5

Uncertainty and Consumer Behavior

So far, we have assumed that prices, incomes, and other variables are known with certainty. However, many of the choices that people make involve considerable uncertainty. Most people, for example, borrow to finance large purchases, such as a house or a college education, and plan to pay for them out of future income. But for most of us, future incomes are uncertain. Our earnings can go up or down; we can be promoted or demoted, or even lose our jobs. And if we delay buying a house or investing in a college education, we risk price increases that could make such purchases less affordable. How should we take these uncertainties into account when making major consumption or investment decisions?

Sometimes we must choose how much *risk* to bear. What, for example, should you do with your savings? Should you invest your money in something safe, such as a savings account, or something riskier but potentially more lucrative, such as the stock market? Another example is the choice of a job or career. Is it better to work for a large, stable company with job security but slim chance for advancement, or is it better to join (or form) a new venture that offers less job security but more opportunity for advancement?

To answer such questions, we must examine the ways that people can compare and choose among risky alternatives. We will do this by taking the following steps:

1. In order to compare the riskiness of alternative choices, we need to quantify risk. We therefore begin this chapter by discussing measures of risk.
2. We will examine people's preferences toward risk. Most people find risk undesirable, but some people find it more undesirable than others.
3. We will see how people can sometimes reduce or eliminate risk. Sometimes risk can be reduced by diversification, by buying insurance, or by investing in additional information.
4. In some situations, people must choose the amount of risk they wish to bear. A good example is investing in stocks or bonds. We will see that such investments involve tradeoffs between the monetary gain that one can expect and the riskiness of that gain.

In a world of uncertainty, individual behavior may sometimes seem unpredictable, even irrational, and perhaps contrary to the basic

CHAPTER OUTLINE

LIST OF EXAMPLES

assumptions of consumer theory. In the final section of this chapter, we offer an overview of the flourishing field of behavioral economics, which, by introducing important ideas from psychology, has broadened and enriched the study of microeconomics.

5.1 DESCRIBING RISK

To describe risk quantitatively, we begin by listing all the possible outcomes of a particular action or event, as well as the likelihood that each outcome will occur.[1] Suppose, for example, that you are considering investing in a company that explores for offshore oil. If the exploration effort is successful, the company's stock will increase from \$30 to \$40 per share; if not, the price will fall to \$20 per share. Thus there are two possible future outcomes: a \$40-per-share price and a \$20-per-share price.

Probability

• **probability** Likelihood that a given outcome will occur.

Probability is the likelihood that a given outcome will occur. In our example, the probability that the oil exploration project will be successful might be 1/4 and the probability that it is unsuccessful 3/4. (Note that the probabilities for all possible events must add up to 1.)

Our interpretation of probability can depend on the nature of the uncertain event, on the beliefs of the people involved, or both. One *objective* interpretation of probability relies on the frequency with which certain events tend to occur. Suppose we know that of the last 100 offshore oil explorations, 25 have succeeded and 75 failed. In that case, the probability of success of 1/4 is objective because it is based directly on the frequency of similar experiences.

But what if there are no similar past experiences to help measure probability? In such instances, objective measures of probability cannot be deduced and more subjective measures are needed. *Subjective probability* is the perception that an outcome will occur. This perception may be based on a person's judgment or experience, but not necessarily on the frequency with which a particular outcome has actually occurred in the past. When probabilities are subjectively determined, different people may attach different probabilities to different outcomes and thereby make different choices. For example, if the search for oil were to take place in an area where no previous searches had ever occurred, I might attach a higher subjective probability than you to the chance that the project will succeed: Perhaps I know more about the project or I have a better understanding of the oil business and can therefore make better use of our common information. Either different information or different abilities to process the same information can cause subjective probabilities to vary among individuals.

Regardless of the interpretation of probability, it is used in calculating two important measures that help us describe and compare risky choices. One measure tells us the *expected value* and the other the *variability* of the possible outcomes.

[1]Some people distinguish between uncertainty and risk along the lines suggested some 60 years ago by economist Frank Knight. *Uncertainty* can refer to situations in which many outcomes are possible but the likelihood of each is unknown. *Risk* then refers to situations in which we can list all possible outcomes and know the likelihood of each occurring. In this chapter, we will always refer to risky situations, but will simplify the discussion by using *uncertainty* and *risk* interchangeably.

Expected Value

The **expected value** associated with an uncertain situation is a weighted average of the **payoffs** or values associated with all possible outcomes. The probabilities of each outcome are used as weights. Thus the expected value measures the *central tendency*—the payoff or value that we would expect on average.

• **expected value** Probability-weighted average of the payoffs associated with all possible outcomes.

• **payoff** Value associated with a possible outcome.

Our offshore oil exploration example had two possible outcomes: Success yields a payoff of \$40 per share, failure a payoff of \$20 per share. Denoting "probability of" by Pr, we express the expected value in this case as

$$\text{Expected value} = \Pr(\text{success})(\$40/\text{share}) + \Pr(\text{failure})(\$20/\text{share})$$
$$= (1/4)(\$40/\text{share}) + (3/4)(\$20/\text{share}) = \$25/\text{share}$$

More generally, if there are two possible outcomes having payoffs X_1 and X_2 and if the probabilities of each outcome are given by $\Pr_1$ and $\Pr_2$, then the expected value is

$$E(X) = \Pr_1 X_1 + \Pr_2 X_2$$

When there are n possible outcomes, the expected value becomes

$$E(X) = \Pr_1 X_1 + \Pr_2 X_2 + \cdots + \Pr_n X_n$$

Variability

Variability is the extent to which the possible outcomes of an uncertain situation differ. To see why variability is important, suppose you are choosing between two part-time summer sales jobs that have the same expected income (\$1500). The first job is based entirely on commission—the income earned depends on how much you sell. There are two equally likely payoffs for this job: \$2000 for a successful sales effort and \$1000 for one that is less successful. The second job is salaried. It is very likely (.99 probability) that you will earn \$1510, but there is a .01 probability that the company will go out of business, in which case you would earn only \$510 in severance pay. Table 5.1 summarizes these possible outcomes, their payoffs, and their probabilities.

• **variability** Extent to which possible outcomes of an uncertain event differ.

Note that these two jobs have the same expected income. For Job 1, expected income is .5(\$2000) + .5(\$1000) = \$1500; for Job 2, it is .99(\$1510) + .01(\$510) = \$1500. However, the *variability* of the possible payoffs is different. We measure variability by recognizing that large differences between actual and expected payoffs (whether positive or negative) imply greater risk. We call these differences **deviations**. Table 5.2 shows the deviations of the possible income from the expected income from each job.

• **deviation** Difference between expected payoff and actual payoff.

By themselves, deviations do not provide a measure of variability. Why? Because they are sometimes positive and sometimes negative, and as you can see

TABLE 5.1 Income from Sales Jobs

	OUTCOME 1		OUTCOME 2		
	Probability	Income (\$)	Probability	Income (\$)	Expected Income (\$)
Job 1: Commission	.5	2000	.5	1000	1500
Job 2: Fixed Salary	.99	1510	.01	510	1500

TABLE 5.2 Deviations from Expected Income ($)

	Outcome 1	Deviation	Outcome 2	Deviation
Job 1	2000	500	1000	−500
Job 2	1510	10	510	−990

from Table 5.2, the average of the probability-weighted deviations is always 0.[2] To get around this problem, we square each deviation, yielding numbers that are always positive. We then measure variability by calculating the **standard deviation**: the square root of the average of the *squares* of the deviations of the payoffs associated with each outcome from their expected values.[3]

• **standard deviation** Square root of the weighted average of the squares of the deviations of the payoffs associated with each outcome from their expected values.

Table 5.3 shows the calculation of the standard deviation for our example. Note that the average of the squared deviations under Job 1 is given by

$$.5(\$250{,}000) + .5(\$250{,}000) = \$250{,}000$$

The standard deviation is therefore equal to the square root of $250,000, or $500. Likewise, the probability-weighted average of the squared deviations under Job 2 is

$$.99(\$100) + .01(\$980{,}100) = \$9900$$

The standard deviation is the square root of $9900, or $99.50. Thus the second job is much less risky than the first; the standard deviation of the incomes is much lower.[4]

The concept of standard deviation applies equally well when there are many outcomes rather than just two. Suppose, for example, that the first summer job yields incomes ranging from $1000 to $2000 in increments of $100 that are all equally likely. The second job yields incomes from $1300 to $1700 (again in increments of $100) that are also equally likely. Figure 5.1 shows the alternatives graphically. (If there had been only two equally probable outcomes, then the figure would be drawn as two vertical lines, each with a height of 0.5.)

You can see from Figure 5.1 that the first job is riskier than the second. The "spread" of possible payoffs for the first job is much greater than the spread for the second. As a result, the standard deviation of the payoffs associated with the first job is greater than that associated with the second.

TABLE 5.3 Calculating Variance ($)

	Outcome 1	Deviation Squared	Outcome 2	Deviation Squared	Weighted Average Deviation Squared	Standard Deviation
Job 1	2000	250,000	1000	250,000	250,000	500
Job 2	1510	100	510	980,100	9900	99.50

[2]For Job 1, the average deviation is .5($500) + .5(−$500) = 0; for Job 2 it is .99($10) + .01(−$990) = 0.

[3]Another measure of variability, *variance*, is the square of the standard deviation.

[4]In general, when there are two outcomes with payoffs X_1 and X_2, occurring with probability $\Pr_1$ and $\Pr_2$, and $E(X)$ is the expected value of the outcomes, the standard deviation is given by σ, where

$$\sigma^2 = \Pr_1[(X_1 - E(X))^2] + \Pr_2[(X_2 - E(X))^2]$$

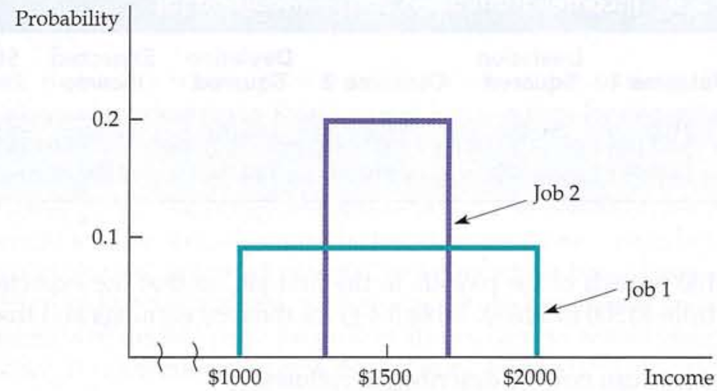

FIGURE 5.1 Outcome Probabilities for Two Jobs

The distribution of payoffs associated with Job 1 has a greater spread and a greater standard deviation than the distribution of payoffs associated with Job 2. Both distributions are flat because all outcomes are equally likely.

In this particular example, all payoffs are equally likely. Thus the curves describing the probabilities for each job are flat. In many cases, however, some payoffs are more likely than others. Figure 5.2 shows a situation in which the most extreme payoffs are the least likely. Again, the salary from Job 1 has a greater standard deviation. From this point on, we will use the standard deviation of payoffs to measure the degree of risk.

Decision Making

Suppose you are choosing between the two sales jobs described in our original example. Which job would you take? If you dislike risk, you will take the second job: It offers the same expected income as the first but with less risk. But suppose

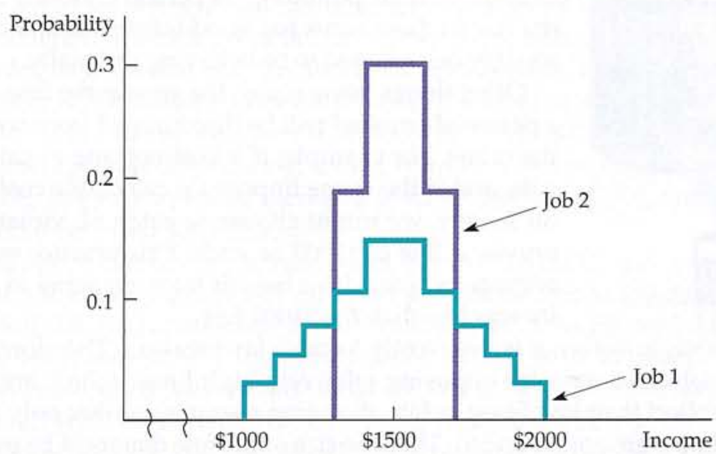

FIGURE 5.2 Unequal Probability Outcomes

The distribution of payoffs associated with Job 1 has a greater spread and a greater standard deviation than the distribution of payoffs associated with Job 2. Both distributions are peaked because the extreme payoffs are less likely than those near the middle of the distribution.

TABLE 5.4 Incomes from Sales Jobs—Modified ($)

	Outcome 1	Deviation Squared	Outcome 2	Deviation Squared	Expected Income	Standard Deviation
Job 1	2100	250,000	1100	250,000	1600	500
Job 2	1510	100	510	980,100	1500	99.50

we add $100 to each of the payoffs in the first job, so that the expected payoff increases from $1500 to $1600. Table 5.4 gives the new earnings and the squared deviations.

The two jobs can now be described as follows:

Job 1: Expected Income = $1600 Standard Deviation = $500

Job 2: Expected Income = $1500 Standard Deviation = $99.50

Job 1 offers a higher expected income but is much riskier than Job 2. Which job is preferred depends on the individual. While an aggressive entrepreneur who doesn't mind taking risks might choose Job 1, with the higher expected income and higher standard deviation, a more conservative person might choose the second job.

People's attitudes toward risk affect many of the decisions they make. In Example 5.1 we will see how attitudes toward risk affect people's willingness to break the law, and how this has implications for the fines that should be set for various violations. Then in Section 5.2, we will further develop our theory of consumer choice by examining people's risk preferences in greater detail.

EXAMPLE 5.1 Deterring Crime

Fines may be better than incarceration in deterring certain types of crimes, such as speeding, double-parking, tax evasion, and air polluting.[5] A person choosing to violate the law in these ways has good information and can reasonably be assumed to be behaving rationally.

Other things being equal, the greater the fine, the more a potential criminal will be discouraged from committing the crime. For example, if it cost nothing to catch criminals, and if the crime imposed a calculable cost of $1000 on society, we might choose to catch all violations and impose a fine of $1000 on each. This practice would discourage people whose benefit from engaging in the activity was less than the $1000 fine.

In practice, however, it is very costly to catch lawbreakers. Therefore, we save on administrative costs by imposing relatively high fines (which are no more costly to collect than low fines), while allocating resources so that only a fraction of the violators are apprehended. Thus the size of the fine that must be imposed to

[5]This discussion builds indirectly on Gary S. Becker, "Crime and Punishment: An Economic Approach," *Journal of Political Economy* (March/April 1968): 169–217. See also A. Mitchell Polinsky and Steven Shavell, "The Optimal Tradeoff Between the Probability and the Magnitude of Fines," *American Economic Review* 69 (December 1979): 880–91.

discourage criminal behavior depends on the attitudes toward risk of potential violators.

Suppose that a city wants to deter people from double-parking. By double-parking, a typical resident saves $5 in terms of his own time for engaging in activities that are more pleasant than searching for a parking space. If it costs nothing to catch a double-parker, a fine of just over $5—say, $6—should be assessed every time he double-parks. This policy will ensure that the net benefit of double-parking (the $5 benefit less the $6 fine) would be less than zero. Our citizen will therefore choose to obey the law. In fact, all potential violators whose benefit was less than or equal to $5 would be discouraged, while a few whose benefit was greater than $5 (say, someone who double-parks because of an emergency) would violate the law.

In practice, it is too costly to catch all violators. Fortunately, it's also unnecessary. The same deterrence effect can be obtained by assessing a fine of $50 and catching only one in ten violators (or perhaps a fine of $500 with a one-in-100 chance of being caught). In each case, the expected penalty is $5, i.e., [$50][.1] or [$500][.01]. A policy that combines a high fine and a low probability of apprehension is likely to reduce enforcement costs. This approach is especially effective if drivers don't like to take risks. In our example, a $50 fine with a .1 probability of being caught might discourage most people from violating the law. We will examine attitudes toward risk in the next section.

5.2 PREFERENCES TOWARD RISK

We used a job example to show how people might evaluate risky outcomes, but the principles apply equally well to other choices. In this section, we concentrate on consumer choices generally and on the *utility* that consumers obtain from choosing among risky alternatives. To simplify things, we'll consider the utility that a consumer gets from his or her income—or, more appropriately, the market basket that the consumer's income can buy. We now measure payoffs, therefore, in terms of utility rather than dollars.

Figure 5.3(a) shows how we can describe one woman's preferences toward risk. The curve $0E$, which gives her utility function, tells us the level of utility (on the vertical axis) that she can attain for each level of income (measured in thousands of dollars on the horizontal axis). The level of utility increases from 10 to 16 to 18 as income increases from $10,000 to $20,000 to $30,000. But note that *marginal utility* is diminishing, falling from 10 when income increases from 0 to $10,000, to 6 when income increases from $10,000 to $20,000, and to 2 when income increases from $20,000 to $30,000.

In §3.1, we explained that a utility function assigns a level of utility to each possible market basket.

In §3.5, marginal utility is described as the additional satisfaction obtained by consuming an additional amount of a good.

Now suppose that our consumer has an income of $15,000 and is considering a new but risky sales job that will either double her income to $30,000 or cause it to fall to $10,000. Each possibility has a probability of .5. As Figure 5.3 (a) shows, the utility level associated with an income of $10,000 is 10 (at point A) and the utility level associated with an income of $30,000 is 18 (at E). The risky job must be compared with the current $15,000 job, for which the utility is 13.5 (at B).

To evaluate the new job, she can calculate the expected value of the resulting income. Because we are measuring value in terms of her utility, we must calculate the **expected utility** E(u) that she can obtain. The expected utility is *the sum of the utilities associated with all possible outcomes, weighted by the probability that each outcome will occur*. In this case expected utility is

• **expected utility** Sum of the utilities associated with all possible outcomes, weighted by the probability that each outcome will occur.

$$E(u) = (1/2)u(\$10{,}000) + (1/2)u(\$30{,}000) = 0.5(10) + 0.5(18) = 14$$

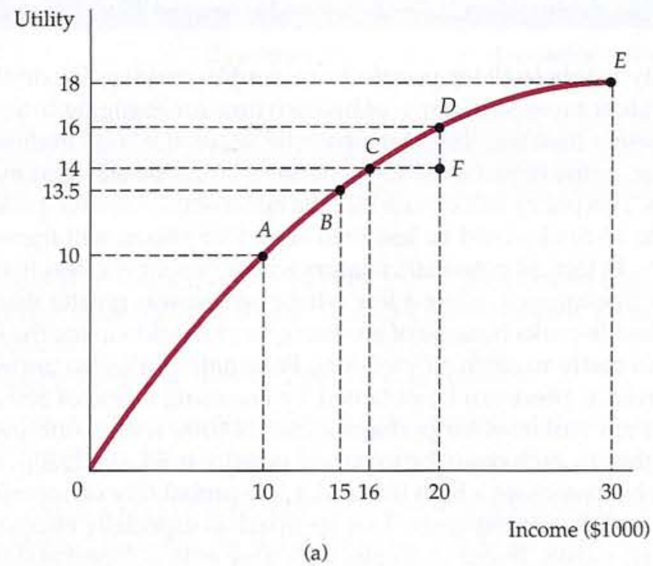

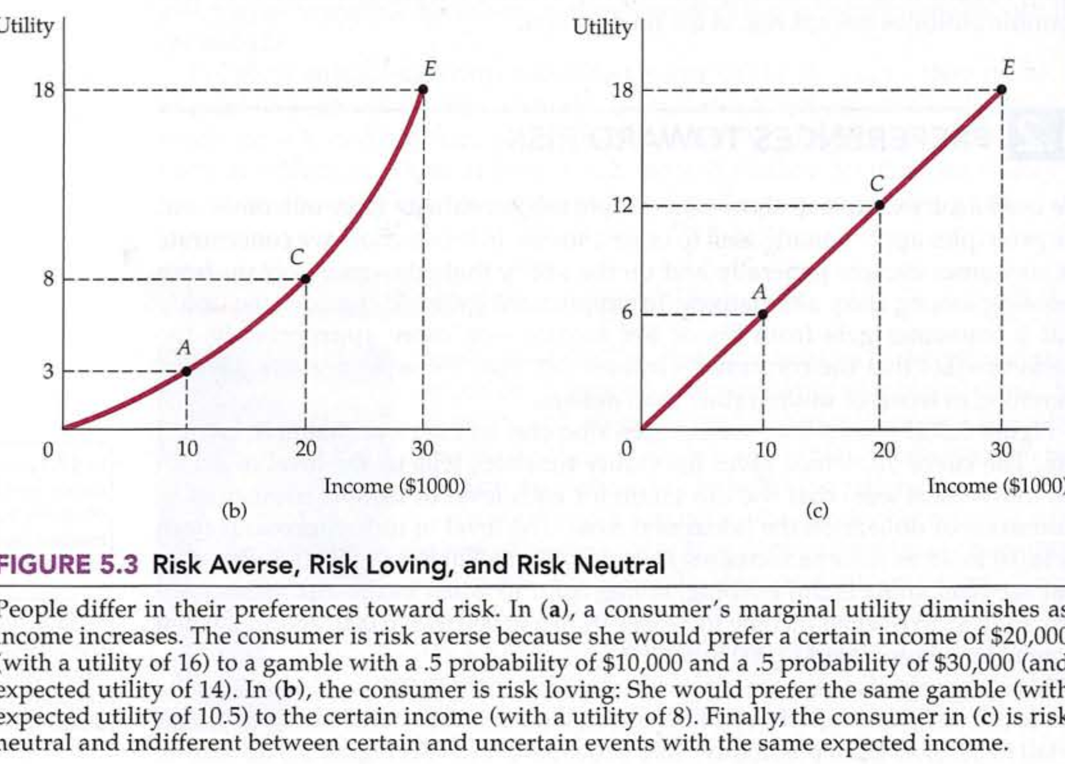

FIGURE 5.3 Risk Averse, Risk Loving, and Risk Neutral

People differ in their preferences toward risk. In (**a**), a consumer's marginal utility diminishes as income increases. The consumer is risk averse because she would prefer a certain income of $20,000 (with a utility of 16) to a gamble with a .5 probability of $10,000 and a .5 probability of $30,000 (and expected utility of 14). In (**b**), the consumer is risk loving: She would prefer the same gamble (with expected utility of 10.5) to the certain income (with a utility of 8). Finally, the consumer in (**c**) is risk neutral and indifferent between certain and uncertain events with the same expected income.

The risky new job is thus preferred to the original job because the expected utility of 14 is greater than the original utility of 13.5.

The old job involved no risk—it guaranteed an income of $15,000 and a utility level of 13.5. The new job is risky but offers both a higher expected income ($20,000) and, more importantly, a higher expected utility. If the woman wishes to increase her expected utility, she will take the risky job.

Different Preferences Toward Risk

People differ in their willingness to bear risk. Some are risk averse, some risk loving, and some risk neutral. An individual who is **risk averse** prefers a certain given income to a risky income with the same expected value. (Such a person has a diminishing marginal utility of income.) Risk aversion is the most common attitude toward risk. To see that most people are risk averse most of the time, note that most people not only buy life insurance, health insurance, and car insurance, but also seek occupations with relatively stable wages.

• **risk averse** Condition of preferring a certain income to a risky income with the same expected value.

Figure 5.3(a) applies to a woman who is risk averse. Suppose hypothetically that she can have either a certain income of $20,000, or a job yielding an income of $30,000 with probability .5 and an income of $10,000 with probability .5 (so that the expected income is also $20,000). As we saw, the expected utility of the uncertain income is 14—an average of the utility at point $A(10)$ and the utility at $E(18)$—and is shown by F. Now we can compare the expected utility associated with the risky job to the utility generated if $20,000 were earned without risk. This latter utility level, 16, is given by D in Figure 5.3(a). It is clearly greater than the expected utility of 14 associated with the risky job.

For a risk-averse person, losses are more important (in terms of the change in utility) than gains. Again, this can be seen from Figure 5.3(a). A $10,000 increase in income, from $20,000 to $30,000, generates an increase in utility of two units; a $10,000 decrease in income, from $20,000 to $10,000, creates a loss of utility of six units.

A person who is **risk neutral** is indifferent between a certain income and an uncertain income with the same expected value. In Figure 5.3(c) the utility associated with a job generating an income of either $10,000 or $30,000 with equal probability is 12, as is the utility of receiving a certain income of $20,000. As you can see from the figure, the marginal utility of income is constant for a risk-neutral person.[6]

• **risk neutral** Condition of being indifferent between a certain income and an uncertain income with the same expected value.

Finally, an individual who is **risk loving** prefers an uncertain income to a certain one, even if the expected value of the uncertain income is less than that of the certain income. Figure 5.3(b) shows this third possibility. In this case, the expected utility of an uncertain income, which will be either $10,000 with probability .5 or $30,000 with probability .5, is *higher* than the utility associated with a certain income of $20,000. Numerically,

• **risk loving** Condition of preferring a risky income to a certain income with the same expected value.

$$E(u) = .5u(\$10{,}000) + .5u(\$30{,}000) = .5(3) + .5(18) = 10.5 > u(\$20{,}000) = 8$$

Of course, some people may be averse to some risks and act like risk lovers with respect to others. For example, many people purchase life insurance and are conservative with respect to their choice of jobs, but still enjoy gambling. Some criminologists might describe criminals as risk lovers, especially if they commit crimes despite a high prospect of apprehension and punishment. Except for such special cases, however, few people are risk loving, at least with respect to major purchases or large amounts of income or wealth.

Risk Premium The **risk premium** is the maximum amount of money that a risk-averse person will pay to avoid taking a risk. In general, the magnitude of the

• **risk premium** Maximum amount of money that a risk-averse person will pay to avoid taking a risk.

[6]Thus, when people are risk neutral, the income they earn can be used as an indicator of well-being. A government policy that doubles incomes would then also double their utility. At the same time, government policies that alter the risks that people face, without changing their expected incomes, would not affect their well-being. Risk neutrality allows a person to avoid the complications that might be associated with the effects of governmental actions on the riskiness of outcomes.

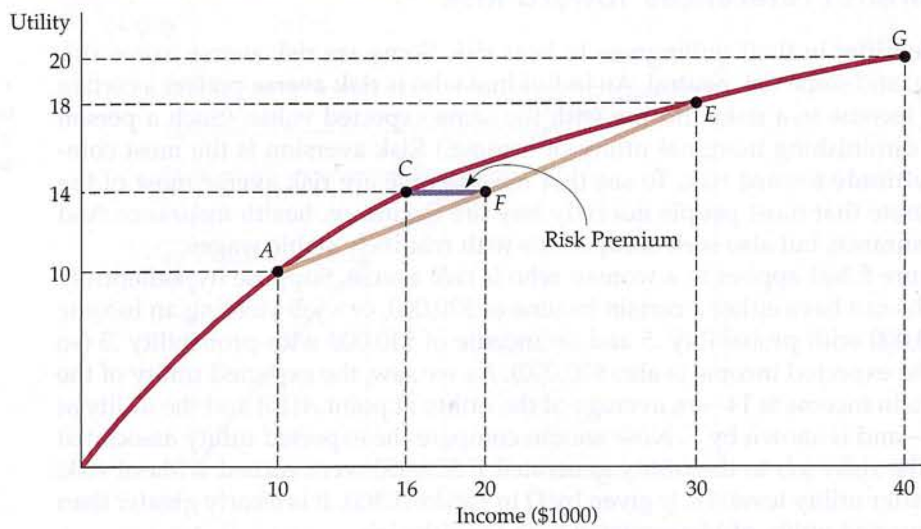

FIGURE 5.4 Risk Premium

The risk premium, *CF*, measures the amount of income that an individual would give up to leave her indifferent between a risky choice and a certain one. Here, the risk premium is \$4000 because a certain income of \$16,000 (at point C) gives her the same expected utility (14) as the uncertain income (a .5 probability of being at point *A* and a .5 probability of being at point *E*) that has an expected value of \$20,000.

risk premium depends on the risky alternatives that the person faces. To determine the risk premium, we have reproduced the utility function of Figure 5.3(a) in Figure 5.4 and extended it to an income of \$40,000. Recall that an expected utility of 14 is achieved by a woman who is going to take a risky job with an expected income of \$20,000. This outcome is shown graphically by drawing a horizontal line to the vertical axis from point *F*, which bisects straight line *AE* (thus representing an average of \$10,000 and \$30,000). But the utility level of 14 can also be achieved if the woman has a *certain* income of \$16,000, as shown by dropping a vertical line from point C. Thus, the risk premium of \$4000, given by line segment *CF*, is the amount of expected income (\$20,000 minus \$16,000) that she would give up in order to remain indifferent between the risky job and a hypothetical job that would pay her a certain income of \$16,000.

Risk Aversion and Income The extent of an individual's risk aversion depends on the nature of the risk and on the person's income. Other things being equal, risk-averse people prefer a smaller variability of outcomes. We saw that when there are two outcomes—an income of \$10,000 and an income of \$30,000—the risk premium is \$4000. Now consider a second risky job, also illustrated in Figure 5.4. With this job, there is a .5 probability of receiving an income of \$40,000, with a utility level of 20, and a .5 probability of getting an income of \$0, with a utility level of 0. The expected income is again \$20,000, but the expected utility is only 10:

$$\text{Expected utility} = .5u(\$0) + .5u(\$40{,}000) = 0 + .5(20) = 10$$

Compared to a hypothetical job that pays \$20,000 with certainty, the person holding this risky job gets 6 fewer units of expected utility: 10 rather than 16 units. At the same time, however, this person could also get 10 units of utility from a job that pays \$10,000 with certainty. Thus the risk premium in this case is \$10,000,

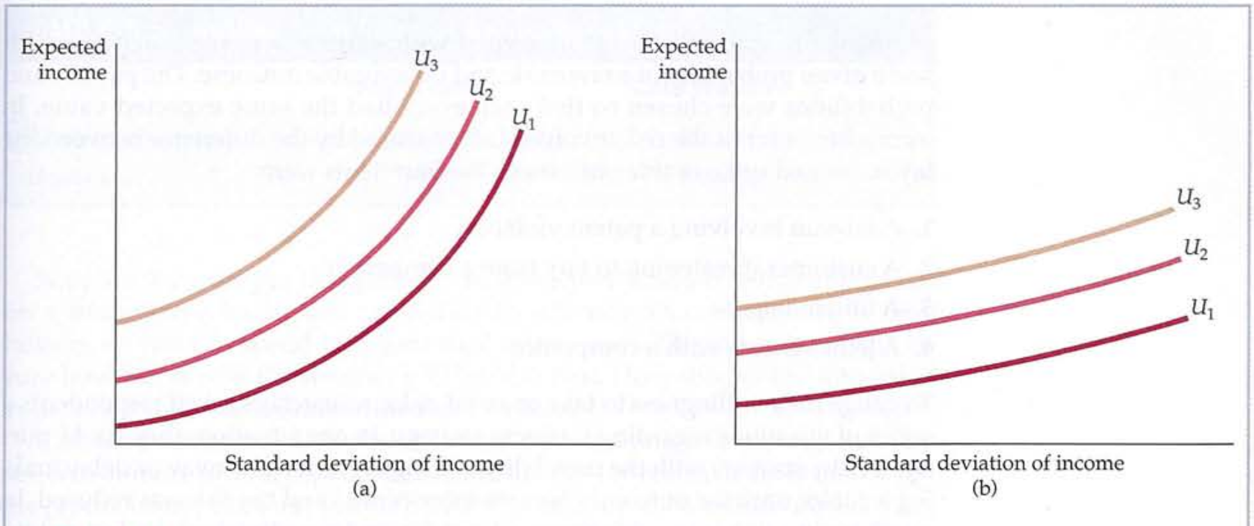

FIGURE 5.5 Risk Aversion and Indifference Curves

Part **(a)** applies to a person who is highly risk averse: An increase in this individual's standard deviation of income requires a large increase in expected income if he or she is to remain equally well off. Part **(b)** applies to a person who is only slightly risk averse: An increase in the standard deviation of income requires only a small increase in expected income if he or she is to remain equally well off.

because this person would be willing to give up $10,000 of her $20,000 expected income to avoid bearing the risk of an uncertain income. The greater the variability of income, the more the person would be willing to pay to avoid the risky situation.

Risk Aversion and Indifference Curves We can also describe the extent of a person's risk aversion in terms of indifference curves that relate expected income to the variability of income, where the latter is measured by the standard deviation. Figure 5.5 shows such indifference curves for two individuals, one who is highly risk averse and another who is only slightly risk averse. Each indifference curve shows the combinations of expected income and standard deviation of income that give the individual the same amount of utility. Observe that all of the indifference curves are upward sloping: Because risk is undesirable, the greater the amount of risk, the greater the expected income needed to make the individual equally well off.

In §3.1, we define an indifference curve as all market baskets that generate the same level of satisfaction for a consumer.

Figure 5.5(a) describes an individual who is highly risk averse. Observe that in order to leave this person equally well off, an increase in the standard deviation of income requires a large increase in expected income. Figure 5.5(b) applies to a slightly risk-averse person. In this case, a large increase in the standard deviation of income requires only a small increase in expected income.

EXAMPLE 5.2 Business Executives and the Choice of Risk

Are business executives more risk loving than most people? When they are presented with alternative strategies, some risky, some safe, which do they choose? In one study, 464 executives were asked to respond to a questionnaire describing risky situations that an individual might face as vice president of a hypothetical company.[7] Respondents were presented with four risky events, each of which had a given probability of a favorable and unfavorable outcome. The payoffs and probabilities were chosen so that each event had the same expected value. In increasing order of the risk involved (as measured by the difference between the favorable and unfavorable outcomes), the four items were:

1. A lawsuit involving a patent violation
2. A customer threatening to buy from a competitor
3. A union dispute
4. A joint venture with a competitor

To gauge their willingness to take or avoid risks, researchers asked respondents a series of questions regarding business strategy. In one situation, they could pursue a risky strategy with the possibility of a high return right away or delay making a choice until the outcomes became more certain and the risk was reduced. In another situation, respondents could opt for an immediately risky but potentially profitable strategy that could lead to a promotion, or they could delegate the decision to someone else, which would protect their job but eliminate the promotion possibility.

The study found that executives vary substantially in their preferences toward risk. Roughly 20 percent indicated that they were relatively risk neutral; 40 percent opted for the more risky alternatives; and 20 percent were clearly risk averse (20 percent did not respond). More importantly, executives (including those who chose risky alternatives) typically made efforts to reduce or eliminate risk, usually by delaying decisions and collecting more information.

We will return to the use of indifference curves as a means of describing risk aversion in Section 5.4, where we discuss the demand for risky assets. First, however, we will turn to the ways in which an individual can reduce risk.

5.3 REDUCING RISK

As the recent growth in state lotteries shows, people sometimes choose risky alternatives that suggest risk-loving rather than risk-averse behavior. Most people, however, spend relatively small amounts on lottery tickets and casinos. When more important decisions are involved, they are generally risk averse. In this section, we describe three ways by which both consumers and businesses commonly reduce risks: *diversification*, *insurance*, and *obtaining more information* about choices and payoffs.

Diversification

Recall the old saying, "Don't put all your eggs in one basket." Ignoring this advice is unnecessarily risky: If your basket turns out to be a bad bet, all will be lost. Instead, you can reduce risk through **diversification**: allocating your resources to a variety of activities whose outcomes are not closely related.

• **diversification** Practice of reducing risk by allocating resources to a variety of activities whose outcomes are not closely related.

[7]This example is based on Kenneth R. MacCrimmon and Donald A. Wehrung, "The Risk In-Basket," *Journal of Business* 57 (1984): 367–87.

TABLE 5.5 Income from Sales of Appliances ($)

	Hot Weather	Cold Weather
Air conditioner sales	30,000	12,000
Heater sales	12,000	30,000

Suppose, for example, that you plan to take a part-time job selling appliances on a commission basis. You can decide to sell only air conditioners or only heaters, or you can spend half your time selling each. Of course, you can't be sure how hot or cold the weather will be next year. How should you apportion your time in order to minimize the risk involved?

Risk can be minimized by *diversification*—by allocating your time so that you sell two or more products (whose sales are not closely related) rather than a single product. Suppose there is a 0.5 probability that it will be a relatively hot year, and a 0.5 probability that it will be cold. Table 5.5 gives the earnings that you can make selling air conditioners and heaters.

If you sell only air conditioners or only heaters, your actual income will be either $12,000 or $30,000, but your expected income will be $21,000 (.5[$30,000] + .5[$12,000]). But suppose you diversify by dividing your time evenly between the two products. In that case, your income will certainly be $21,000, regardless of the weather. If the weather is hot, you will earn $15,000 from air conditioner sales and $6000 from heater sales; if it is cold, you will earn $6000 from air conditioners and $15,000 from heaters. In this instance, diversification eliminates all risk.

Of course, diversification is not always this easy. In our example, heater and air conditioner sales are **negatively correlated variables**—they tend to move in opposite directions; whenever sales of one are strong, sales of the other are weak. But the principle of diversification is a general one: As long as you can allocate your resources toward a variety of activities whose outcomes are *not* closely related, you can eliminate some risk.

• **negatively correlated variables** Variables having a tendency to move in opposite directions.

The Stock Market Diversification is especially important for people who invest in the stock market. On any given day, the price of an individual stock can go up or down by a large amount, but some stocks rise in price while others fall. An individual who invests all her money in a single stock (i.e., puts all her eggs in one basket) is therefore taking much more risk than necessary. Risk can be reduced—although not eliminated—by investing in a portfolio of ten or twenty different stocks. Likewise, you can diversify by buying shares in **mutual funds**: organizations that pool funds of individual investors to buy a large number of different stocks. There are thousands of mutual funds available today for both stocks and bonds. These funds are popular because they reduce risk through diversification and because their fees are typically much lower than the cost of assembling one's own portfolio of stocks.

• **mutual fund** Organization that pools funds of individual investors to buy a large number of different stocks or other financial assets.

In the case of the stock market, not all risk is diversifiable. Although some stocks go up in price when others go down, stock prices are to some extent **positively correlated variables**: They tend to move in the same direction in response to changes in economic conditions. For example, the onset of a severe recession, which is likely to reduce the profits of many companies, may be accompanied by a decline in the overall market. Even with a diversified portfolio of stocks, therefore, you still face some risk.

• **positively correlated variables** Variables having a tendency to move in the same direction.

Insurance

We have seen that risk-averse people are willing to pay to avoid risk. In fact, if the cost of insurance is equal to the expected loss (e.g., a policy with an expected loss of $1000 will cost $1000), risk-averse people will buy enough insurance to recover fully from any financial losses they might suffer.

Why? The answer is implicit in our discussion of risk aversion. Buying insurance assures a person of having the same income whether or not there is a loss. Because the insurance cost is equal to the expected loss, this certain income is equal to the expected income from the risky situation. For a risk-averse consumer, the guarantee of the same income regardless of the outcome generates more utility than would be the case if that person had a high income when there was no loss and a low income when a loss occurred.

To clarify this point, let's suppose a homeowner faces a 10-percent probability that his house will be burglarized and he will suffer a $10,000 loss. Let's assume he has $50,000 worth of property. Table 5.6 shows his wealth in two situations—with insurance costing $1000 and without insurance.

Note that expected wealth is the same ($49,000) in both situations. The variability, however, is quite different. As the table shows, with no insurance the standard deviation of wealth is $3000; with insurance, it is 0. If there is no burglary, the uninsured homeowner gains $1000 relative to the insured homeowner. But with a burglary, the uninsured homeowner loses $9000 relative to the insured homeowner. Remember: for a risk-averse individual, losses count more (in terms of changes in utility) than gains. A risk-averse homeowner, therefore, will enjoy higher utility by purchasing insurance.

The Law of Large Numbers Consumers usually buy insurance from companies that specialize in selling it. Insurance companies are firms that offer insurance because they know that when they sell a large number of policies, they face relatively little risk. The ability to avoid risk by operating on a large scale is based on the *law of large numbers*, which tells us that although single events may be random and largely unpredictable, the average outcome of many similar events can be predicted. For example, I may not be able to predict whether a coin toss will come out heads or tails, but I know that when many coins are flipped, approximately half will turn up heads and half tails. Likewise, if I am selling automobile insurance, I cannot predict whether a particular driver will have an accident, but I can be reasonably sure, judging from past experience, what fraction of a large group of drivers will have accidents.

Actuarial Fairness By operating on a large scale, insurance companies can be sure that over a sufficiently large number of events, total premiums paid in will be equal to the total amount of money paid out. Let's return to our burglary example. A man knows that there is a 10-percent probability that his house will be burgled; if it is, he will suffer a $10,000 loss. Prior to facing this risk, he calculates the expected loss to be $1000 (.10 × $10,000). The risk involved is considerable,

TABLE 5.6 The Decision to Insure ($)

Insurance	Burglary (Pr = .1)	No Burglary (Pr = .9)	Expected Wealth	Standard Deviation
No	40,000	50,000	49,000	3000
Yes	49,000	49,000	49,000	0

however, because there is a 10-percent probability of a large loss. Now suppose that 100 people are similarly situated and that all of them buy burglary insurance from the same company. Because they all face a 10-percent probability of a $10,000 loss, the insurance company might charge each of them a premium of $1000. This $1000 premium generates an insurance fund of $100,000 from which losses can be paid. The insurance company can rely on the law of large numbers, which holds that the expected loss to the 100 individuals as a whole is likely to be very close to $1000 each. The total payout, therefore, will be close to $100,000, and the company need not worry about losing more than that.

When the insurance premium is equal to the expected payout, as in the example above, we say that the insurance is **actuarially fair**. But because they must cover administrative costs and make some profit, insurance companies typically charge premiums *above* expected losses. If there are a sufficient number of insurance companies to make the market competitive, these premiums will be close to actuarially fair levels. In some states, however, insurance premiums are regulated in order to protect consumers from "excessive" premiums. We will examine government regulation of markets in detail in Chapters 9 and 10 of this book.

• **actuarially fair** Characterizing a situation in which an insurance premium is equal to the expected payout.

In recent years, some insurance companies have come to the view that catastrophic disasters such as earthquakes are so unique and unpredictable that they cannot be viewed as diversifiable risks. Indeed, as a result of losses from past disasters, these companies do not feel that they can determine actuarially fair insurance rates. In California, for example, the state itself has had to enter the insurance business to fill the gap created when private companies refused to sell earthquake insurance. The state-run pool offers less insurance coverage at higher rates than was previously offered by private insurers.

EXAMPLE 5.3 The Value of Title Insurance When Buying a House

Suppose you are buying your first house. To close the sale, you will need a deed that gives you clear "title." Without such a clear title, there is always a chance that the seller of the house is not its true owner. Of course, the seller could be engaging in fraud, but it is more likely that the seller is unaware of the exact nature of his or her ownership rights. For example, the owner may have borrowed heavily, using the house as "collateral" for a loan. Or the property might carry with it a legal requirement that limits the use to which it may be put.

Suppose you are willing to pay $300,000 for the house, but you believe there is a one-in-twenty chance that careful research will reveal that the seller does not actually own the property. The property would then be worth nothing. If there were no insurance available, a risk-neutral person would bid at most $285,000 for the property (.95[$300,000] + .05[0]). However, if you expect to tie up most of your assets in the house, you would probably be risk averse and, therefore, bid much less—say, $230,000.

In situations such as this, it is clearly in the interest of the buyer to be sure that there is no risk of a lack of full ownership. The buyer does this by purchasing "title insurance." The title insurance company researches the history of the property, checks to see whether any legal liabilities are attached to it, and generally assures itself that there is no ownership problem. The insurance company then agrees to bear any remaining risk that might exist.

Because the title insurance company is a specialist in such insurance and can collect the relevant information relatively easily, the cost of title insurance is often less than the expected value of the loss involved. A fee of $1500 for title insurance is not unusual, even though the expected loss can be much higher. It is also in the interest of sellers to provide title insurance, because all but the most risk-loving buyers will pay much more for a house when it is insured than when it is not. In fact, most states require sellers to provide title insurance before a sale can be completed. In addition, because mortgage lenders are all concerned about such risks, they usually require new buyers to have title insurance before issuing a mortgage.

The Value of Information

• **value of complete information** Difference between the expected value of a choice when there is complete information and the expected value when information is incomplete.

People often make decisions based on limited information. If more information were available, one could make better predictions and reduce risk. Because information is a valuable commodity, people will pay for it. The **value of complete information** is the difference between the expected value of a choice when there is complete information and the expected value when information is incomplete.

To see how information can be valuable, suppose you manage a clothing store and must decide how many suits to order for the fall season. If you order 100 suits, your cost is $180 per suit. If you order only 50 suits, your cost increases to $200. You know that you will be selling suits for $300 each, but you are not sure how many you can sell. All suits not sold can be returned, but for only half of what you paid for them. Without additional information, you will act on your belief that there is a .5 probability that you will sell 100 suits and a .5 probability that you will sell 50. Table 5.7 gives the profit that you would earn in each of these two cases.

Without additional information, you would choose to buy 100 suits if you were risk neutral, taking the chance that your profit might be either $12,000 or $1500. But if you were risk averse, you might buy 50 suits: In that case, you would know for sure that your profit would be $5000.

With complete information, you can place the correct order regardless of future sales. If sales were going to be 50 and you ordered 50 suits, your profits would be $5000. If, on the other hand, sales were going to be 100 and you ordered 100 suits, your profits would be $12,000. Because both outcomes are equally likely, your expected profit with complete information would be $8500. The value of information is computed as

	Expected value with complete information:	$8500
Less:	Expected value with uncertainty (buy 100 suits):	–6750
	Value of complete information	$1750

Thus it is worth paying up to $1750 to obtain an accurate prediction of sales. Even though forecasting is inevitably imperfect, it may be worth investing in a marketing study that provides a reasonable forecast of next year's sales.

TABLE 5.7 Profits from Sales of Suits ($)

	Sales of 50	Sales of 100	Expected Profit
Buy 50 suits	5000	5000	5000
Buy 100 suits	1500	12,000	6750

EXAMPLE 5.4 The Value of Information in the Dairy Industry

Historically, the U.S. dairy industry has allocated its advertising expenditures more or less uniformly throughout the year.[8] But per-capita consumption of milk has declined over the years—a situation that has stirred producers to look for new strategies to encourage milk consumption. One strategy would be to increase advertising expenditures and to continue advertising at a uniform rate throughout the year. A second strategy would be to invest in market research in order to obtain more information about the seasonal demand for milk; marketers could then reallocate expenditures so that advertising was most intense when the demand for milk was greatest.

Research into milk demand shows that sales follow a seasonal pattern, with demand being greatest during the spring and lowest during the summer and early fall. The price elasticity of milk demand is negative but small and the income elasticity positive and large. Most important is the fact that milk advertising has the most effect on sales when consumers have the strongest preference for the product (March, April, and May) and least when preferences are weakest (August, September, and October).

In §2.4, we define the price elasticity of demand as the percentage change in quantity demanded resulting from a 1-percent change in the price of a good.

In this case, the cost of obtaining seasonal information about milk demand is relatively low and the value of the information substantial. To estimate this value, we can compare the actual sales of milk during a typical year with sales levels that would have been reached had advertising expenditures been made in proportion to the strength of seasonal demand. In the latter case, 30 percent of the advertising budget would be allocated in the first quarter of the year and only 20 percent in the third quarter.

Applying these calculations to the New York metropolitan area, we discover that the value of information—the value of the additional annual milk sales—is about \$4 million. This figure corresponds to a 9-percent increase in the profit to producers.

You might think that more information is always a good thing. As the following example shows, however, that is not always the case.

EXAMPLE 5.5 Doctors, Patients, and the Value of Information

Suppose you were seriously ill and required major surgery. Assuming you wanted to get the best care possible, how would you go about choosing a surgeon and a hospital to provide that care? Many people would ask their friends or their primary-care physician for a recommendation. Although this might be helpful, a truly informed decision would probably require more detailed information. For example, how successful has a recommended surgeon and her affiliated hospital been in performing the particular operation that you need? How many of her patients have died or had serious complications from the operation, and how do

[8] This example is based on Henry Kinnucan and Olan D. Forker, "Seasonality in the Consumer Response to Milk Advertising with Implications for Milk Promotion Policy," *American Journal of Agricultural Economics* 68 (1986): 562–71.

these numbers compare with those for other surgeons and hospitals? This kind of information is likely to be difficult or impossible for most patients to obtain. Would patients be better off if detailed information about the performance records of doctors and hospitals were readily available?

Not necessarily. More information is often, but not always, better. Interestingly in this case, access to performance information could actually lead to worse health outcomes. Why? Because access to such information would create two different incentives that would affect the behavior of both doctors and patients. First, it would allow patients to choose doctors with better performance records, which creates an incentive for doctors to perform better. That is a good thing. But second, it would encourage doctors to limit their practices to patients who are in relatively good health. The reason is that very old or very sick patients are more likely to have complications or die as a result of treatment; doctors who treat such patients are likely to have worse performance records (other factors being equal). To the extent that doctors would be judged according to performance, they would have an incentive to avoid treating very old or sick patients. As a result, such patients would find it difficult or impossible to obtain treatment.

Whether more information is better depends on which effect dominates—the ability of patients to make more informed choices versus the incentive for doctors to avoid very sick patients. In a recent study, economists examined the impact of the mandatory "report cards" introduced in New York and Pennsylvania in the early 1990s to evaluate outcomes of coronary bypass surgeries.[9] They analyzed hospital choices and outcomes for all elderly heart attack patients and patients receiving coronary bypass surgery in the United States from 1987 through 1994. By comparing trends in New York and Pennsylvania to the trends in other states, they could determine the effect of the increased information made possible by the availability of report cards. They found that although report cards improved matching of patients with hospitals and doctors, they also caused a shift in treatment from sicker patients towards healthier ones. Overall, this led to worse outcomes, especially among sicker patients. Thus the study concluded that report cards reduced welfare.

More information often improves welfare because it allows people to reduce risk and to take actions that might reduce the effect of bad outcomes. However, as this example makes clear, information can cause people to change their behavior in undesirable ways. We will discuss this issue further in Chapter 17.

*5.4 THE DEMAND FOR RISKY ASSETS

Most people are risk averse. Given a choice, they prefer fixed monthly incomes to those which, though equally large on average, fluctuate randomly from month to month. Yet many of these same people will invest all or part of their savings in stocks, bonds, and other assets that carry some risk. Why do risk-averse people invest in the stock market and thereby risk losing part or all of

[9]David Dranove, Daniel Kessler, Mark McClennan, and Mark Satterthwaite, "Is More Information Better? The Effects of 'Report Cards' on Health Care Providers," *Journal of Political Economy* 3 (June 2003).

their investments?[10] How do people decide how much risk to bear when making investments and planning for the future? To answer these questions, we must examine the demand for risky assets.

Assets

An **asset** is *something that provides a flow of money or services to its owner*. A home, an apartment building, a savings account, or shares of General Motors stock are all assets. A home, for example, provides a flow of housing services to its owner, and, if the owner did not wish to live there, could be rented out, thereby providing a monetary flow. Likewise, apartments can be rented out, providing a flow of rental income to the owner of the building. A savings account pays interest (usually every day or every month), which is usually reinvested in the account.

• **asset** Something that provides a flow of money or services to its owner.

The monetary flow that one receives from asset ownership can take the form of an explicit payment, such as the rental income from an apartment building: Every month, the landlord receives rent checks from the tenants. Another form of explicit payment is the dividend on shares of common stock: Every three months, the owner of a share of General Motors stock receives a quarterly dividend payment.

But sometimes the monetary flow from ownership of an asset is implicit: It takes the form of an increase or decrease in the price or value of the asset. An increase in the value of an asset is a *capital gain;* a decrease is a *capital loss*. For example, as the population of a city grows, the value of an apartment building may increase. The owner of the building will then earn a capital gain beyond the rental income. The capital gain is *unrealized* until the building is sold because no money is actually received until then. There is, however, an implicit monetary flow because the building *could* be sold at any time. The monetary flow from owning General Motors stock is also partly implicit. The price of the stock changes from day to day, and each time it does, owners gain or lose.

Risky and Riskless Assets

A **risky asset** provides *a monetary flow that is at least in part random*. In other words, the monetary flow is not known with certainty in advance. A share of General Motors stock is an obvious example of a risky asset: You cannot know whether the price of the stock will rise or fall over time, nor can you even be sure that the company will continue to pay the same (or any) dividend per share. Although people often associate risk with the stock market, most other assets are also risky.

• **risky asset** Asset that provides an uncertain flow of money or services to its owner.

An apartment building is one example. You cannot know how much land values will rise or fall, whether the building will be fully rented all the time, or even whether the tenants will pay their rents promptly. Corporate bonds are another example—the issuing corporation could go bankrupt and fail to pay bond owners their interest and principal. Even long-term U.S. government bonds that mature in 10 or 20 years are risky. Although it is highly unlikely that the federal government will go bankrupt, the rate of inflation could unexpectedly increase and make future interest payments and the eventual repayment of principal worth less in real terms, thereby reducing the value of the bonds.

[10]Most Americans have at least some money invested in stocks or other risky assets, though often indirectly. For example, many people who hold full-time jobs have shares in pension funds underwritten in part by their own salary contributions and in part by employer contributions. Usually such funds are partly invested in the stock market.

• **riskless (or risk-free) asset** Asset that provides a flow of money or services that is known with certainty.

In contrast, a **riskless (or risk-free) asset** pays a monetary flow that is known with certainty. Short-term U.S. government bonds—called Treasury bills—are riskless, or almost riskless. Because they mature in a few months, there is very little risk from an unexpected increase in the rate of inflation. You can also be reasonably confident that the U.S. government will not default on the bond (i.e., refuse to pay back the holder when the bond comes due). Other examples of riskless or almost riskless assets include passbook savings accounts and short-term certificates of deposit.

Asset Returns

• **return** Total monetary flow of an asset as a fraction of its price.

People buy and hold assets because of the monetary flows they provide. To compare assets with each other, it helps to think of this monetary flow relative to an asset's price or value. The **return** on an asset is *the total monetary flow it yields—including capital gains or losses—as a fraction of its price*. For example, a bond worth \$1000 today that pays out \$100 this year (and every year) has a return of 10 percent.[11] If an apartment building was worth \$10 million last year, increased in value to \$11 million this year, and also provided rental income (after expenses) of \$0.5 million, it would have yielded a return of 15 percent over the past year. If a share of General Motors stock was worth \$80 at the beginning of the year, fell to \$72 by the end of the year, and paid a dividend of \$4, it will have yielded a return of –5 percent (the dividend yield of 5 percent less the capital loss of 10 percent).

• **real return** Simple (or nominal) return on an asset, less the rate of inflation.

When people invest their savings in stocks, bonds, land, or other assets, they usually hope to earn a return that exceeds the rate of inflation. Thus, by delaying consumption, they can buy more in the future than they can by spending all their income now. Consequently, we often express the return on an asset in *real*—i.e., *inflation-adjusted*—terms. The **real return** on an asset is its simple (or nominal) return *less* the rate of inflation. For example, with an annual inflation rate of 5 percent, our bond, apartment building, and share of GM stock have yielded real returns of 5 percent, 10 percent, and –10 percent, respectively.

• **expected return** Return that an asset should earn on average.

• **actual return** Return that an asset earns.

Expected versus Actual Returns Because most assets are risky, an investor cannot know in advance what returns they will yield over the coming year. For example, our apartment building might have depreciated in value instead of appreciating, and the price of GM stock might have risen instead of fallen. However, we can still compare assets by looking at their expected returns. The **expected return** on an asset is *the expected value of its return*, i.e., the return that it should earn on average. In some years, an asset's **actual return** may be much higher than its expected return and in some years much lower. Over a long period, however, the average return should be close to the expected return.

Different assets have different expected returns. Table 5.8, for example, shows that while the expected real return of a U.S. Treasury bill has been less than 1 percent, the expected real return on a group of representative stocks on the New York Stock Exchange has been more than 9 percent.[12] Why would anyone

[11]The price of a bond often changes during the course of a year. If the bond appreciates (or depreciates) in value during the year, its return will be greater (or less) than 10 percent. In addition, the definition of *return* given above should not be confused with the "internal rate of return," which is sometimes used to compare monetary flows occurring over a period of time. We discuss other return measures in Chapter 15, when we deal with present discounted values.

[12]For some stocks, the expected return is higher, and for some it is lower. Stocks of smaller companies (e.g., some of those traded on the NASDAQ) have higher expected rates of return—and higher return standard deviations.

TABLE 5.8 Investments—Risk and Return (1926–2006*)

	Average Rate of Return (%)	Average Real Rate of Return (%)	Risk (Standard Deviation, %)
Common stocks (S&P 500)	12.3	9.2	20.1
Long-term corporate bonds	6.2	3.1	8.5
U.S. Treasury bills	3.8	0.7	3.1

*Source: *Stocks, Bonds, Bills, and Inflation: 2007 Yearbook*, Morningstar, Inc.

buy a Treasury bill when the expected return on stocks is so much higher? Because the demand for an asset depends not just on its expected return, but also on its *risk*: Although stocks have a higher expected return than Treasury bills, they also carry much more risk. One measure of risk, the standard deviation of the real annual return, is equal to 20.1 percent for common stocks, 8.5 percent for corporate bonds, and only 3.1 percent for U.S. Treasury bills.

The numbers in Table 5.8 suggest that the higher the expected return on an investment, the greater the risk involved. Assuming that one's investments are well diversified, this is indeed the case.[13] As a result, the risk-averse investor must balance expected return against risk. We examine this trade-off in more detail in the next section.

The Trade-Off Between Risk and Return

Suppose a woman wants to invest her savings in two assets—Treasury bills, which are almost risk free, and a representative group of stocks. She must decide how much to invest in each asset. She might, for instance, invest only in Treasury bills, only in stocks, or in some combination of the two. As we will see, this problem is analogous to the consumer's problem of allocating a budget between purchases of food and clothing.

Let's denote the risk-free return on the Treasury bill by R_f. Because the return is risk free, the expected and actual returns are the same. In addition, let the *expected* return from investing in the stock market be R_m and the actual return be r_m. The actual return is risky. At the time of the investment decision, we know the set of possible outcomes and the likelihood of each, but we do not know what particular outcome will occur. The risky asset will have a higher expected return than the risk-free asset ($R_m > R_f$). Otherwise, risk-averse investors would buy only Treasury bills and no stocks would be sold.

The Investment Portfolio To determine how much money the investor should put in each asset, let's set b equal to the fraction of her savings placed in the stock market and $(1 - b)$ the fraction used to purchase Treasury bills. The

[13]It is *nondiversifiable* risk that matters. An individual stock may be very risky but still have a low expected return because most of the risk could be diversified away by holding a large number of such stocks. *Nondiversifiable risk*, which arises from the fact that individual stock prices are correlated with the overall stock market, is the risk that remains even if one holds a diversified portfolio of stocks. We discuss this point in detail in the context of the *capital asset pricing model* in Chapter 15.

expected return on her total portfolio, R_p, is a weighted average of the expected return on the two assets:[14]

$$R_p = bR_m + (1 - b)R_f \tag{5.1}$$

Suppose, for example, that Treasury bills pay 4 percent ($R_f = .04$), the stock market's expected return is 12 percent ($R_m = .12$), and $b = 1/2$. Then $R_p = 8$ percent. How risky is this portfolio? One measure of riskiness is the standard deviation of its return. We will denote the *standard deviation* of the risky stock market investment by σ_m. With some algebra, we can show that the *standard deviation of the portfolio*, σ_p (with one risky and one risk-free asset) is the fraction of the portfolio invested in the risky asset times the standard deviation of that asset:[15]

$$\sigma_p = b\sigma_m \tag{5.2}$$

The Investor's Choice Problem

In §3.2 we explain how a budget line is determined from an individual's income and the prices of the available goods.

We have still not determined how the investor should choose this fraction b. To do so, we must first show that she faces a risk-return trade-off analogous to a consumer's budget line. To identify this trade-off, note that equation (5.1) for the expected return on the portfolio can be rewritten as

$$R_p = R_f + b(R_m - R_f)$$

Now, from equation (5.2) we see that $b = \sigma_p/\sigma_m$, so that

$$\boxed{R_p = R_f + \frac{(R_m - R_f)}{\sigma_m}\sigma_p} \tag{5.3}$$

• **Price of risk** Extra risk that an investor must incur to enjoy a higher expected return.

Risk and the Budget Line This equation is a *budget line* because it describes the trade-off between risk (σ_p) and expected return (R_p). Note that it is the equation for a straight line: Because R_m, R_f, and σ_m are constants, the slope $(R_m - R_f)/\sigma_m$ is a constant, as is the intercept, R_f. The equation says that *the expected return on the portfolio R_p increases as the standard deviation of that return σ_p increases*. We call the slope of this budget line, $(R_m - R_f)/\sigma_m$, the **price of risk,** because it tells us how much extra risk an investor must incur to enjoy a higher expected return.

The budget line is drawn in Figure 5.6. If our investor wants no risk, she can invest all her funds in Treasury bills ($b = 0$) and earn an expected return R_f. To receive a higher expected return, she must incur some risk. For example, she could invest all her funds in stocks ($b = 1$), earning an expected return R_m but incurring a standard deviation σ_m. Or she might invest some fraction of her

[14]The expected value of the sum of two variables is the sum of the expected values. Therefore

$$R_p = E[br_m] + E[(1-b)R_f] = bE[r_m] + (1-b)R_f = bR_m + (1-b)R_f$$

[15]To see why, we observe from footnote 4 that we can write the variance of the portfolio return as

$$\sigma_p^2 = E[br_m + (1-b)R_f - R_p]^2$$

Substituting equation (5.1) for the expected return on the portfolio, R_p, we have

$$\sigma_p^2 = E[br_m + (1-b)R_f - bR_m - (1-b)R_f]^2 = E[b(r_m - R_m)]^2 = b^2\sigma_m^2$$

Because the standard deviation of a random variable is the square root of its variance, $\sigma_p = b\sigma_m$.

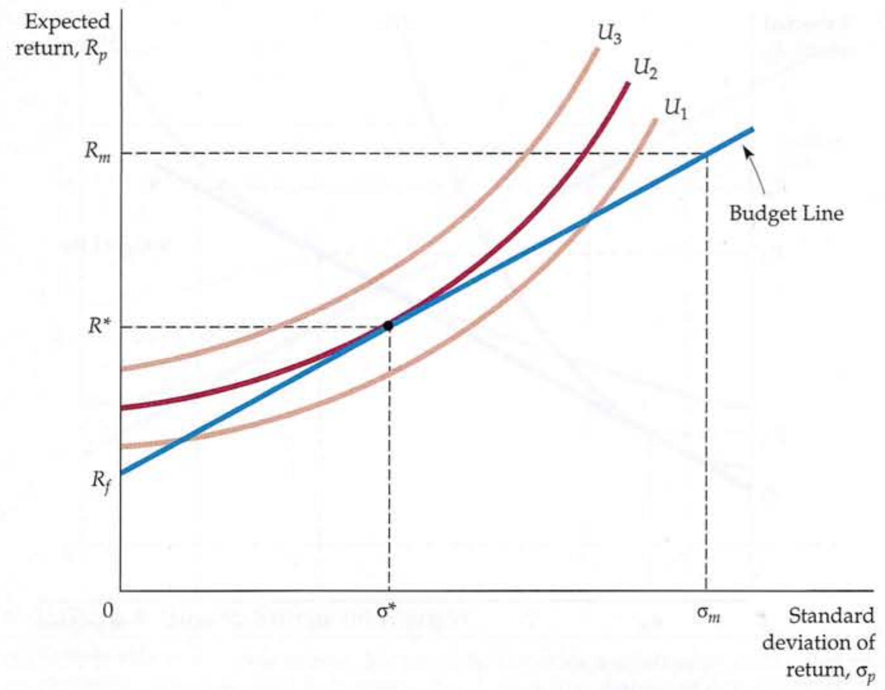

FIGURE 5.6 Choosing Between Risk and Return

An investor is dividing her funds between two assets—Treasury bills, which are risk free, and stocks. The budget line describes the trade-off between the expected return and its riskiness, as measured by the standard deviation of the return. The slope of the budget line is $(R_m - R_f)/\sigma_m$, which is the price of risk. Three indifference curves are drawn, each showing combinations of risk and return that leave an investor equally satisfied. The curves are upward-sloping because a risk-averse investor will require a higher expected return if she is to bear a greater amount of risk. The utility-maximizing investment portfolio is at the point where indifference curve U_2 is tangent to the budget line.

funds in each type of asset, earning an expected return somewhere between R_f and R_m and facing a standard deviation less than σ_m but greater than zero.

Risk and Indifference Curves Figure 5.6 also shows the solution to the investor's problem. Three indifference curves are drawn in the figure. Each curve describes combinations of risk and return that leave the investor equally satisfied. The curves are upward-sloping because risk is undesirable. Thus, with a greater amount of risk, it takes a greater expected return to make the investor equally well-off. Curve U_3 yields the greatest amount of satisfaction and U_1 the least amount: For a given amount of risk, the investor earns a higher expected return on U_3 than on U_2 and a higher expected return on U_2 than on U_1.

Of the three indifference curves, the investor would prefer to be on U_3. This position, however, is not feasible, because U_3 does not touch the budget line. Curve U_1 is feasible, but the investor can do better. Like the consumer choosing quantities of food and clothing, our investor does best by choosing a combination of risk and return at the point where an indifference curve (in this case U_2) is tangent to the budget line. At that point, the investor's return has an expected value R^* and a standard deviation σ^*.

Naturally, people differ in their attitudes toward risk. This fact is illustrated in Figure 5.7, which shows how two different investors choose their portfolios.

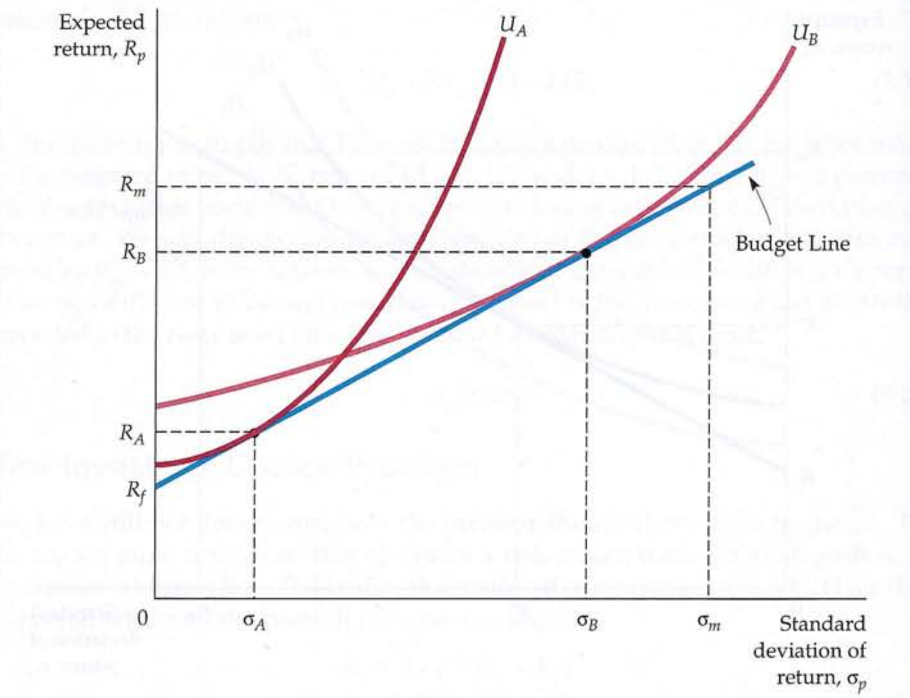

FIGURE 5.7 The Choices of Two Different Investors

Investor *A* is highly risk averse. Because his portfolio will consist mostly of the risk-free asset, his expected return R_A will be only slightly greater than the risk-free return. His risk σ_A, however, will be small. Investor *B* is less risk averse. She will invest a large fraction of her funds in stocks. Although the expected return on her portfolio R_B will be larger, it will also be riskier.

Investor *A* is quite risk averse. Because his indifference curve U_A is tangent to the budget line at a point of low risk, he will invest almost all of his funds in Treasury bills and earn an expected return R_A just slightly larger than the risk-free return R_f. Investor *B* is less risk averse. She will invest most of her funds in stocks, and while the return on her portfolio will have a higher expected value R_B, it will also have a higher standard deviation σ_B.

If Investor *B* has a sufficiently low level of risk aversion, she might buy stocks on *margin*: that is, she would borrow money from a brokerage firm in order to invest more than she actually owns in the stock market. In effect, a person who buys stocks on margin holds a portfolio with more than 100 percent of the portfolio's value invested in stocks. This situation is illustrated in Figure 5.8, which shows indifference curves for two investors. Investor *A*, who is relatively risk-averse, invests about half of his funds in stocks. Investor *B*, however, has an indifference curve that is relatively flat and tangent with the budget line at a point where the expected return on the portfolio exceeds the expected return on the stock market. In order to hold this portfolio, the investor must borrow money because she wants to invest *more* than 100 percent of her wealth in the stock market. Buying stocks on margin in this way is a form of *leverage*: the investor increases her expected return above that for the overall stock market, but at the cost of increased risk.

In Chapters 3 and 4, we simplified the problem of consumer choice by assuming that the consumer had only two goods from which to choose—food and

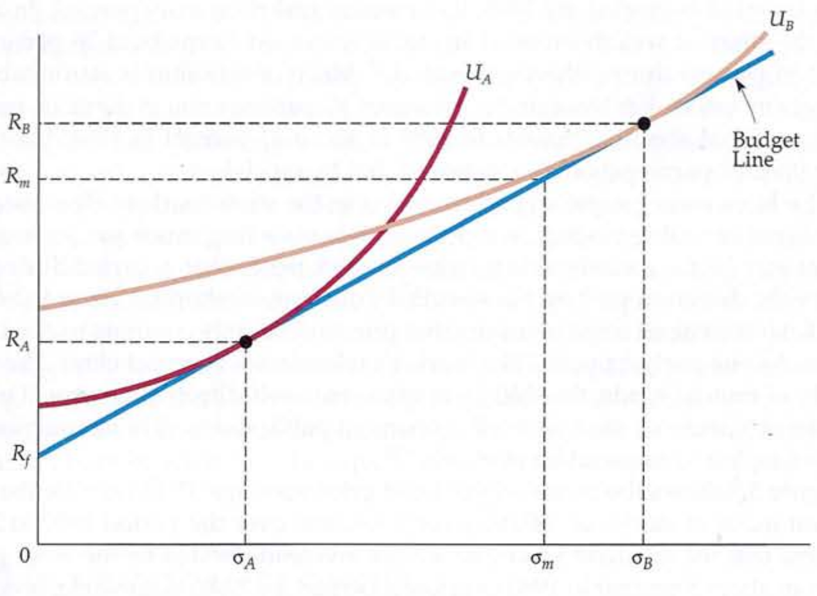

FIGURE 5.8 Buying Stocks on Margin

Because Investor *A* is risk averse, his portfolio contains a mixture of stocks and risk-free Treasury bills. Investor *B*, however, has a very low degree of risk aversion. Her indifference curve, U_B, is tangent to the budget line at a point where the expected return and standard deviation for her portfolio exceed those for the stock market overall. This implies that she would like to invest *more* than 100 percent of her wealth in the stock market. She does so by buying stocks *on margin*—i.e., by borrowing from a brokerage firm to help finance her investment.

clothing. In the same spirit, we have simplified the investor's choice by limiting it to Treasury bills and stocks. The basic principles, however, would be the same if we had more assets (e.g., corporate bonds, land, and different types of stocks). Every investor faces a trade-off between risk and return.[16] The degree of extra risk that each is willing to bear in order to earn a higher expected return depends on how risk averse he or she is. Less risk-averse investors tend to include a larger fraction of risky assets in their portfolios.

EXAMPLE 5.6 Investing in the Stock Market

The 1990s witnessed a shift in the investing behavior of Americans. First, many people started investing in the stock market for the first time. In 1989, about 32 percent of families in the United States had part of their wealth invested in the stock market, either directly (by owning individual stocks) or indirectly (through mutual funds or pension

[16]As mentioned earlier, what matters is nondiversifiable risk, because investors can eliminate diversifiable risk by holding many different stocks (e.g., via mutual funds). We discuss diversifiable versus nondiversifiable risk in Chapter 15.

plans invested in stocks). By 1998, that fraction had risen to 49 percent. In addition, the share of wealth invested in stocks increased from about 26 percent to about 54 percent during the same period.[17] Much of this shift is attributable to younger investors. For those under the age of 35, participation in the stock market increased from about 22 percent in 1989 to about 41 percent in 1998. For those older than 35, participation also increased, but by much less.

Why have more people started investing in the stock market? One reason is the advent of online trading, which has made investing much easier. Another reason may be the considerable increase in stock prices that occurred during the late 1990s, driven in part by the so-called "dot com euphoria." These increases may have convinced some investors that prices could only continue to rise in the future. As one analyst put it, "The market's relentless seven-year climb, the popularity of mutual funds, the shift by employers to self-directed retirement plans, and the avalanche of do-it-yourself investment publications all have combined to create a nation of financial know-it-alls."[18]

Figure 5.9 shows the dividend yield and price/earnings (P/E) ratio for the S&P 500 (an index of stocks of 500 large corporations) over the period 1980 to 2007. Observe that the dividend yield (the annual dividend divided by the stock price) fell from about 5 percent in 1980 to below 2 percent by 2000. Meanwhile, however, the price/earnings ratio (the share price divided by annual earnings per share)

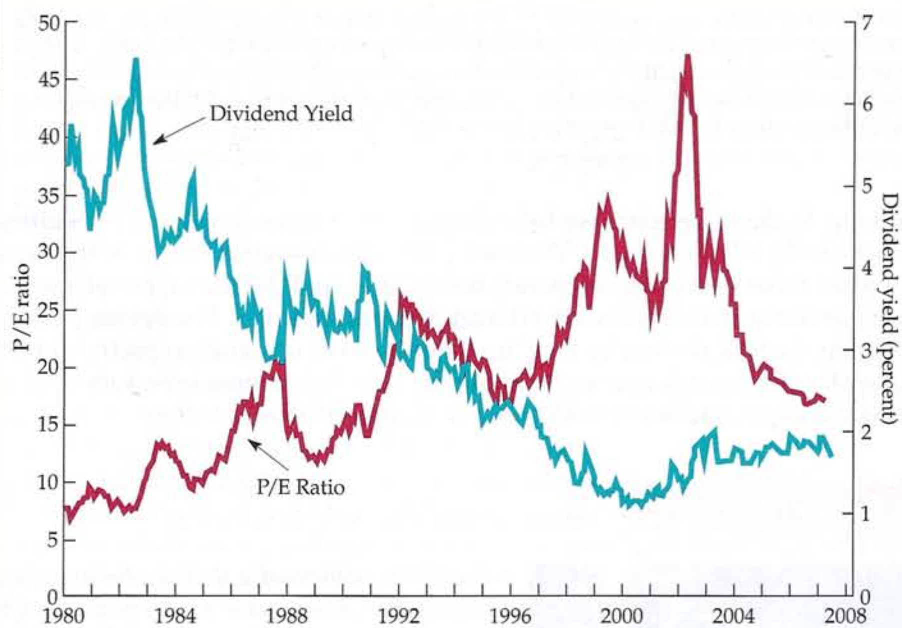

FIGURE 5.9 Dividend Yield and P/E Ratio for S&P 500

The dividend yield for the S&P 500 (the annual dividend divided by the stock price) has fallen dramatically, while the price/earnings ratio (the stock price divided by the annual earnings-per-share) rose from 1980 to 2002 and then dropped.

[17]Data are from the *Federal Reserve Bulletin*, January 2000.

[18]"We're All Bulls Here: Strong Market Makes Everybody an Expert," *Wall Street Journal*, September 12, 1997.

increased from about 8 in 1980 to over 40 in 2002, before falling to around 20 between 2005 and 2007. In retrospect, the increase in the P/E ratio through 2002 could only have occurred if investors believed that corporate profits would continue to grow rapidly in the coming decade. This suggests that in the late 1990s, many investors had a low degree of risk aversion, were quite optimistic about the economy, or both. Alternatively, some economists have argued that the run-up of stock prices during the 1990s was the result of "herd behavior," in which investors rushed to get into the market after hearing of the successful experiences of others.[19]

The psychological motivations that explain herd behavior can help to explain stock market bubbles. However, they go far beyond the stock market. They also apply to the behavior of consumers and firm managers in a wide variety of settings. Such behavior cannot always be captured by the simplified assumptions that we have made up to this point about consumer choice. In the next section, we will discuss these aspects of behavior in detail, and we will see how the traditional models of Chapters 3 and 4 can be expanded to help us understand this behavior.

5.5 BEHAVIORAL ECONOMICS

Recall that the basic theory of consumer demand is based on three assumptions: (1) consumers have clear preferences for some goods over others; (2) consumers face budget constraints; and (3) given their preferences, limited incomes, and the prices of different goods, consumers choose to buy combinations of goods that maximize their satisfaction (or utility). These assumptions, however, are not always realistic: Preferences are not always clear or might vary depending on the context in which choices are made, and consumer choices are not always utility-maximizing.

Perhaps our understanding of consumer demand (as well as the decisions of firms) would be improved if we incorporated more realistic and detailed assumptions regarding human behavior. This has been the objective of the newly flourishing field of *behavioral economics*, which has broadened and enriched the study of microeconomics.[20] We introduce this topic by highlighting some examples of consumer behavior that cannot be easily explained with the basic utility-maximizing assumptions that we have relied on so far:

- There has just been a big snowstorm, so you stop at the hardware store to buy a snow shovel. You had expected to pay $20 for the shovel—the price that the store normally charges. However, you find that the store has suddenly raised the price to $40. Although you would expect a price increase because of the storm, you feel that a doubling of the price is unfair and that the store is trying to take advantage of you. Out of spite, you do not buy the shovel.[21]
- Tired of being snowed in at home you decide to take a vacation in the country. On the way, you stop at a highway restaurant for lunch. Even though you are

[19]See, for example, Robert Shiller, *Irrational Exuberance*, Princeton University Press, 2000.

[20]For more detailed discussion of the material presented in this section, see Stefano DellaVigna, "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature* (forthcoming); Colin Camerer and George Loewenstein, "Behavioral Economics: Past, Present, Future," in Colin Camerer, George Loewenstein, and Matthew Rabin (eds.), *Advances in Behavioral Economics*, Princeton University Press, 2003.

[21]This example is based on Daniel Kahneman, Jack Knetsch, and Richard Thaler, "Fairness as a Constraint on Profit Seeking: Entitlements in the Market," *American Economic Review* 76 (September 1986): 728–741.

unlikely to return to that restaurant, you believe that it is fair and appropriate to leave a 15-percent tip in appreciation of the good service that you received.

- You buy this textbook from an Internet bookseller because the price is lower than the price at your local bookstore. However, you ignore the shipping cost when comparing prices.

Each of these examples illustrates plausible behavior that cannot be explained by a model based solely on the basic assumptions described in Chapters 3 and 4. Instead, we need to draw on insights from psychology and sociology to augment our basic assumptions about consumer behavior. These insights will enable us to account for more complex consumer preferences, for the use of simple rules in decision-making, and for the difficulty that people often have in understanding the laws of probability.

More Complex Preferences

The standard model of consumer behavior assumes that consumers place unique values on the goods and services that they purchase. However, psychologists have found that perceived value depends on the circumstances. Consider, for example, apartment prices in Pittsburgh and San Francisco. In Pittsburgh, the median monthly rent in 2006 for a two-bedroom apartment was about $650, while in San Francisco the rent for a similar apartment was $2,125. For someone accustomed to San Francisco housing prices, Pittsburgh might seem like a bargain. On the other hand, someone moving from Pittsburgh to San Francisco might feel "gouged"—thinking it unfair for housing to cost that much.[22]

• **reference point** The point from which an individual makes a consumption decision.

In this example, the **reference point**—the point from which the individual makes the consumption decision—is clearly different for long-time residents of San Francisco and Pittsburgh. Reference points can develop for many reasons: our past consumption, our experience in a market, our expectations about how prices should behave, and even the context in which we consume a good. Reference points can strongly affect the way people approach economic decisions.

• **endowment effect** Tendency of individuals to value an item more when they own it than when they do not.

A well-known example of a reference point is the **endowment effect**—the fact that individuals tend to value an item more when they happen to own it than when they do not. One way to think about this effect is to consider the gap between the price that a person is willing to pay for a good and the price at which she is willing to sell the same good to someone else. Our basic theory of consumer behavior says that this price should be the same, but many experiments suggest that is not what happens in practice.[23]

In one classroom experiment, half of the students chosen at random were given a free coffee mug with a market value of $5; the other half got nothing.[24] Students with the mug were asked the price at which they would sell it back to the professor; the second group was asked the minimum amount of money that they would accept in lieu of a mug. The decision faced by both groups is similar but their reference points are different. For the first group, whose reference point was possession of a mug, the average selling price was $7. For the second

[22]This example is based on Uri Simonsohn and George Loewenstein, "Mistake #37: The Effects of Previously Encountered Prices on Current Housing Demand," *The Economic Journal* 116 (January 2006): 175–199.

[23]Experimental work such as this has been important to the development of behavioral economics. It is for this reason that the 2002 Nobel Prize in economics was shared by Vernon Smith, who did much of the pioneering work in the use of experiments to test economic theories.

[24]Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler, "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy 98*, (December 1990): 1925–48.

group, which did not have a mug, the average amount desired in lieu of a mug was $3.50. This gap in prices shows that giving up the mug was perceived to be a greater "loss" to those who had one than the "gain" from obtaining a mug for those without one. Such a result, aptly called **loss aversion**, has been apparent in many experimental studies.

• **loss aversion** Tendency for individuals to prefer avoiding losses over acquiring gains.

As another example of loss aversion, people are sometimes hesitant to sell stocks at a loss, even if they could invest the proceeds in other stocks that they think are better investments. Why? Because the original price paid for the stock—which turned out to be too high given the realities of the market—acts as a reference point, and people are averse to losses. (A $1000 loss on an investment seems to "hurt" more than the perceived benefit from a $1000 gain.) While there are a variety of circumstances in which endowment effects arise, we now know that these effects tend to disappear as consumers gain relevant experience. We would not expect to see stockbrokers or other investment professionals exhibit the loss aversion described above.[25]

Many people do things because they think it is appropriate to do so, even though there is no financial or other material benefit. Examples include charitable giving, volunteering time, or tipping in a restaurant. And, as in our examples of renting an apartment, buying a snow shovel, and tipping pointed out, there are occasions in which consumers' views about fairness also affect their behavior.

Our basic consumer theory does not appear to account for fairness, at least at first glance. The so-called *ultimatum game* illustrates this supposition. Imagine that, under the following rules, you are offered a chance to divide 100 one-dollar bills with a stranger whom you will never meet again: You first propose a division of the money between you and the stranger. The stranger will respond by either accepting or rejecting your proposal. If he accepts, you each get the share that you proposed. If he rejects, you both get nothing. What should you do?

Because more money means more utility, our basic theory provides a clear answer to this question. You should propose that you get $99 while the other person gets only $1. Moreover, the responder should be happy to accept this proposal, because $1 is more than he had before and more than he would get if he rejected your offer (in both cases zero). This is a beneficial deal for both of you.

Yet most people facing this choice hesitate to make such an offer because they think it unfair, and many "strangers" would reject the offer. Why? The stranger might believe that because you both received the windfall opportunity to divide $100, a simple and fair division would be 50/50 or something close to that. Maybe the stranger will turn down the $1 offer to teach you that greediness is not appropriate behavior. Indeed, if you believe that the stranger will feel this way, it will be rational for you to offer a greater amount. In fact, when this game is played experimentally, typical sharing proposals range between 67/33 and 50/50, and such offers are normally accepted.

The ultimatum game shows how fairness can affect economic decisions. Not surprisingly, fairness concerns can also affect negotiations between firms and their workers. A firm may offer a higher wage to employees because the managers believe that workers deserve a comfortable standard of living or because they want to foster a pleasant working environment. Moreover, workers who do not get a wage that they feel is fair may not put much effort into their work.[26]

[25]John A. List, "Does Market Experience Eliminate Market Anomalies?" *Quarterly Journal of Economics* 118 (January 2003): 41–71.

[26]For a general discussion of behavioral economics and the theory of wages and employment, see George Akerlof, "Behavioral Macroeconomics and Macroeconomic Behavior," *American Economic Review* 92 (June 2002): 411–33.

(In Section 17.6, we will see that paying workers higher-than-market wages can also be explained by the "efficiency wage theory" of labor markets, in which fairness concerns do not apply.) Fairness also affects the ways in which firms set prices and can explain why firms can more easily raise prices in response to higher costs than to increases in demand.[27]

Fortunately, fairness concerns can be taken into account in the basic model of consumer behavior. If individuals moving to San Francisco believe that high apartment rents are unfair, their maximum willingness to pay for rental housing will be reduced. If a sufficient number of individuals feel this way, the resulting reduction in demand will lead to lower rental prices. Similarly, if enough workers do not feel that their wages are fair, there will be a reduction in the supply of labor, and wage rates will increase.

Rules of Thumb and Biases in Decision Making

Many economic (and everyday) decisions can be quite complex, especially if they involve choices about matters in which we have little experience. In such cases, people often resort to rule of thumb or other mental shortcuts to help them make decisions. In the tipping example, you took a mental shortcut when you decided to offer a 15-percent tip. The use of such rules of thumb, however, can introduce a bias into our economic decision making—something that our basic model does not allow.[28]

The mental rules that we use in making decisions frequently depend on both the context in which the decisions are made and the information available. For example, imagine that you just received a solicitation from a new local charity to make a donation. Rather than asking for a gift of any amount, the charity asks you to choose: \$20, \$50, \$100, \$250, or "other." The purpose of these suggestions is to induce you to anchor your final donation. **Anchoring** refers to the impact that a suggested (perhaps unrelated) piece of information may have on your final decision. Rather than trying to decide precisely how much to donate—say \$44.52—and not wanting to appear miserly, one might simply write a check for the next higher category—\$50. Another individual wishing to make only a token donation of \$10 might choose the lowest stated amount, \$20. In both cases, anchoring can bias individual choices toward larger donations.

• **anchoring** Tendency to rely heavily on one prior (suggested) piece of information when making a decision.

A common way to economize on the effort involved in making decisions is to ignore seemingly unimportant pieces of information. For example, goods purchased over the Internet often involve shipping costs. Although small, these costs should be included as part of the good's final price when making a consumption decision. However, a recent study has shown that shipping costs are typically ignored by many consumers when deciding to buy things online. Their decisions are biased because they view the price of goods to be lower than they really are.[29]

Whereas depending on rules of thumb can introduce biases in decision making, it is important to understand that they do serve a useful purpose. Frequently, rules of thumb help to save time and effort and result in only small biases. Thus, they should not be dismissed outright.

[27]See, for example, Julio J. Rotemberg, "Fair Pricing," NBER Working Paper No. W10915, 2004.

[28]For an introduction to this topic see Amos Tversky and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases," *Science* 185 (September 1974): 1124–31.

[29]Tankim Hossain and John Morgan, ". . . Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay," *Advances in Economic Analysis* & Policy 6: 2 (2006).

Probabilities and Uncertainty

An important part of decision making under uncertainty is the calculation of expected utility, which requires two pieces of information: a utility value for each outcome (from the utility function) and the probability of each outcome. Although the expected-utility approach may seem simple, in practice we tend to have difficulty making such calculations. In part, this is because many of us lack a basic understanding of probability.

People are sometimes prone to a bias called the *law of small numbers*: They tend to overstate the probability that certain events will occur when faced with relatively little information from recent memory. For example, many people tend to overstate the likelihood that they or someone they know will die in a plane crash or win the lottery. Recall the roulette player who bets on black after seeing red come up three times in a row: He has ignored the laws of probability.

Research has shown that investors in the stock market are often subject to a small-numbers bias, believing that high returns over the past few years are likely to be followed by more high returns over the next few years—thereby contributing to the kind of "herd behavior" that we discussed in the previous section. In this case, investors assess the likely payoff from investing by observing the market over a short period of time. In fact, one would have to study stock market returns for many decades in order to estimate accurately the expected return on equity investments. Similarly when people assess the likelihood that housing prices will rise based on several years of data, the resulting misperceptions can result in housing price bubbles.[30]

Although individuals may have some understanding of true probabilities (as when flipping a coin), complications arise when probabilities are unknown. For instance, few people have an idea about the probability that they or a friend will be in a car or airplane accident. In such cases, we form subjective probability assessments about such events. Our estimation of subjective probabilities may be close to true probabilities, but often they are not.

Forming subjective probabilities is not always an easy task and people are generally prone to several biases in the process. For instance, when evaluating the likelihood of an event, the context in which the evaluation is made can be very important. If a tragedy such as a plane crash has occurred recently, many people will tend to overestimate the probability of it happening to them. Likewise, when a probability for a particular event is very, very small, many people simply ignore that possibility in their decision making.

Summing Up

Where does this leave us? Should we dispense with the traditional consumer theory discussed in Chapters 3 and 4? Not at all. In fact, the basic theory that we learned up to now works quite well in many situations. It helps us to understand and evaluate the characteristics of consumer demand and to predict the impact on demand of changes in prices or incomes. Although it does not explain all consumer decisions, it sheds light on many of them. The developing field of behavioral economics tries to explain and to elaborate on those situations that are not well explained by the basic consumer model.

---

[30]See Charles Himmelberg, Christopher Mayer, and Todd Sinai, "Assessing High House Prices: Bubbles, Fundamentals and Misperceptions," *Journal of Economic Perspectives* 19 (Fall 2005).

EXAMPLE 5.7 New York City Taxicab Drivers

Most cab drivers rent their taxicabs for a fixed daily fee from a company that owns a fleet of cars. They can then choose to drive the cab as little or as much as they want during a 12-hour period. As with many services, business is highly variable from day to day, depending on the weather, subway breakdowns, holidays, and so on. How do cabdrivers respond to these variations, many of which are largely unpredictable?

In many cities, taxicab rates are fixed by regulation and do not change from day to day. However, on busy days drivers can earn a higher income because they do not have to spend as much time searching for riders. Traditional economic theory would predict that drivers will work longer hours on busy days than on slow days; an extra hour on a busy day might bring in $20, whereas an extra hour on a slow day might yield only $10. Does traditional theory explain the actual behavior of taxicab drivers?

A recent study analyzed actual taxicab trip records obtained from the New York Taxi and Limousine Commission for the spring of 1994.[31] The daily fee to rent a taxi was then $76, and gasoline cost about $15 per day. Surprisingly, the researchers found that most drivers drive *more* hours on slow days and *fewer* hours on busy days. In other words, there is a *negative relationship* between the effective hourly wage and the number of hours worked each day; the higher the wage, the sooner the cabdrivers quit for the day. Behavioral economics can explain this result. Suppose that most taxicab drivers have an income target for each day. That target effectively serves as a reference point. Daily income targeting makes sense from a behavioral perspective. An income target provides a simple decision rule for drivers because they need only keep a record of their fares for the day. A daily target also helps drivers with potential self-control problems; without a target, a driver may choose to quit earlier on many days just to avoid the hassles of the job. The target in the 1994 study appeared to be about $150 per day.

Still other studies challenge this "behavioral" explanation of behavior. A different study, also of New York City cab drivers who rented their taxis, concluded that the traditional economic model does indeed offer important insights into drivers' behavior.[32] The study concluded that daily income had only a small effect on a driver's decision as to when to quit for the day. Rather, the decision to stop appears to be based on the cumulative number of hours already worked that day and not on hitting a specific income target.

What can account for these two seemingly contradictory results? The two studies used different techniques in analyzing and interpreting the taxicab trip records. Although behavioral models, such as those with reference points or targeted goals, often lead to interesting implications for economic theory, the traditional model can indeed go a long way in explaining what we frequently observe.

---

[31]Colin Camerer, Linda Babcock, George Loewenstein, and Richard Thaler, "Labor Supply of New York City Cabdrivers: One Day at a Time," *Quarterly Journal of Economics* (May 1997): 404–41.

[32]Henry S. Farber, "Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers," *Journal of Political Economy* 113 (2005): 46–82.

SUMMARY

1. Consumers and managers frequently make decisions in which there is uncertainty about the future. This uncertainty is characterized by the term *risk*, which applies when each of the possible outcomes and its probability of occurrence is known.
2. Consumers and investors are concerned about the expected value and the variability of uncertain outcomes. The expected value is a measure of the central tendency of the values of risky outcomes. Variability is frequently measured by the standard deviation of outcomes, which is the square root of the probability-weighted average of the squares of the deviation from the expected value of each possible outcome.
3. Facing uncertain choices, consumers maximize their expected utility—an average of the utility associated with each outcome—with the associated probabilities serving as weights.
4. A person who would prefer a certain return of a given amount to a risky investment with the same expected return is risk averse. The maximum amount of money that a risk-averse person would pay to avoid taking a risk is called the *risk premium*. A person who is indifferent between a risky investment and the certain receipt of the expected return on that investment is risk neutral. A risk-loving consumer would prefer a risky investment with a given expected return to the certain receipt of that expected return.
5. Risk can be reduced by (a) diversification, (b) insurance, and (c) additional information.
6. The *law of large numbers* enables insurance companies to provide insurance for which the premiums paid equal the expected value of the losses being insured against. We call such insurance *actuarially fair*.
7. Consumer theory can be applied to decisions to invest in risky assets. The budget line reflects the price of risk, and consumers' indifference curves reflect their attitudes toward risk.
8. Individual behavior sometimes seems unpredictable, even irrational, and contrary to the assumptions that underlie the basic model of consumer choice. The study of behavioral economics enriches consumer theory by accounting for *reference points*, *endowment effects*, *anchoring*, fairness considerations, and deviations from the laws of probability.

QUESTIONS FOR REVIEW

1. What does it mean to say that a person is *risk averse*? Why are some people likely to be risk averse while others are risk lovers?
2. Why is the variance a better measure of variability than the range?
3. George has $5000 to invest in a mutual fund. The expected return on mutual fund *A* is 15 percent and the expected return on mutual fund B is 10 percent. Should George pick mutual fund *A* or fund *B*?
4. What does it mean for consumers to maximize expected utility? Can you think of a case in which a person might *not* maximize expected utility?
5. Why do people often want to insure fully against uncertain situations even when the premium paid exceeds the expected value of the loss being insured against?
6. Why is an insurance company likely to behave as if it were risk neutral even if its managers are risk-averse individuals?
7. When is it worth paying to obtain more information to reduce uncertainty?
8. How does the diversification of an investor's portfolio avoid risk?
9. Why do some investors put a large portion of their portfolios into risky assets while others invest largely in risk-free alternatives? (*Hint*: Do the two investors receive exactly the same return on average? If so, why?)
10. What is an endowment effect? Give an example of such an effect.
11. Jennifer is shopping and sees an attractive shirt. However, the price of $50 is more than she is willing to pay. A few weeks later, she finds the same shirt on sale for $25 and buys it. When a friend offers her $50 for the shirt, she refuses to sell it. Explain Jennifer's behavior.

EXERCISES

1. Consider a lottery with three possible outcomes:
   - $125 will be received with probability .2
   - $100 will be received with probability .3
   - $50 will be received with probability .5

   a. What is the expected value of the lottery?

   b. What is the variance of the outcomes?

   c. What would a risk-neutral person pay to play the lottery?
2. Suppose you have invested in a new computer company whose profitability depends on two factors: (1) whether the U.S. Congress passes a tariff raising the

cost of Japanese computers and (2) whether the U.S. economy grows slowly or quickly. What are the four mutually exclusive states of the world that you should be concerned about?

3. Richard is deciding whether to buy a state lottery ticket. Each ticket costs $1, and the probability of winning payoffs is given as follows:

Probability	Return
.5	$0.00
.25	$1.00
.2	$2.00
.05	$7.50

   a. What is the expected value of Richard's payoff if he buys a lottery ticket? What is the variance?
   b. Richard's nickname is "No-Risk Rick" because he is an extremely risk-averse individual. Would he buy the ticket?
   c. Richard has been given 1000 lottery tickets. Discuss how you would determine the smallest amount for which he would be willing to sell all 1000 tickets.
   d. In the long run, given the price of the lottery tickets and the probability/return table, what do you think the state would do about the lottery?

4. Suppose an investor is concerned about a business choice in which there are three prospects—the probability and returns are given below:

Probability	Return
.4	$100
.3	30
.3	−30

What is the expected value of the uncertain investment? What is the variance?

5. You are an insurance agent who must write a policy for a new client named Sam. His company, Society for Creative Alternatives to Mayonnaise (SCAM), is working on a low-fat, low-cholesterol mayonnaise substitute for the sandwich-condiment industry. The sandwich industry will pay top dollar to the first inventor to patent such a mayonnaise substitute. Sam's SCAM seems like a very risky proposition to you. You have calculated his possible returns table as follows:

Probability	Return	Outcome
.999	−$1,000,000	(he fails)
.001	$1,000,000,000	(he succeeds and sells his formula)

   a. What is the expected return of Sam's project? What is the variance?
   b. What is the most that Sam is willing to pay for insurance? Assume Sam is risk neutral.
   c. Suppose you found out that the Japanese are on the verge of introducing their own mayonnaise substitute next month. Sam does not know this and has just turned down your final offer of $1000 for the insurance. Assume that Sam tells you SCAM is only six months away from perfecting its mayonnaise substitute *and* that you know what you know about the Japanese. Would you raise or lower your policy premium on any subsequent proposal to Sam? Based on his information, would Sam accept?

6. Suppose that Natasha's utility function is given by $u(I) = \sqrt{10I}$, where I represents annual income in thousands of dollars.
   a. Is Natasha risk loving, risk neutral, or risk averse? Explain.
   b. Suppose that Natasha is currently earning an income of $40,000 ($I = 40$) and can earn that income next year with certainty. She is offered a chance to take a new job that offers a .6 probability of earning $44,000 and a .4 probability of earning $33,000. Should she take the new job?
   c. In (b), would Natasha be willing to buy insurance to protect against the variable income associated with the new job? If so, how much would she be willing to pay for that insurance? (*Hint*: What is the risk premium?)

7. Suppose that two investments have the same three payoffs, but the probability associated with each payoff differs, as illustrated in the table below:

Payoff	Probability (Investment A)	Probability (Investment B)
$300	0.10	0.30
$250	0.80	0.40
$200	0.10	0.30

   a. Find the expected return and standard deviation of each investment.
   b. Jill has the utility function $U = 5I$, where I denotes the payoff. Which investment will she choose?
   c. Ken has the utility function $U = 5\sqrt{I}$. Which investment will he choose?
   d. Laura has the utility function $U = 5I^2$. Which investment will she choose?

8. As the owner of a family farm whose wealth is $250,000, you must choose between sitting this season out and investing last year's earnings ($200,000) in a safe money market fund paying 5.0 percent or planting summer corn. Planting costs $200,000, with a six-month time to harvest. If there is rain, planting summer corn will yield $500,000 in revenues at harvest. If there is a drought, planting will yield $50,000 in revenues.

As a third choice, you can purchase AgriCorp drought-resistant summer corn at a cost of $250,000 that will yield $500,000 in revenues at harvest if there is rain, and $350,000 in revenues if there is a drought. You are risk averse, and your preference for family wealth (W) is specified by the relationship $U(W) = \sqrt{W}$. The probability of a summer drought is 0.30, while the probability of summer rain is 0.70.

Which of the three options should you choose? Explain.

9. Draw a utility function over income $u(I)$ that describes a man who is a risk lover when his income is low but risk averse when his income is high. Can you explain why such a utility function might reasonably describe a person's preferences?

10. A city is considering how much to spend to hire people to monitor its parking meters. The following information is available to the city manager:
    - Hiring each meter monitor costs $10,000 per year.
    - With one monitoring person hired, the probability of a driver getting a ticket each time he or she parks illegally is equal to .25.
    - With two monitors, the probability of getting a ticket is .5; with three monitors, the probability is .75; and with four, it's equal to 1.
    - With two monitors hired, the current fine for overtime parking is $20.

    a. Assume first that all drivers are risk neutral. What parking fine would you levy, and how many meter monitors would you hire (1, 2, 3, or 4) to achieve the current level of deterrence against illegal parking at the minimum cost?
    b. Now assume that drivers are highly risk averse. How would your answer to (a) change?
    c. (For discussion) What if drivers could insure themselves against the risk of parking fines? Would it make good public policy to permit such insurance?

11. A moderately risk-averse investor has 50 percent of her portfolio invested in stocks and 50 percent in risk-free Treasury bills. Show how each of the following events will affect the investor's budget line and the proportion of stocks in her portfolio:
    a. The standard deviation of the return on the stock market increases, but the expected return on the stock market remains the same.
    b. The expected return on the stock market increases, but the standard deviation of the stock market remains the same.
    c. The return on risk-free Treasury bills increases.