) two candies in 101 days.

Many people confronted with such choices will answer A to the first question and B to the second. In a sense, they are more patient in the long run than they are in the short run.

This raises the possibility that consumers’ preferences may be *time-inconsistent*: they may alter their decisions simply because time passes. A person confronting question 2 may choose B and wait the extra day for the extra candy. But after 100 days pass, he then confronts question 1. The pull of instant gratification may induce him to change his mind.

We see this kind of behavior in many situations in life. A person on a diet may have a second helping at dinner, while promising himself that he will eat less tomorrow. A person may smoke one more cigarette, while promising himself that this is the last one. And a consumer may splurge at the shopping center, while promising himself that tomorrow he will cut back his spending and start saving more for retirement. But when tomorrow arrives, the promises are in the past, and a new self takes control of the decisionmaking, with its own desire for instant gratification.

These observations raise as many questions as they answer. Will the renewed focus on psychology among economists offer a better understanding of

consumer behavior? Will it offer new prescriptions regarding, for instance, tax policy toward saving? It is too early to say, but without doubt, these questions are on the forefront of the research agenda.⁷

16.7 Conclusion

In the work of six prominent economists, we have seen a progression of views on consumer behavior. Keynes proposed that consumption depends largely on current income. Since then, economists have argued that consumers understand that they face an intertemporal decision. Consumers look ahead to their future resources and needs, implying a more complex consumption function than the one that Keynes proposed. Keynes suggested a consumption function of the form

$$\text{Consumption} = f(\text{Current Income}).$$

Recent work suggests instead that

$$\begin{aligned} \text{Consumption} \\ = f(\text{Current Income, Wealth, Expected Future Income, Interest Rates}). \end{aligned}$$

In other words, current income is only one determinant of aggregate consumption.

Economists continue to debate the importance of these determinants of consumption. There remains disagreement about, for example, the influence of interest rates on consumer spending, the prevalence of borrowing constraints, and the importance of psychological effects. Economists sometimes disagree about economic policy because they assume different consumption functions. For instance, as we saw in the previous chapter, the debate over the effects of government debt is in part a debate over the determinants of consumer spending. The key role of consumption in policy evaluation is sure to maintain economists' interest in studying consumer behavior for many years to come.

Summary

1. Keynes conjectured that the marginal propensity to consume is between zero and one, that the average propensity to consume falls as income rises, and that current income is the primary determinant of consumption. Studies of household data and short time-series confirmed Keynes's conjectures. Yet studies of long time-series found no tendency for the average propensity to consume to fall as income rises over time.
2. Recent work on consumption builds on Irving Fisher's model of the consumer. In this model, the consumer faces an intertemporal budget constraint

⁷ For more on this topic, see David Laibson, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics* 62 (May 1997): 443–477; and George-Marios Angeletos, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, "The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evidence," *Journal of Economic Perspectives*, 15(3) (Summer 2001): 47–68.

and chooses consumption for the present and the future to achieve the highest level of lifetime satisfaction. As long as the consumer can save and borrow, consumption depends on the consumer's lifetime resources.

3. Modigliani's life-cycle hypothesis emphasizes that income varies somewhat predictably over a person's life and that consumers use saving and borrowing to smooth their consumption over their lifetimes. According to this hypothesis, consumption depends on both income and wealth.
4. Friedman's permanent-income hypothesis emphasizes that individuals experience both permanent and transitory fluctuations in their income. Because consumers can save and borrow, and because they want to smooth their consumption, consumption does not respond much to transitory income. Consumption depends primarily on permanent income.
5. Hall's random-walk hypothesis combines the permanent-income hypothesis with the assumption that consumers have rational expectations about future income. It implies that changes in consumption are unpredictable, because consumers change their consumption only when they receive news about their lifetime resources.
6. Laibson has suggested that psychological effects are important for understanding consumer behavior. In particular, because people have a strong desire for instant gratification, they may exhibit time-inconsistent behavior and may end up saving less than they would like.

KEY CONCEPTS

Marginal propensity to consume	Normal good	Precautionary saving
Average propensity to consume	Income effect	Permanent-income hypothesis
Intertemporal budget constraint	Substitution effect	Permanent income
Discounting	Borrowing constraint	Transitory income
Indifference curves	Life-cycle hypothesis	Random walk
Marginal rate of substitution		

QUESTIONS FOR REVIEW

1. What were Keynes's three conjectures about the consumption function?
2. Describe the evidence that was consistent with Keynes's conjectures and the evidence that was inconsistent with them.
3. How do the life-cycle and permanent-income hypotheses resolve the seemingly contradictory pieces of evidence regarding consumption behavior?
4. Use Fisher's model of consumption to analyze an increase in second-period income. Compare the case in which the consumer faces a binding borrowing constraint and the case in which he does not.
5. Explain why changes in consumption are unpredictable if consumers obey the permanent-income hypothesis and have rational expectations.
6. Give an example in which someone might exhibit time-inconsistent preferences.

PROBLEMS AND APPLICATIONS

- The chapter uses the Fisher model to discuss a change in the interest rate for a consumer who saves some of his first-period income. Suppose, instead, that the consumer is a borrower. How does that alter the analysis? Discuss the income and substitution effects on consumption in both periods.
- Jack and Jill both obey the two-period Fisher model of consumption. Jack earns \$100 in the first period and \$100 in the second period. Jill earns nothing in the first period and \$210 in the second period. Both of them can borrow or lend at the interest rate r .
 - You observe both Jack and Jill consuming \$100 in the first period and \$100 in the second period. What is the interest rate r ?
 - Suppose the interest rate increases. What will happen to Jack's consumption in the first period? Is Jack better off or worse off than before the interest rate rise?
 - What will happen to Jill's consumption in the first period when the interest rate increases? Is Jill better off or worse off than before the interest-rate increase?
- The chapter analyzes Fisher's model for the case in which the consumer can save or borrow at an interest rate of r and for the case in which the consumer can save at this rate but cannot borrow at all. Consider now the intermediate case in which the consumer can save at rate r_s and borrow at rate r_b , where $r_s < r_b$.
 - What is the consumer's budget constraint in the case in which he consumes less than his income in period one?
 - What is the consumer's budget constraint in the case in which he consumes more than his income in period one?
 - Graph the two budget constraints and shade the area that represents the combination of first-period and second-period consumption the consumer can choose.
 - Now add to your graph the consumer's indifference curves. Show three possible outcomes: one in which the consumer saves, one in which he borrows, and one in which he neither saves nor borrows.
- What determines first-period consumption in each of the three cases?
- Explain whether borrowing constraints increase or decrease the potency of fiscal policy to influence aggregate demand in each of the following two cases:
 - A temporary tax cut.
 - An announced future tax cut.
- In the discussion of the life-cycle hypothesis in the text, income is assumed to be constant during the period before retirement. For most people, however, income grows over their lifetimes. How does this growth in income influence the lifetime pattern of consumption and wealth accumulation shown in Figure 16-12 under the following conditions?
 - Consumers can borrow, so their wealth can be negative.
 - Consumers face borrowing constraints that prevent their wealth from falling below zero.

Do you consider case (a) or case (b) to be more realistic? Why?
- Demographers predict that the fraction of the population that is elderly will increase over the next 20 years. What does the life-cycle model predict for the influence of this demographic change on the national saving rate?
- One study found that the elderly who do not have children dissave at about the same rate as the elderly who do have children. What might this finding imply about the reason the elderly do not dissave as much as the life-cycle model predicts?
- Consider two savings accounts that pay the same interest rate. One account lets you take your money out on demand. The second requires that you give 30-day advance notification before withdrawals. Which account would you prefer? Why? Can you imagine a person who might make the opposite choice? What do these choices say about the theory of the consumption function?