W''_M on a lower indifference curve (U_1^M) In general, households that generate the average amount of trash (2,000 pounds) or less will benefit from changing to a price-per-bag arrangement, while households that generate more than the average amount are likely to be harmed. Under the annual fee system, households that generate little trash tend to pay more than the cost of disposing of their trash, and that excess implicitly subsidizes heavy users of the trash disposal service. The price-per-bag system makes each household pay the cost of its own trash disposal, thereby removing the implicit subsidy arrangement of the annual fee and harming those who were on balance being subsidized.

5.5 The Consumer's Choice to Save or Borrow

Saving involves consuming less than one's current income, which makes it possible to consume more at a later date. Borrowing makes it possible to consume more than current income, but consumption in the future must fall below future income to repay the loan. A decision to save (or borrow) is therefore a decision to rearrange consumption between various time periods. By suitably adapting the theory of consumer choice, we can examine the factors that influence decisions to save or borrow.

Let's confront this topic in the simplest possible way. Imagine a short-lived individual, Ms. Cher Noble, whose lifetime spans two time periods, year 1 (this year) and year 2 (next year). Ms. Noble's earnings in year 1 (I_1) are $10,000, but they will fall to $2,200 in year 2 ($I_2$). The interest rate, r, at which she can borrow or lend is 10 percent per year. We assume that there is no inflation in the general price level so that $1 will purchase the same quantity of goods both years. (If there is inflation, the earnings in each year can simply be expressed in dollars of constant purchasing power, and the real rate of interest can be used instead of the nominal rate.)

This information allows us to plot Ms. Noble's budget line. In Figure 5.9, consumption in year 1 is measured on the vertical axis and consumption in year 2 (next year's consumption) on the horizontal axis. Any point in the diagram therefore represents a certain level of

Figure 5.9

Consumer Choice over Two Time Periods

With an interest rate, r, of 10 percent and with earnings of $10,000 in year 1 and $2,200 in year 2, the budget line relating consumption in the two years is AZ with a slope of $1/(1 + r)$. The optimal point is W, with saving of I_1C_1 in year 1. In year 2, consumption exceeds that year's income by I_2C_2, which is equal to the amount saved ($3,000) plus interest on this sum ($300).

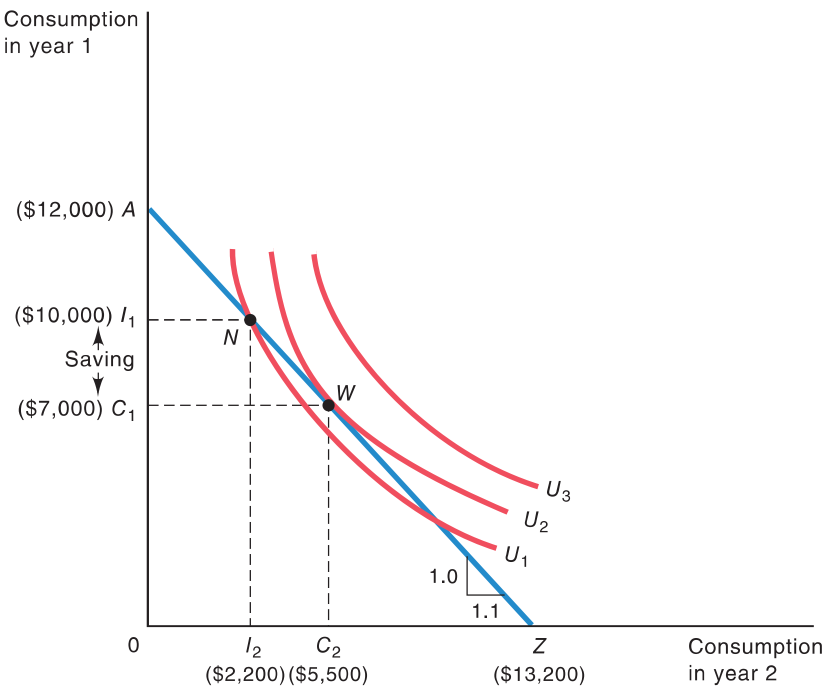

consumption in each year. The budget line indicates what combinations are available to the consumer. Point *N*, for example, identifies the consumption mix where the individual's entire earnings are spent in each year: a market basket containing $10,000 in consumption in year 1 (equal to year 1 earnings of $10,000) and $2,200 in consumption in year 2 (equal to year 2 earnings). Point *N* is sometimes called the **endowment point**, showing the consumption mix available to the individual if no saving or borrowing takes place. Alternatively, by saving or borrowing, the consumer can choose a different market basket.

endowment point
the consumption mix available to the individual if no saving or borrowing takes place

To identify another point on the budget line, suppose that Ms. Noble's entire year 1 income of $10,000 is saved. In this case consumption in year 1 is zero, but in year 2 she could consume $13,200, equal to the sum saved the year before ($10,000), plus interest on that sum at a 10 percent rate ($1,000 in interest), plus earnings in year 2 ($2,200). Thus, point *Z* shows the horizontal intercept of the budget line; if consumption in year 1 is zero, Ms. Noble can consume $13,200 in year 2. The vertical intercept, *A*, shows the maximum possible consumption in year 1. This maximum is achieved by borrowing as much as possible, limited by how much can be repaid in year 2. That is, if $2,000 is borrowed, year 1 consumption can be $12,000 ($2,000 plus year 1 earnings). Ms. Noble can borrow $2,000 at a maximum because $2,000 plus 10 percent interest, or $2,200, must be repaid the next year. This amount equals total year 2 earnings with nothing left over for consumption. So the budget line's vertical intercept is $12,000.

Points *A*, *N*, and *Z* represent three points on the consumer's budget line. Connecting these points yields *AZ* as the entire budget line. If the consumer, Ms. Noble, chooses a point along the *NZ* portion of the line, she will be consuming less than earnings in year 1, or saving, and consuming more than earnings in year 2 (by an amount equal to the previous year's saving plus interest). Along the *AN* portion of the budget line, Ms. Noble is borrowing in year 1 and repaying the loan in year 2.

Notice how the budget line's slope relates to the interest rate. In fact, the slope is equal to $1/(1 + r)$, where r is the interest rate. Thus, if Ms. Noble reduces consumption by $1.00 in year 1 (saves $1.00), she can increase consumption in year 2 by $1.00 plus the interest of $0.10, or by $1.10. With an interest rate of 10 percent the slope is equal to $1/(1 + 0.10)$, or 0.91 (rounded). This result tells us that the present cost of $1.00 consumed in year 2 is $0.91 in year 1, since $0.91 saved today grows to $1.00 a year later at a 10 percent interest rate—or, conversely, to have $1.00 to spend in year 2, the consumer must save $0.91 in year 1.

Now let's bring Ms. Noble's preferences into the picture. Because consumption in both years is desirable (more is preferred to less), the indifference curves have the usual shape. The slope of an indifference curve at any point is the marginal rate of substitution between consumption in year 1 and consumption in year 2, and it shows the willingness of Ms. Noble to reduce consumption in year 1 to have greater consumption in year 2.

For the indifference curves shown in Figure 5.9, the consumer's optimal choice is point *W*. Consumption in year 1 is $7,000, and consumption in year 2 is $5,500. Note that Ms. Noble is saving some of her year 1 earnings, as indicated by the choice to consume less than her income of $10,000 in year 1. The amount of saving in year 1 is the difference between income and consumption in that year, which is shown by the distance I_1C_1, or $3,000. In year 2, the individual's consumption is $3,300 greater than year 2 earnings; this sum is equal to the amount saved, $3,000, plus interest on the saving.

Thus, the consumer's optimal consumption point is once again characterized by a tangency between an indifference curve and the budget line. The only novel feature here is that the "commodities" consumed refer to consumption in different time periods, but that does not change the substance of the analysis.

A Change in Endowment

The budget line relevant for consumption choices over time depends on current and future income as well as the interest rate. A change in one or more of these variables will alter the budget line and affect the market basket chosen. Let's examine how a change in year 2 earnings will affect consumption and saving. Continuing with the example just discussed, suppose that Ms. Noble expects year 2 earnings to be zero rather than $2,200.

A change in earnings in either year moves the endowment point in the graph. In Figure 5.10, budget line *AZ* is reproduced from Figure 5.9; it shows the opportunities available when earnings are $10,000 in year 1 and $2,200 in year 2. When year 2 earnings fall to zero, the endowment point moves from *N* to I_1 on the vertical axis: if Ms. Noble's consumption equals earnings in each year, consumption will be $10,000 in year 1 and nothing in year 2. The new budget line is I_1Z', and it is parallel to *AZ* because the interest rate remains unchanged [both lines have a slope of $1/(1 + r)$].

A reduction in future income will not alter the relative cost of future and present consumption, but it will influence behavior through its income effect. If consumption in each year is a normal good—which is almost certain to be true because consumption is a very broadly defined good—then the shift in the budget line will lead to reduced consumption in both years. Ms. Noble spreads the loss over both years, cutting back on consumption in year 1 and year 2. The new optimal consumption point, *W′*, illustrates this situation with consumption reduced from $7,000 to $6,000 in year 1 and from $5,500 to $4,400 in year 2.

This analysis implies that a reduction in expected future earnings causes saving in year 1 to increase. Recall that saving is the difference between consumption and income. Before the income loss, saving was I_1C_1, or $3,000; after the loss, saving increases to I_1C_1', or $4,000. Current saving, therefore, doesn't depend exclusively on current income (which is unchanged in this example); it is also affected by the expected level of future income.

So far we have been looking at a person who saves in year 1. Yet some people borrow in the present and repay the loan later. Under different circumstances the individual shown to be a saver in Figure 5.10 could become a borrower. For example, suppose that instead of year 1 income of $10,000 and year 2 income of $2,200 (point *N*), earnings are $2,000 in year 1 and $11,000 in year 2 (point *Y*). *The budget line doesn't change*: only the endowment point changes, from *N* to *Y* on *AZ*. The optimal consumption point remains *W*, but to reach that point, the individual borrows $5,000 in year 1 and repays the loan plus interest in year 2.

Figure 5.10

An Income Change and Intertemporal Choice
If year 2 income is zero instead of $2,200, the budget line shifts from *AZ* to I_1Z', a parallel shift. The result is an increase in saving in year 1, from I_1C_1 to I_1C_1'. Consumption in both years falls, assuming that consumption in each year is a normal good.

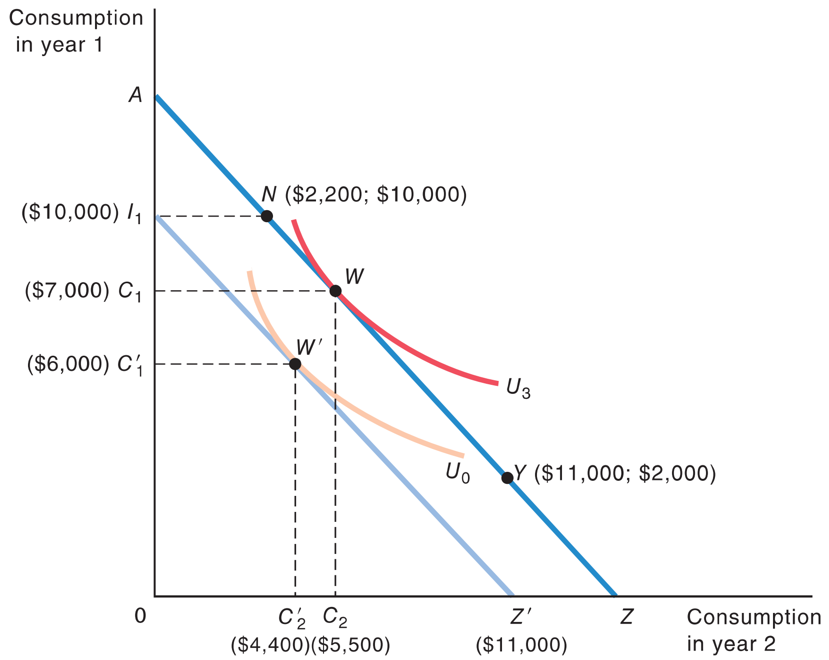

This analysis shows how the pattern of earnings over time is likely to affect saving and borrowing decisions. A relatively high present income but sharply reduced future income (such as endowment point *N*) is typical of middle-aged persons approaching retirement, and we expect to see them save part of their current income. A low present income but a higher expected future income (such as endowment point *Y*) is typical of students and young workers, and we often see such persons acquiring debt and consuming above their present income.

APPLICATION 5.4

Social Security and Saving

Our intertemporal choice model can easily be adapted to examine the important issue of how the Social Security program in the United States affects American workers saving for retirement. To do so, we make our two time periods represent the years before and after retirement. Thus, we get the lifetime budget line *AZ* for Justin, a typical American worker, in Figure 5.11, which shows his options regarding consumption before and after retirement. His lifetime earnings are *OA*, and these must finance his consumption during his working years as well as his retired years. With preferences shown by U_1, he chooses to consume C_1 before retirement, saving AC_1 of his lifetime earnings, and this provides a retirement level of consumption of C_2.

Now introduce Social Security, a program which we will discuss in greater detail in a later chapter, but here we will simplify to two key parts: a tax on earnings and a pension provided by the government during retirement. The tax of *AA′* by itself shifts the budget line inward to *A′Z′*, but that is not the whole story because in return for paying the tax Justin is promised a pension when he retires. Assume he expects the pension to be C_G, which is the amount he would have gotten if he saved the amount of the tax for himself. Then the expected government pension shifts the after-tax budget line outward to *A′RZ*, with *A′R* equal to the government provided pension, C_G. Confronted with *A′RZ*, Justin chooses point *W*, the same point as before the introduction of Social Security.

It may look like Social Security has no effect, but a closer examination reveals an important result. Justin is now saving $A'C_1$ to supplement the government pension, and since he was previously saving AC_1, we see that Social Security has caused him to reduce his saving by *AA′*. In fact, he has reduced his saving by exactly the amount of the Social Security tax he pays. This need not always be the outcome; it is the result of our assumption that he expects the government pension to be the same as he would have achieved by saving the amount of the tax. (You will find it instructive to go through the analysis for the cases when he expects the government pension to be larger or smaller than shown in our diagram.)

Figure 5.11

Social Security and Saving

Social Security is composed of two key parts: a tax on earnings (*AA′* in this diagram); and a pension provided by the government during retirement (C_G). Social Security reduces saving from $S(AC_1)$ to $S'(A'C_1)$.

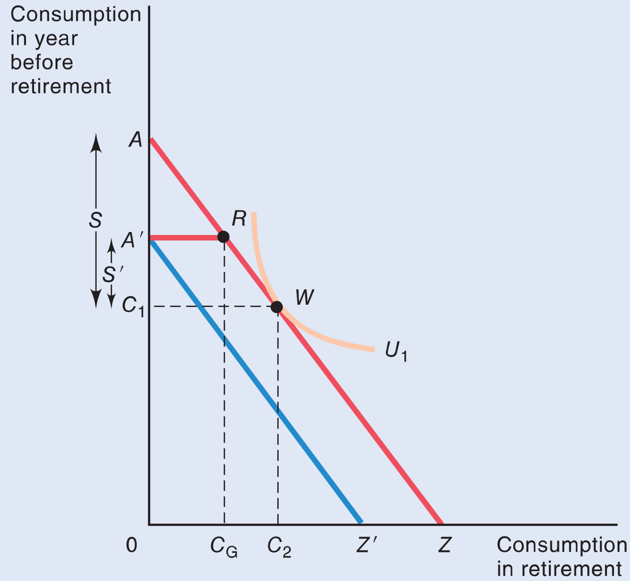

In interpreting this analysis, it is important to understand that the government is not taking Justin's tax dollars and saving them for him (which would leave total saving unchanged). Social Security is financed on a pay-as-you-go basis, which means that the government takes Justin's tax dollars and transfers them to people already retired. So when Justin reduces his saving, there is no offsetting increase in government saving. This analysis thus indicates that pay-as-you-go Social Security decreases national saving by reducing the incentive workers have to save for their retirement. In a later chapter we will examine how this reduction in saving affects the overall economy.

Changes in the Interest Rate

Will people save more at a higher interest rate? The answer to this question may not be the yes response commonly expected. Let's see why.

A higher interest rate changes the relative cost of present versus future consumption, which is reflected in a change in the budget line's slope. If the interest rate rises from 10 to 20 percent, the budget line's slope, $1/(1 + r)$, changes from $1/(1 + 0.1)$ to $1/(1 + 0.2)$, so the new budget line has a slope of 1/1.2, or 0.83. Reducing consumption by \$1.00 in year 1 (saving \$1.00) permits consumption of \$1.20 more in year 2. Put somewhat differently, the present cost of consuming \$1.00 in year 2 is \$0.83 in year 1, because \$0.83 will grow to \$1.00 in one year at a 20 percent interest rate. Or, conversely, to have \$1.00 to spend in year 2, the consumer need save only \$0.83 now at the higher interest rate. Thus, a higher interest rate reduces the cost of future consumption in terms of the present sacrifice required.

Figure 5.12a shows the way the budget line changes. The initial budget line *AZ* once again reflects the 10 percent interest rate, and the endowment point is *N*. *When the interest rate rises to 20 percent, the budget line rotates about point* N *and becomes* A′Z′. Point *N* is also on the new budget line because the individual can still consume I_1 in year 1 and I_2 in year 2 if no borrowing or saving takes place. The budget line's slope becomes flatter through point *N* because an increase in the interest rate from 10 to 20 percent increases the cost of present versus future consumption. At the higher interest rate, that is, increasing present consumption by \$1.00 (through borrowing) now requires \$1.20 to be paid back in year 2. And increasing future consumption by \$1.00 (through saving) now requires only \$0.83 to be banked in year 1.

Incorporating the indifference curves of the individual, we can determine the preferred market basket associated with the new budget line. For the indifference curves shown in Figure 5.12a, the initial optimal levels of consumption are C_1 and C_2, with I_1C_1 saving in year 1. When the interest rate rises to 20 percent, the new optimal point W' involves consumption of C_2' in year 2 and C_1' in year 1. A lower consumption in year 1 means an increase in saving from I_1C_1 to I_1C_1'. Consequently, this individual will save more when the interest rate rises from 10 to 20 percent.

The Case of a Higher Interest Rate Leading to Less Saving

The Figure 5.12a case, however, is not the only possible outcome. For example, if the individual's indifference curves are as shown in Figure 5.12b, an increase in the interest rate from 10 to 20 percent leads to less saving. With the new budget line $A'Z'$, the optimal point is W', involving more consumption in both years than at *W*. (Note that a higher interest rate allows a saver to increase consumption in the present and still consume more in the future.) Because consumption increases in year 1, saving falls when the interest rate rises from 10 to 20 percent for an individual with indifference curves as shown in Figure 5.12b.

What factors determine whether saving rises or falls when the interest rate increases from 10 to 20 percent? Once again, they are the familiar income and substitution effects. (To avoid complicating the diagram, we do not show these effects explicitly in Figure 5.12 but, instead, give a verbal explanation.) The substitution effect associated with a higher interest rate results from the change in the relative cost of present versus future consumption.

A Change in Interest Rates and Intertemporal Choice

A change in the interest rate rotates the budget line at endowment point *N*. When the interest rate rises from 10 to 20 percent, the budget line rotates from *AZ* to *A′Z′*. Whether saving rises or falls depends on the magnitude of the income effect (encourages less saving provided that year 1 consumption is a normal good) relative to the substitution effect (encourages more saving). (a) If the substitution effect outweighs the income effect, saving rises with the interest rate increase. (b) If the income effect outweighs the substitution effect, saving falls with the interest rate increase.

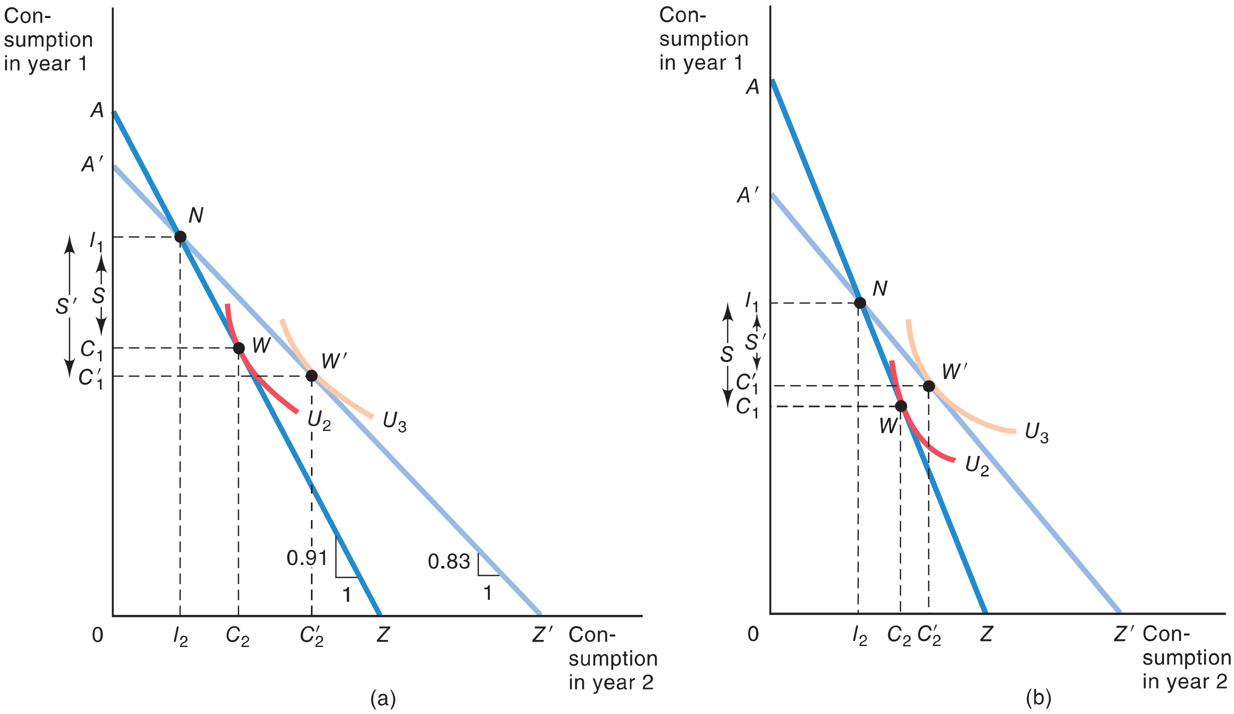

A higher interest rate reduces the cost of future consumption, which implies a substitution effect that favors future consumption at the expense of present consumption. So the substitution effect encourages future consumption instead of present consumption. In contrast, the income effect associated with a higher interest rate enriches the saver, who is able to attain a higher indifference curve. A higher real income enables the individual to consume more in both periods, so the income effect favors increased consumption in both periods if consumption in each period is a normal good.

Because both the substitution and income effects favor more consumption in year 2, it will definitely increase. However, substitution and income effects for year 1 consumption are in opposing directions, so the outcome depends on their relative sizes. If the income effect is greater, year 1 consumption will rise, which implies that saving will fall.

Intuitively, we can see why some people might save less at a higher interest rate. Think of a person who is saving for a specific good, like a cruise around the world. A higher interest rate means the consumer can purchase the cruise without committing as many dollars to present saving. If the cruise costs $6,000, a person would have to save $5,454 at a 10 percent interest rate but only $5,000 at a 20 percent interest rate to purchase the cruise one year later. For such a focused saver the income effect of a higher interest rate (favoring greater current consumption and thus less saving) may outweigh the substitution effect (favoring greater future consumption through more saving). Saving for such an individual thus may decrease as the interest rate rises.

5.6 Investor Choice

The theory of consumer choice not only helps to illuminate the decision to save or borrow but also can be applied to explain what types of financial assets an individual intent on saving for the future should purchase, or invest in. That is, should such an individual invest in stocks, U.S. Treasury bills, gold, cattle futures or the local bank? At first glance, investing in stocks would appear to be quite desirable given the returns generated historically by such financial assets. For example, a portfolio composed of U.S. common stocks has averaged an annual real rate of return of roughly 6 percent since 1802.[8] (The real rate of return is the nominal rate of return less the inflation rate.) By contrast, short-term U.S. Treasury bills have provided an average annual real rate of return of only 0.6 percent.

The historical data on annual rates of return also indicate that a dollar invested in U.S. common stocks in 1802 would have grown, in real terms, to nearly $800,000 today—even when accounting for the Great Depression and the most recent significant economic downturn. A dollar invested in short-term U.S. Treasury bills ("T-bills" for short) in 1802 would have grown to only about $300 after adjusting for inflation.

Given the historical data on the rates of return of various financial assets, a natural question arises. Specifically, why would a rational investor ever choose to purchase T-bills as opposed to stocks? It would seem logical to place all one's financial eggs in the asset basket offering the highest return.

The reason investors do not typically allocate all of their portfolios to the asset with the highest rate of return involves the risk associated with such an action. If higher rates of return are associated with greater risk and investors are averse to risk (that is, risk is a "bad"), a trade-off exists between return and risk. While it may be true that *in retrospect* stocks display higher returns than T-bills, a person contemplating purchasing stocks does not know *beforehand* what next year's return will be.

Stocks may average a 6 percent annual real rate of return. The average, however, may be arrived at by a return of –6 percent in one year and 18 percent in another, with the dramatic and unpredictable gyrations averaging out to 6 percent over a two-year period. By contrast, T-bills have a lower average return but may be a very sure thing from period to period. There is very little chance, that is, that the U.S. government will default on its borrowing obligations by not paying the promised amount to purchasers of T-bills that mature (must be paid off) in a fairly short time period such as three months. An average annual real rate of return on T-bills of 0.6 percent thus may be arrived at by a constant and more predictable return of 0.6 percent in year 1 and 0.6 percent in year 2. Indeed, the historical volatility of returns has been significantly lower for T-bills than for stocks.

Figure 5.13 displays what an investor's indifference map looks like if higher expected returns are associated with greater volatility or risk, and risk is a bad in the eyes of the investor. As discussed in Chapter 3, the indifference curves are upward sloping if the expected return on an asset is a good while the risk associated with the asset's return is a bad. Holding constant the expected return on an asset at $return_1$, a rise in the risk associated with the return on the asset, such as a move from point *A* to point *B*, places the risk-averse investor on a lower indifference curve (U_1 versus U_2). Furthermore, as the risk of an asset's return rises from $risk_1$ to $risk_2$, the expected return on the asset must rise from $return_1$ to $return_2$, as at point *C*, to compensate for the added risk and leave the investor's utility level unchanged at U_2. Greater levels of well-being are shown by indifference curves above and to the northwest: U_2 is preferred to U_1.

[8]Jeremy J. Siegel, *Stocks for the Long Run* 4th ed., (New York: McGraw-Hill, 2014).

The Return–Risk Trade-off
Since volatility or risk is a bad and expected return is a good, the investor's indifference curves are upward sloping.

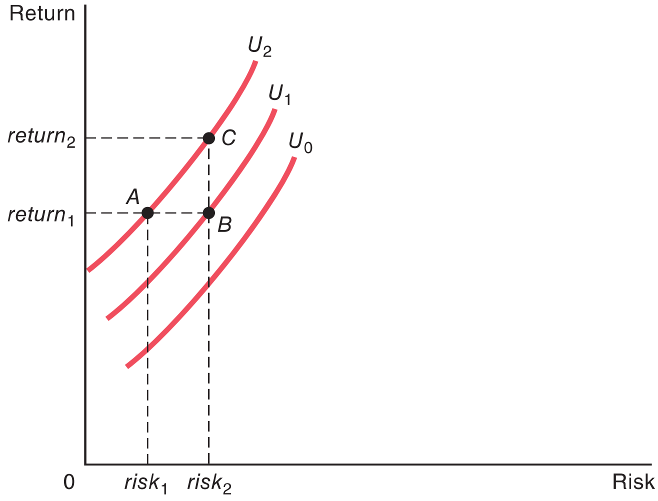

APPLICATION 5.5

Entrepreneurs and Their Risk–Return Preferences

Studies indicate that small, entrepreneurial ventures created more than 90 percent of the net new jobs in the United States in the past few decades.[9] Concurrently, interest in entrepreneurship has grown. More than 1,000 colleges and universities now offer courses in entrepreneurship, and a recent Gallup poll indicates that 7 out of 10 high school students in the United States want to start and own their own businesses in their adult years.

Entrepreneurs are characterized by their willingness to take risks and their ability to pursue and seize opportunities to create value, notwithstanding apparent resource constraints, through a new or existing company. Figure 5.14

[9]Mickie P. Slaughter, "Entrepreneurship: Economic Impact and Public Policy Implications," Kauffman Center for Entrepreneurial Leadership, March 1996.

Figure 5.14 Differences in Individuals' Risk–Return Preferences
An entrepreneur (a) has flatter indifference curves than a nonentrepreneur (b) because the entrepreneur is less risk averse.

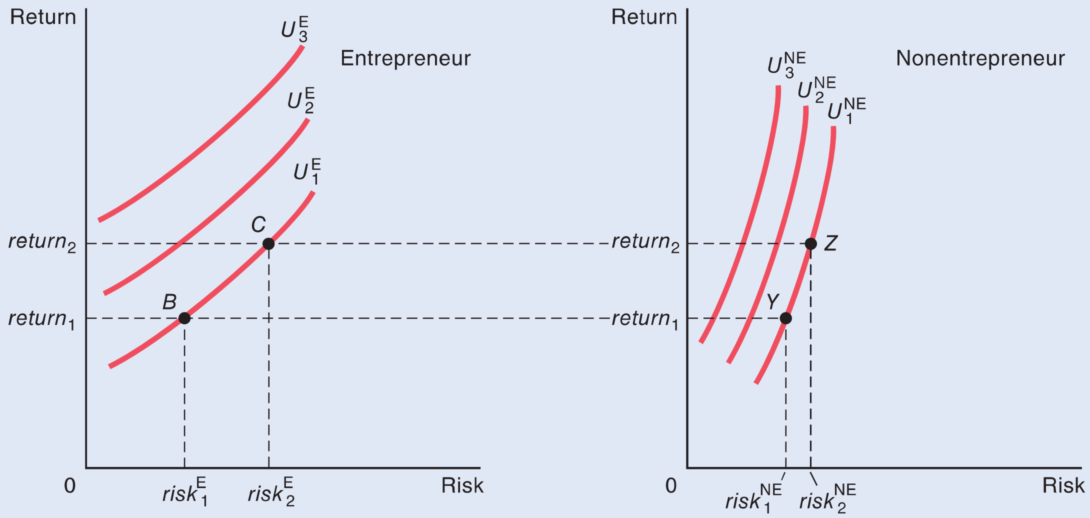

depicts the observed differences in risk-return preferences for a representative entrepreneur and nonentrepreneur. The indifference curves are flatter in the case of the entrepreneur because the entrepreneur is willing to take on more risk for a given increase in the expected return (for example, from $return_1$ to $return_2$). That is, the difference between $risk_1^E$ and $risk_2^E$ (as the entrepreneur moves between points *B* and *C* along indifference curve U_1^E in Figure 5.14a) is greater than the difference between $risk_1^{NE}$ and $risk_2^{NE}$ (as the nonentrepreneur moves between points *Y* and *Z* along indifference curve U_1^{NE} in Figure 5.14b).

Investor Preferences toward Risk: Risk Aversion

We have so far mentioned that, other things being equal, risk-averse investors prefer a sure thing—an asset promising the same return at lower volatility. As the volatility of an asset's return increases, risk-averse investors must be compensated with a greater expected return to remain equally well off (that is, to remain on the same indifference curve, as shown in Figure 5.13). In addition to looking at an indifference map, we can also examine what is meant by "risk aversion" by graphing the total utility curve of an investor as a function of the expected return per dollar invested in the stock of a company such as IBM.

As depicted in Figure 5.15a, the total utility curve has a positive but diminishing slope for risk-averse investors. To see why this is the case, let's investigate the properties of such a total utility curve. The height of the curve at each possible return indicates the utility of that return to the investor. For example, as drawn, the utility of a return of 2 percent equals 100 while the utility of a return of 8 percent equals 300.

Figure 5.15 Investor Preferences toward Risk

(a) A risk-averse investor's total utility curve has an upward, but diminishing slope. For every return between the equally likely returns of 2 and 8 percent, the height of the total utility curve exceeds the height of the expected utility chord, *AZ*. (b) A risk-neutral investor has a total utility curve that is an upward-sloping straight line. The total utility curve and the expected utility chord have the same height at every return between 2 and 8 percent. (c) A risk-loving investor's total utility curve has an upward and increasing slope. The height of the expected utility chord exceeds the height of the total utility curve over the relevant range of returns.

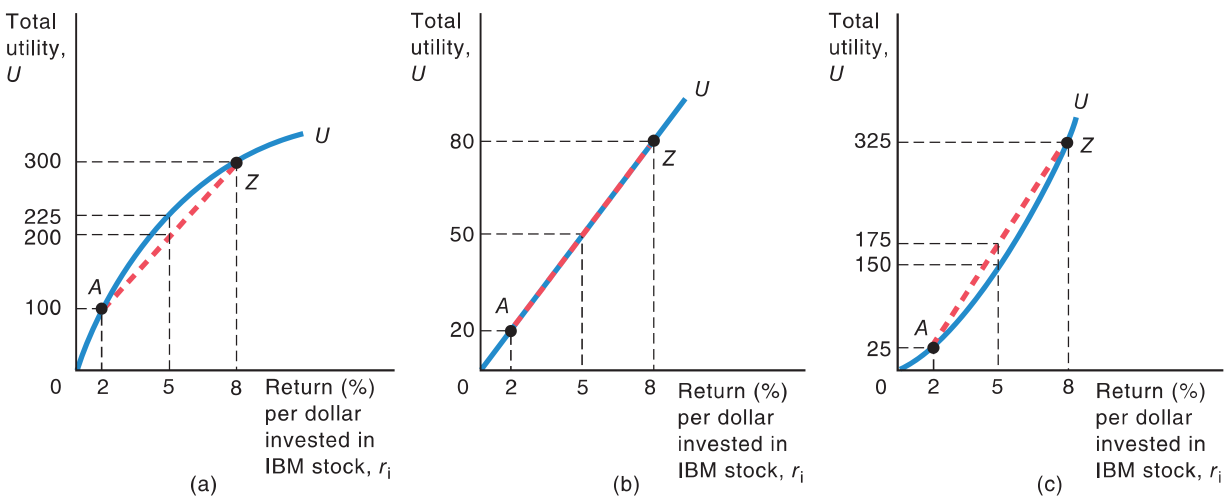

expected return
the summed value of each possible rate of return weighted by its probability

expected utility
the summed value of each possible utility weighted by its probability

Suppose that, per dollar invested, IBM stock will provide either a return of 2 or 8 percent and that the probability of either outcome is 0.5. The **expected return** on IBM stock can be designated as $E(r_i)$ where E is shorthand for "expected," r represents return, and the subscript i reflects the fact that the return can take on different possible values. $E(r_i)$ is the average return that the investor can anticipate receiving from IBM stock; it is equal to the summed value of each possible rate of return weighted by its probability: $E(r_i) = 2(0.5) + 8(0.5) = 5$.

The **expected utility** from holding IBM stock, $E[U(r_i)]$, is the average utility the investor can anticipate receiving from among the different possible returns on IBM stock; it is equal to the summed value of each possible utility weighted by its probability. The expected utility from holding IBM stock in our example is consequently 200, $E[U(r_i)] = 100(0.5) + 300(0.5) = 200$.

Graphically, the expected utility from investing in IBM stock with an equally likely return of 2 and 8 percent can be determined by constructing a dashed chord, AZ, that connects the heights of the total utility curve at returns of 2 and 8 percent. The height of chord AZ at the expected return of IBM stock (5 percent in our example) indicates the average or expected utility from investing in IBM stock, $E[U(r_i)]$. In Figure 5.15a, total utility rises from 100 to 300 between returns of 2 and 8 percent. Thus, to achieve a 200-unit increase in total utility (300 − 100 = 200) as the return increases by 6 percent (8 − 2 = 6), the average utility when we are at a return of 5 percent—halfway between a return of 2 and 8 percent—must be half the distance between 100 and 300 units of total utility. As a result, the expected utility of holding IBM stock is given by the height of chord AZ at the expected return of 5 percent: $E[U(r_i)] = 200$ in the case of Figure 5.15a.

Now compare the height of the risk-averse investor's total utility curve with the height of chord AZ at a return of 5 percent. The height of the total utility curve at 5 percent in Figure 5.15a (225) indicates the utility the investor would derive were IBM stock to provide a sure return (with no variance or risk) of 5 percent per dollar invested. Because the height of the total utility curve (225) exceeds the height of the straight-line segment AZ (200) at a return of 5 percent, this shows that the risk-averse investor gets more utility from a sure return of 5 percent (and no variance) than from investing in IBM stock that generates the same average return of 5 percent but through varying between returns of 2 and 8 percent. In general, an individual is deemed to be **risk averse** if he or she prefers a certain return to an uncertain prospect generating the same expected return.

risk averse
a state of preferring a certain return to an uncertain prospect that generates the same expected return

We can define risk aversion more formally by using some symbols for the heights of the total utility curve and the expected utility chord AZ. While the return on IBM's stock varies in our example, the average or expected return, $E(r_i)$, is 5 percent and is a fixed number. The height of the total utility curve at 5 percent thus can be written as $U[E(r_i)]$—the utility that would be generated by a sure return of 5 percent. Since a risk-averse investor prefers a sure return to an uncertain prospect generating the same expected return, this implies that the height of the total utility curve exceeds the height of the expected utility chord AZ at the expected return of the risky investment:

$$U[E(r_i)] > E[U(r_i)].$$

The preceding inequality will be true whenever an investor's total utility curve has a positive, but diminishing slope. Since the slope of the total utility curve is the marginal utility associated with various returns (see Chapter 3), risk aversion is present whenever the marginal utility of payoffs to an individual declines as a function of the size of those payoffs. With diminishing marginal utility, an investor will prefer a sure payoff of 5 percent per dollar invested to the prospect of earning 8 percent (3 percent above the average) half of the time and earning 2 percent (3 percent below the average) the other half. The marginal utility associated with each unit increase in the return between 5 and 8 percent will not be as great as the marginal utility associated with each unit decrease in the return between 5 and 2 percent if diminishing marginal utility applies to an asset's payoffs.

Among investors making choices involving uncertain payoffs, aversion appears to be the most common attitude toward risk, especially when the size of the payoffs involved is significant. For example, consider the risk of losing one's home to a fire. Most families would rather pay a $1,000 annual premium for an insurance policy covering loss due to fire than to face the prospect that in any given year the probability of losing their home valued at $100,000 to fire is 0.01 while the probability of there being no fire is 0.99. Most families, that is, would prefer a sure loss of $1,000 to an insurance company in the form of a premium than to take the gamble of not having fire insurance, even though the latter gamble involves the same expected return, $E(r_i)$, of –$1,000 per year, $E(r_i) = (-100{,}000)(0.01) + 0(0.99) = -1{,}000$.

Investor Preferences toward Risk: Risk Neutral and Risk Loving

Although risk aversion may be the most common attitude toward risk, it is not the only possible attitude. Indeed, there are two other views people can take toward risk. They can be neutral toward it or they can seek it out. These two alternatives are depicted by Figures 5.15b and 5.15c, respectively.

In the case of a risk-neutral investor (Figure 5.15b), the total utility curve is an upward-sloping straight line. For the same example involving IBM stock whose return is equally likely to be 2 or 8 percent, a risk-neutral investor will be indifferent between a sure return of 5 percent and investing in the risky IBM stock generating an expected return of 5 percent. Graphically, the height of the investor's total utility curve (50) equals the height of the expected utility chord, *AZ*, at the expected return of 5 percent. In general, an individual who is **risk neutral** gets the same utility from a certain return, $U[E(r_i)]$, as from an uncertain prospect generating the same expected return, $E[U(r_i)]$. Formally, that is, the following equation holds for a risk-neutral investor:

risk neutral
a state of deriving the same utility from a certain return as from an uncertain prospect generating the same expected return

$$U[E(r_i)] = E[U(r_i)].$$

For a risk-loving investor (Figure 5.15c), the total utility curve has an upward, but ever-increasing slope. In our IBM example, a risk-lover will prefer investing in the risky IBM stock generating an expected return of 5 percent over a sure return of 5 percent. Graphically, the height of the investor's total utility curve (150) is less than the height (175) of the expected utility chord, *AZ*, at the expected return of 5 percent. In general, an individual who is **risk loving** gets less utility from a certain return, $U[E(r_i)]$, than from an uncertain prospect generating the same expected return, $E[U(r_i)]$. Formally, the following equation holds for a risk lover:

risk loving
a state of deriving less utility from a certain return than from an uncertain prospect generating the same expected return

$$U[E(r_i)] < E[U(r_i)].$$

An example of risk-loving behavior involves the purchase of state lottery tickets even when one knows that the return on each dollar invested in such tickets is –50 percent. That is, the typical dollar invested in a state lottery has an expected payback of 50 cents. The return of –50 percent is much lower than the sure return of 0 percent one could obtain by keeping the money spent on lottery tickets in one's pockets instead.[10] While such risk-loving behavior certainly does occur, it tends to diminish as the gambling stakes increase. For example, more individuals at the casinos in Las Vegas play the $0.25 slot machines than the $100 blackjack tables.

[10]State lotteries make most of their money, and impose such a highly negative return on investors, by stretching out payments to lottery winners. For example, you may win $10 million in a state lottery drawing but that winning will be paid to you in equal installments over a period of 20 years. As we saw in Section 5.5, any dollars earned in the future are equal to less than their face value in terms of today's dollars. By contrast, the major casinos in Las Vegas pay gamblers an average return of –2 percent per dollar invested. That is, the typical dollar bet in a Las Vegas casino has an expected payoff of 98 cents. The less negative return in Las Vegas is due to the fact that casinos pay immediately when an "investor" gets lucky at a slot machine or blackjack table.

APPLICATION 5.6

Risk Aversion While Standing in Line

Many airline ticket counters, department of motor vehicles offices, federal customs checkpoints, banks, postal branches, college financial aid departments, and fast-food restaurants have a single line feeding to multiple clerks as opposed to separate lines for each clerk. This is the case even though a single-line system is unlikely to alter the average time a customer must spend waiting to see a clerk. After all, the line is not likely to affect either the total number of clerks or the number of customers waiting to see them.

A single line, however, does reduce the variance of a customer's waiting time since a customer is less likely to get into a "slow" or "fast" line. Holding constant the expected waiting time, the reduction in the variance of the waiting time will be appealing to customers if they are risk averse to spending time in lines. If a sure wait of 5 minutes is preferred to the prospect of a 0.5 probability each of a 2-minute and an 8-minute wait, customers will be happier under the single-line system. According to operations researchers, at least part of the reason customers prefer less variance when waiting in line is their sense of social justice. Specifically, relative to an expected wait of 5 minutes, experiencing only a 2-minute wait because one is fortunate enough to get into a "fast" line appears to add less to total utility than experiencing an 8-minute wait in a "slow" line subtracts from total utility, since the longer wait brings with it the added aggravation of seeing more recent arrivals who get into a fast line receive service first.[11]

[11]Richard C. Larson, "Perspectives on Queues: Social Justice and the Psychology of Queuing," *Operations Research*, 35 No. 6 (November–December 1987), pp. 895–905.

Minimizing Exposure to Risk

insurance
an arrangement by which the consumer pays a premium in return for the promise that the insurer will provide compensation for losses due to an accident, illness, fire, and so on

Living in a world where risk abounds, how can risk-averse individuals minimize their exposure to it? One method already mentioned is **insurance**. By paying a premium in return for the promise that the insurer will provide compensation for losses due to an accident, illness, fire, and so on, one effectively exchanges a gamble for a sure return. And risk-averse individuals would prefer a sure loss of $1,000 per year in the form of a premium paid to an insurance company over being uninsured and confronting the gamble of a 0.01 probability they will suffer a $100,000 loss versus the 0.99 probability they will not.

What is the maximum price that a risk-averse individual would be willing to pay for an insurance premium against a 0.01 probability of suffering a $100,000 loss? We can show that it is larger than $1,000, the expected value of the loss: $100,000(0.01) + $0(0.99) = $1,000. Figure 5.16 displays the total utility curve for an individual who owns a $100,000 home and faces the chance that the home will burn down with probability 0.01 in any given year. The expected utility from such a prospect is given by the height (U^*) of the expected utility chord, *AZ*, at an income of $99,000. (The expected payoff or income is $99,000 since with probability 0.99 there will be no fire and the individual's $100,000 home will be undamaged while with probability 0.01 the home will be reduced to ash and a value of $0.) Note that the same utility of U^* would also be provided by a certain income of $98,300 since the height of the total utility curve equals U^* at an income of $98,300. Thus, the individual whose total utility curve is depicted in Figure 5.16 would be willing to pay up to $1,700 per year to insure against a fire loss. The homeowner is indifferent between (a) paying $1,700 ($PP''$) per year in exchange for a certain income of $98,300 (the insurance company will fully compensate the individual for the $100,000 loss in case of a fire), and (b) the prospect of remaining uninsured and facing an expected, but not certain, loss of $1,000 ($P'P''$) and an income of $99,000.

Will an insurance company be willing to supply coverage against fire loss to the homeowner for less than the maximum the homeowner is willing to pay for such coverage? The answer, in all likelihood, is yes. This is because insurance companies typically insure a large pool of individuals with respect to the same risk. While the probability that any one

Figure 5.16

Pricing Insurance
In our example, a risk-averse homeowner is willing to pay up to $1,700 ($PP''$) per year to insure against a 0.01 probability of losing a $100,000 home to fire. By pooling a large enough number of equal risks, the minimum the insurance company will be willing to write an insurance policy for is $1,000 ($P'P''$) plus any costs associated with administering the policy.

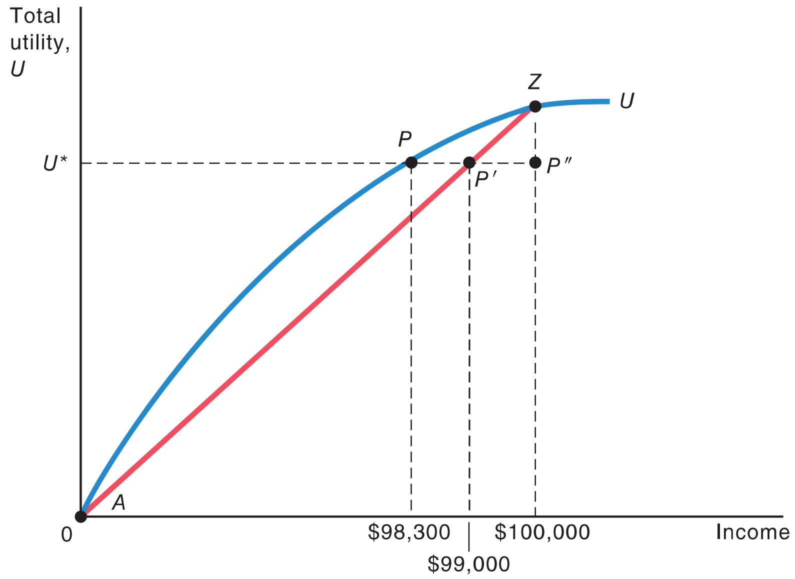

home will burn down in a given year may be subject to significant variability and therefore be difficult to predict, it is possible to predict much more precisely the total number of similarly situated homes in a large pool that will burn down in a given year. If the probability of any one house burning down is 0.01, for example, then an insurance company can be reasonably confident that, out of a pool of 10 million homes it insures, 100,000 will burn down in any given year. An apt analogy involves flipping coins. The deviation in the actual number of heads from the theoretically predicted probability will be much greater if a coin is flipped only once versus if the coin is flipped a large number of times.

By pooling a large enough number of similar fire risks, an insurance company will be able to predict reasonably well (that is, with very low variance) the total expected payments it will have to make on the fire insurance policies that it writes—$1,000 per home insured in our example. At the bare minimum, therefore, the insurance company should be willing to supply an insurance policy for $1,000 (plus any cost associated with administering the policy)—distance $P'P''$ in Figure 5.16. And since the maximum the risk-averse homeowner is willing to pay for such insurance (PP'') exceeds the minimum the insurance company needs to be paid ($P'P''$), an opportunity exists for mutually beneficial exchange.

In addition to insurance, another method of minimizing one's exposure to risk is through diversifying one's asset holdings. **Diversification** involves investing a given amount of resources in numerous independent projects instead of one single project. What may be an unacceptably large risk can thus be translated into more palatable small ones. The larger the number of independent projects, the more predictable the expected return on investment (as in our earlier coin-flipping example) and the lower the overall risk.

diversification
investing a given amount of resources in numerous independent projects instead of a single project in order to minimize exposure to risk

Organizational structures that promote the risk-spreading advantages attendant to diversification include partnerships and corporations whose stock is publicly traded. Suppose, for example, that a promising start-up company requires $100,000 in funds to get off the ground but there is some probability that the venture will not succeed and that any invested funds will be lost. If you are risk averse, you are likely to find investing in the venture more palatable if you can do it jointly with 10 other equal partners each chipping in $10,000 than if you have to underwrite the entire venture by yourself.

As with business partnerships, public trading of corporate stocks allows investors to diversify their asset portfolio by buying up a small number of shares of stock in a large number of companies. By not putting their financial eggs in a single asset basket, investors are able to mitigate their exposure to risk. This is because certain risks uniquely affect a single stock or a small group of stocks. And the investor's overall vulnerability to such unique risks will be lower the more diffuse a portfolio's holdings.

For example, take the case of the unexpected death in 1983 of Henry "Scoop" Jackson, a senator from the state of Washington with considerable clout on defense-spending issues. Jackson's death had no impact on most stocks but produced a decline in the price of Boeing Aircraft Company stock—Senator Jackson had been known as the "Senator from Boeing" for his lobbying efforts on behalf of the Washington-state-based firm.[12] Jackson's death also led to a rise in Lockheed Corporation's stock price since Lockheed was the largest defense manufacturer in Georgia, home to the senator expected to replace Jackson as ranking member of the Senate Armed Services Committee, Sam Nunn. An investor holding only Boeing *or* Lockheed stock at the time of Jackson's death would have confronted much greater risk (more volatility in asset returns) than would an investor with a more diversified portfolio. A well-diversified portfolio would include (a) many stocks that were not affected at all by Jackson's death, and (b) other stocks whose movement in the wake of Jackson's passing would work to cancel each other out (such as Boeing and Lockheed), thereby lowering the volatility of the overall portfolio.

[12]Brian E. Roberts, "A Dead Senator Tells No Lies: Seniority and the Distribution of Federal Benefits," *American Journal of Political Science*, 34, No. 1 (February 1990), pp. 31–58.

SUMMARY

- Consumer choice theory can be applied to a wide range of interesting and important policy questions.
- An excise subsidy is a form of subsidy in which the government pays part of a good's per-unit price and allows the consumer to buy as many units as desired at the subsidized price. Subsidies can also be made in the form of cash, as a lump-sum transfer. Consumer choice theory helps us discover that based on their own preferences, consumers will be better off if they receive cash.
- In general, subsidies cannot harm the recipients of the subsidies, at least in the case where they do not have to pay taxes to finance the subsidies. In the case of the ObamaCare legislation signed into law in 2010, however, subsidies can harm recipients to the extent that there is a mandate to purchase the subsidy. The subsidies associated with ObamaCare, moreover, provide disincentives to work and are associated with significant deadweight losses once the costs, to taxpayers, of funding the subsidies are taken into account.
- Public schools offer another common form of subsidy in which the government makes a certain quantity of a good available at no cost, or below the market price. Voucher programs allow parents to purchase education (with vouchers) at any school they choose. A variety of possible consequences of a school voucher program can be identified with the tools of consumer choice theory.
- When consumers directly bear the cost of such services as trash collection instead of paying a fixed annual fee, they have an incentive to cut down on the amount of the service they use, eliminating what was in effect a subsidy for heavy users of the service.
- A decision to save (or to borrow) is a decision to rearrange consumption between various time periods. Adaptations of consumer choice theory allow us to examine this decision and to see how it is affected by changes in earnings and in the interest rate.
- Consumer choice theory can help explain what types of financial assets an individual who is saving for the future should buy—stocks, Treasury bills, gold, and so on. Indifference curves demonstrate the tradeoff between expected risk and return. Total utility curves shed light on the way in which insurance operates to reduce individuals' risk.