Chapter 5

Understanding Risk

Risk may be a four-letter word, but it's one we can't avoid. Every day we make decisions that involve financial and economic risk. How much car insurance should we buy? Should we refinance the mortgage now or a year from now? Should we save more for retirement, or spend the extra money on a new car? Making any decision that has more than one possible outcome is similar to gambling: We put the money on the roulette table and take our chances.

Interestingly enough, the tools we use today to measure and analyze risk were first developed to help players analyze games of chance like roulette and blackjack. For thousands of years, people have played games based on a throw of the dice, but they had little understanding of how those games actually worked. In ancient times, dice of various sorts were used to consult the gods, so any effort to analyze the odds was thought improper. But even those who ignored religious concerns could not correctly analyze a single throw of a die because they did not understand the concept of zero. That meant that the complex computations necessary to develop a theory of probability were impossible.[1]

By the mid-17th century, the power of religion had waned and mathematical tools had developed to the point that people could begin to make sense out of cards, dice, and other games.[2] Since the invention of probability theory, we have come to realize that many everyday events, including those in economics, finance, and even weather forecasting, are best thought of as analogous to the flip of a coin or the throw of a die. For better or worse, we no longer treat these random events as if they were divinely ordained.

Still, while experts can make educated guesses about the future path of interest rates, inflation, or the stock market, their predictions are really only that—guesses. And while meteorologists are fairly good at forecasting the weather a day or two ahead, economists, financial advisors, and business gurus have dismal records.[3] So understanding the possibility of various occurrences should allow everyone to make better choices. While risk cannot be eliminated, it can often be managed effectively.

Finally, while most people view risk as a curse to be avoided whenever possible, risk also creates opportunities. The payoff from a winning bet on one hand of cards can often erase the losses on a losing hand. Thus the importance of probability theory to

[1]For further details on this history, see F.N. David, *Games, Gods and Gambling: The Origins and History of Probability and Statistical Ideas from the Earliest Times to the Newtonian Era,* (New York, Hafner Publishing Co.,1962); P.L. Bernstein, *Against the Gods: The Remarkable Story of Risk,* (New York, John Wiley & Sons, 1998); and I. Hacking, *The Emergence of Probability: A Philosophical Study of Early Ideas about Probability, Induction and Statistical Inference,* (London, Cambridge University Press, 1975).

[2]Blaise Pascal's invention of probability theory in the mid-17th century is described in some detail in the supplementary material on the website www.mhhe.com/economics/cecchetti1e. This includes a description of the problem he tackled in developing probability theory and an analysis of the problem using the tools presented here.

[3]William A. Sherden's book *The Fortune Sellers,* (New York, NY: John Wiley & Sons, 1998) suggests that meteorologists are the only people who can legitimately claim to predict anything.

the development of modern financial markets is hard to overemphasize. People require compensation for taking risks. Without the capacity to measure risk, we could not calculate a fair price for transferring risk from one person to another, nor could we price stocks and bonds, much less sell insurance. The market for options didn't exist until economists learned how to compute the price of an option using probability theory.

In this chapter, we will learn how to measure risk and assess whether it will increase or decrease. We will also come to understand why changes in risk lead to changes in the demand for particular financial instruments and to corresponding changes in the price of those instruments.

Defining Risk

The dictionary definition of *risk*, the "possibility of loss or injury,"[4] focuses on the perils of putting oneself in a situation in which the outcome is unknown. But this common use of the word doesn't quite fit our purposes because we care about gains as well as losses. We need a definition of risk that focuses on the fact that the outcomes of financial and economic decisions are almost always unknown at the time the decisions are made. Here is the definition we will use:

Risk *is a measure of uncertainty about the future payoff to an investment*, measured over some *time horizon* and *relative to a benchmark*.

This definition has several important elements. First, risk is a *measure* that can be quantified. In comparing two potential investments, we want to know which one is riskier and by how much. All other things held equal, we expect a riskier investment to be less desirable than others and to command a lower price. Uncertainties that are not quantifiable cannot be priced.

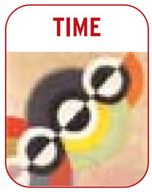

Second, risk arises from *uncertainty about the future*. We know that the future will follow one and only one of many possible courses, but we don't know which one. This statement is true of even the simplest random event—more things can happen than will happen. If you flip a coin, it can come up either heads or tails. It cannot come up both heads and tails or neither heads nor tails; only one of two possibilities will occur.

Third, risk has to do with the *future payoff* of an investment, which is unknown. Though we do not know for certain what is going to happen to our investment, we must be able to list all the possibilities. Imagining all the possible payoffs and the likelihood of each one is a difficult but indispensable part of computing risk.

Fourth, our definition of risk refers to an *investment* or group of investments. We can use the term *investment* very broadly here to include everything from the balance in a bank account to shares of a mutual fund to lottery tickets and real estate.

Fifth, risk must be measured over some *time horizon*. Every investment has a time horizon. We hold some investments for a day or two and others for many years. In most cases, the risk of holding an investment over a short period is smaller than the risk of holding it over a long one, but there are important exceptions to the rule that we will discuss later.[5]

Finally, risk must be measured *relative to a benchmark* rather than in isolation. If someone tells you that an investment is risky, you should immediately ask: "Relative to what?" The simplest answer is "Relative to an investment with no risk at all," called a *risk-free investment*. But there are other possibilities, often more appropriate. For example, in considering the performance of a particular investment advisor or money manager, a good **benchmark** is the performance of a group of experienced

[4]*Merriam-Webster's Collegiate Dictionary*, 11th ed. (Springfield, MA: Merriam-Webster, Inc., 2003).

[5]In Chapter 8 we will consider evidence that holding stock for one year is riskier than holding it for 20 years.

APPLYING THE CONCEPT
IT'S NOT JUST EXPECTED RETURN THAT MATTERS

Your life seems to be going well. You enjoy your job, and it pays well enough that you can put a little aside each month. You can't resist the dollar-for-dollar match your employer is offering on contributions to your retirement account, so you're slowly building up some long-term savings. But every so often, you wonder if you're saving enough. One day you go home and fire up the financial planning program on your computer, just to check.

Going through the retirement planner, you enter all the standard information: your age now and when you hope to retire; your salary and the value of all your assets; the monthly contribution to your retirement account and the monthly income you want at retirement. When you finish, the program asks what rate of return to assume. That is, how fast do you expect your savings to grow from now until your retirement? Following the suggestion on the screen and adjusting for inflation, you enter 7 percent, which is the average *real* return on the stock market over the last 75 years.* The light flashes green, signaling that you're on track to meet your financial goals. But are you?

*Inflation complicates computations over very long time periods. Price increases of 2 or 3 percent per year may not seem like much, but over 40 years they add up. At 2 percent inflation, prices double every 36 years. The simplest approach is to ignore inflation and measure income, wealth, and savings in current dollars. Then use a real rate of interest to compute future and present value.

continued on next page

investment advisors or money managers. If you want to know the risk associated with a specific investment strategy, the most appropriate benchmark would be the risk associated with other strategies.

Now that we know what risk is, how do we measure it? We use some rudimentary tools of probability theory, as we will see in the next section.

Measuring Risk

Armed with our definition of risk, we are now ready to quantify and measure it. In this section we will become familiar with the mathematical concepts useful in thinking about random events. We have already used some of these concepts. Recall from the last chapter that the *real* interest rate equals the *nominal* interest rate minus *expected* inflation. Without the proper tools, we weren't able to be explicit about what the term *expected inflation* means. The same is true of the term *expected return*. We see now that the best way to think about expected inflation and expected return is as the average or best guess—the *expected value*—of inflation, or the investment's return out of all the possible values.

Possibilities, Probabilities, and Expected Value

Probability theory tells us that in considering any uncertainty, the first thing we must do is to *list all the possible outcomes* and then *figure out the chance of each one occurring*. When you toss a coin, what are all the *possible* outcomes? There are two and only two.

continued from previous page

Maybe. The program did a series of future- and present-value calculations like the ones described in Applying the Concept: Retirement in Chapter 4. The green light means that if the assumptions you entered are valid, your saving rate is sufficient. That is, *if* your savings grow at 7 percent (adjusted for inflation), you'll be okay. So you need to decide whether you think 7 percent is a reasonable number. While it might be your best guess for the return (that's the *expected return*) over the next few decades, it is surely not the only possibility. You have very little sense of what the average return will be between now and the time that you retire.

To get a 7 percent expected return, you will have to take risk. And risk means that you could end up with less. What if your investment return is only 4 percent per year? Over 40 years that's an enormous difference. At 7 percent annual growth, one dollar today is worth nearly $15 in 40 years, and if you can save $1,000 per year you'll have over $200,000 saved up. Reducing the growth rate to 4 percent means that the future value of a dollar 40 years from now falls to less than $5. The lower return means that with the same $1,000 per year savings, you're left with less than $100,000 after 40 years.† You'll have to save twice as much to meet the same goal. Now that's risk!

You need to know what the possibilities are and how likely each one is. Only then can you assess whether your retirement savings plan is risky or not.

†These numbers are based on future-value calculations. If you save $1,000 per year, after 40 years you will have $\$1{,}000 \times (1.07)^{40} + \$1{,}000 \times (1.07)^{39} + \ldots + \$1{,}000 \times (1.07)^{2} + \$1{,}000 \times (1.07) = \$213{,}610$.

The coin can come down either heads or tails. What is the *chance* of each one of these two outcomes occurring? If the coin is fair, it will come down heads half the time and tails the other half; that's what we mean by *fair*. If we tossed a fair coin over and over again, thousands of times, it would come down heads half the time and tails the other half. But for any individual toss, the coin has an equal chance of coming down heads or tails. To quantify this statement, we can say that the *probability* that the coin will come up heads is one-half.

Probability is a measure of the likelihood that an event will occur. It is always expressed as a number between zero and one. The closer the probability is to zero, the *less* likely it is that an event will occur. If the probability is exactly zero, we are sure that the event will *not* happen. The closer the probability is to one, the *more* likely it is that an event will occur. If the probability is exactly one, the event *will* definitely occur.

Some people prefer to think of random outcomes in terms of frequencies rather than probabilities. Instead of saying that the probability of a coin coming down heads is one-half, we could say that the coin will come down heads once every two tosses on average. Probabilities can always be converted into frequencies in this way.

To grasp these concepts, it is helpful to construct a table. The table lists everything that can happen (all the possibilities) together with their chances of occurring (their probabilities).[6] Let's start with a single coin toss. Table 5.1 lists the possibilities—heads or tails—and the probabilities, both equal to one-half.

[6]In the language of probability and statistics, the first step is to construct the probability distribution for the possible outcomes.

Table 5.1 A Simple Example: All Possible Outcomes of a Single Coin Toss

Possibilities	Probability	Outcome
#1	$\frac{1}{2}$	Heads
#2	$\frac{1}{2}$	Tails

In constructing a table like this one, we must be careful to list *all* possible outcomes. In the case of a coin toss, we know that the coin can come down only two ways, heads or tails. We know that one of these outcomes *must* occur. We just don't know which one.

One important property of probabilities is that we can compute the chance that one *or* the other event will happen by adding the probabilities together. In the case of the coin flip there are only two possibilities; the probability that the coin will come up either heads or tails must be one. If the table is constructed correctly, then, *the values in the probabilities column will sum to one*.

Let's move from a coin toss to something a bit more complicated: an investment that can rise or fall in value. Assume that for $1,000 you can purchase a stock whose value is equally likely to fall to $700 or rise to $1,400. We'll refer to the amount you could get back as the investment's **payoff**. Following the procedure we used to analyze the coin toss, we can construct Table 5.2. Again we list the possibilities and the probability that each will occur, but we add their payoffs (column 3).[7]

We can now go a step further and compute what is called the **expected value** of the investment. We are familiar with the idea of expected value as the **average** or most likely outcome. The expected value is also known as the **mean**. After listing all of the possible outcomes and the probabilities that they will occur, we compute the expected value as the sum of their probabilities times their payoffs. (Another way to say this is that the expected value is the probability-weighted sum of the possible outcomes.)

Computing the expected value of the investment is straightforward. In Table 5.2, the first step is to take the probabilities in the second column and multiply them by their associated payoffs in the third column. The results are in the fourth column. Summing them, we get

$$\text{Expected value} = \tfrac{1}{2}(\$700) + \tfrac{1}{2}(\$1{,}400) = \$1{,}050$$

which appears at the bottom of the table.

[7]As you go through the examples in the chapter, be aware that it is often very difficult to estimate the probabilities needed to do the risk computations. The best way to do it is often to look at history. Investment analysts usually estimate the possibilities and probabilities from what happened in the past.

Table 5.2 Investing $1,000: Case 1

Possibilities	Probability	Payoff	Payoff × Probability
#1	$\frac{1}{2}$	$700	$350
#2	$\frac{1}{2}$	$1,400	$700
Expected Value = Sum of (Probability times Payoff) = $1,050			

The expected value of an investment is a very useful concept, but it can be difficult at first. The problem is that if we make this investment only once, we will obtain either $700 or $1,400, not $1,050. In fact, regardless of the number of times we make this particular investment, the payoff will *never* be $1,050. But what would happen if we were to make this investment 1 million times? About 500,000 of those times the investment would pay off $1,400 and the other 500,000 times it would pay off $700. (Notice that we just converted the probabilities into frequencies.) So the *average* payoff from the 1 million investments would be

$$\frac{500,00}{1,000,000}(\$700) + \frac{500,00}{1,000,000} = (\$1,400) = \$1,050 \text{ (the expected value).}$$

While the world of casino gambling may offer simple bets with just two outcomes, the financial world rarely does. To make the example more realistic, let's double the number of possibilities and look at a case in which the $1,000 investment might pay off $100 or $2,000 in addition to $700 or $1,400. Table 5.3 shows the possibilities, probabilities, and payoffs. We'll assume that the two original possibilities are the most likely; the two new possibilities are much less likely to occur. Note that the probabilities sum to one: 0.1 + 0.4 + 0.4 + 0.1 = 1. Again, we could convert the probabilities to frequencies, so that 0.4 means 4 out of 10. And again, we can compute the expected value by multiplying each probability times its associated payoff and then summing them. So $100 would be the payoff one out of every 10 times, $700 the payoff four out of every 10 times, and so on. To compute the expected value, we would find the average of these 10 investments: $100 + $700 + $700 + $700 + $700 + $1,400 + $1,400 + $1,400 + $1,400 + $2,000 = $10,500; and $10,500/10 = $1,050. Once again the expected value is $1,050.

Because the expected value of this $1,000 investment is $1,050, the expected gain is $50. But most people don't discuss investment payoffs in terms of dollars; instead, they talk about the percentage return. Expressing the return as a percentage allows investors to compute the gain or loss on the investment regardless of the size of the initial investment. In this case, the **expected return** is $50 on a $1,000 investment, or 5 percent. Note that the two $1,000 investments we just discussed are not distinguishable by their expected return, which is 5 percent in both cases. Does that mean an investor would be indifferent between them? Even a casual glance suggests that the answer is no because the second investment has a wider range of payoffs than the first.

Table 5.3 Investing $1,000: Case 2

Possibilities	Probability	Payoff	Payoff × Probability
#1	0.1	$100	10
#2	0.4	$700	280
#3	0.4	$1,400	560
#4	0.1	$2,000	200

Expected Value = Sum of (Probability times Payoff) = $1,050

The highest payoff is higher and the lowest payoff lower than for the first investment. So the two investments carry different levels of risk. The next section discusses measures of risk.

One last word on expected values. Recall from the last chapter that to compute the real interest rate, we need a measure of *expected inflation*. One way to calculate expected inflation is to use the technique we just learned. That is, list all the possibilities for inflation, assign each one a probability, and then calculate the expected value of inflation.

Measures of Risk

Most of us have an intuitive sense of risk and its measurement. For example, we know that walking on a sidewalk is usually a safe activity. But imagine that one day as you are strolling along, you come upon a three-foot hole in the sidewalk. The only way across is to jump over it. If the hole is just a few inches deep, it won't stop you. But the deeper it is, the greater the risk of jumping across because the greater the range of injuries you could sustain. We all have an intuitive sense that the wider the range of

YOUR FINANCIAL WORLD

Choosing the Right Amount of Car Insurance

Car insurance is expensive, especially for young drivers. That should be no surprise, since the younger you are, the more likely you are to have an accident. Only about one in seven drivers is under 25 years old, but more than one quarter of the 10 million accidents that happen each year involve a driver between 16 and 24. Men are worse risks than women. It's hard to fault insurance companies for charging higher premiums to drivers who are more likely than others to file claims.

While you must have some insurance—most states require that you have *liability insurance,* to pay for damage and injuries to others if you cause an accident—you do have some choices. The most important choice is whether or not to buy collision insurance, which pays for damage to your car when the accident is your fault. If you go without it, you'll have to pay for the repairs if you cause a crash.

There is no easy way to figure out how much collision insurance to buy, but there are a few things to think about when you make your decision. First, how much is your car worth? If you do have an accident, the insurance company doesn't promise to fix your car regardless of the costs. Instead, the company will pay you what it is worth. So if your car is old and you crash, the odds are you'll get a check, not a repaired car. Buying collision insurance on old cars is rarely worth it.

What should you do if you have a new car? Here the question is not whether to buy collision insurance but how much. The choice is in something called a *deductible,* the amount you pay for the repair after a crash. If you have a $250 deductible, you'll pay the first $250 and the insurance company will pay the rest. The higher your deductible is, the lower your insurance premium will be.

To see how much your premium can vary, let's look at an example: a 19-year-old male driving a new Saturn (a four-door sedan that cost around $13,000 in 2004). A college student living away from home, he has a good driving record and a good student discount. With a $250 collision deductible, his insurance costs about $2,400 per year. Raising the deductible to $500 would lower the premium by $150 per year, to around $2,250. Can $250 worth of extra insurance possibly be worth paying an extra $150 a year? Only if the driver expects to have an accident once every 20 months. Ideally he won't, so the extra insurance isn't worth paying for.

outcomes, the greater the risk. That's why the investment that has four possible payoffs (Table 5.3) seems riskier than the one with two possible payoffs (Table 5.2).

Thinking about risk in terms of the range of possible outcomes is straightforward. The best way to do it is to start with something that has no risk at all—a sidewalk without a hole in it or an investment with only one possible payoff. We will refer to a financial instrument with no risk at all as a risk-free investment or risk-free asset. A **risk-free asset** *is an investment whose future value is known with certainty and whose return is the* **risk-free rate of return**.[8] The payoff that you will receive from such an investment is guaranteed and cannot vary. For instance, if the risk-free return is 5 percent, a $1,000 risk-free investment will pay $1,050, its expected value, with certainty. If there is a chance that the payoff will be either more or less than $1,050, the investment is risky.

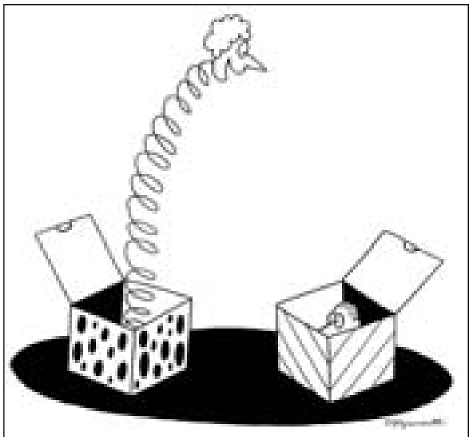

"Come on, Louis. No risk, no reward."

Let's compare this risk-free investment with the first investment we looked at, the one in which $1,000 had an equal chance of turning into $1,400 or $700 (see Table 5.2). That investment had the same expected return as the risk-free investment, 5 percent. The difference is that the payoff wasn't certain, so risk was involved. What caused the risk was the increase in the spread of the potential payoffs. The larger the spread, the higher the risk.

These examples suggest that we can measure risk by quantifying the spread among an investment's possible outcomes. We will look at two such measures. The first is based on a statistical concept called the *standard deviation* and is strictly a measure of spread. The second, called *value at risk,* is a measure of the riskiness of the worst case. When the hole in the sidewalk gets deep enough, you risk being killed if you fall in.

Variance and Standard Deviation The **variance** is defined as the probability-weighted average of the squared deviations of the possible outcomes from their expected value. To calculate the variance of an investment, you can compute the expected value and then subtract it from each of the possible payoffs. Then you square each of the results, multiply it by its probability, and add up the results. In the example of the $1,000 investment that pays either $700 or $1,400, the steps are

1. Compute the expected value: $(\$1{,}400 \times \tfrac{1}{2}) + (\$700 \times \tfrac{1}{2}) = \$1{,}050$.
2. Subtract this from each of the possible payoffs:
   $\$1{,}400 - \$1{,}050 = \$350$
   $\$700 - \$1{,}050 = -\$350$
3. Square each of the results: $\$350^2 = 122{,}500 \text{ dollars}^2$ and $(-\$350)^2 = 122{,}500 \text{ dollars}^2$
4. Multiply each result times its probability and add up the results:

$$\tfrac{1}{2}\,[122{,}500 \text{ dollars}^2] + \tfrac{1}{2}\,[122{,}500 \text{ dollars}^2] = 122{,}500 \text{ dollars}^2$$

Writing this procedure more compactly, we get

$$\begin{aligned}\text{Variance} &= \tfrac{1}{2}(\$1{,}400 - \$1{,}050)^2 + \tfrac{1}{2}(\$700 - \$1{,}050)^2\\ &= 122{,}500 \text{ dollars}^2\end{aligned}$$

[8]In most financial markets, no truly risk-free asset exists, so the risk-free rate of return is not directly observable. Regardless of our inability to measure it exactly, the risk-free rate of return remains a useful concept.

The **standard deviation** is the square root of the variance,[9] or

$$\text{Standard deviation} = \sqrt{Variance} = \sqrt{122{,}500\ dollars^2} = \$350$$

The standard deviation is more useful than the variance because it is measured in the same unit as the payoffs: dollars. (Variance is measured in dollars squared.) That means that we can convert the standard deviation into a percentage of the initial investment of $1,000, or 35 percent. This calculation provides a baseline against

[9] Note that we first find squared deviations of individual outcomes from the expected value. That's the variance. We square the differences when calculating the variance so that payoffs above and below the expected payoff don't cancel each other out. We then take the square root of the variance to get the standard deviation. This gives us a measure of the average difference between the possible payoffs in a way that treats those above and below the expected value equally.

TOOLS OF THE TRADE

The Impact of Leverage on Risk

"Funds Use Leverage to Magnify Returns, But Risks Grow Too" read a headline in *The Wall Street Journal* (November 1, 2002). What is leverage, and how does it affect risk and return? **Leverage** is the practice of borrowing to finance part of an investment. Common examples of leverage are borrowing to buy stock (called a *margin loan*) and borrowing to acquire a house (called a mortgage).* In the case of a margin loan, an investor borrows from a brokerage firm to increase the quantity of stock purchased.

To understand the effects of leverage, let's look at an investment of $1,000 with an expected return of 5 percent (a gain of $50) and a standard deviation of 35 percent ($350). That's the example in Table 5.2. What if in addition to investing $1,000 of your own, you borrow $1,000 and invest a total of $2,000? This investment strategy changes the risk involved. The reason is that the lender wants to be repaid the $1,000 loan regardless of how much your investment returns. If the investment's payoff is high, your $2,000 investment will increase in value to $2,800. After repaying the $1,000 loan, you will be left with $1,800—an increase of $800 over your initial investment of $1,000. If your investment falls in value, the $2,000 will become $1,400. After repaying the $1,000 loan, you will be left with $400—a loss of $600.

Since these two results are equally likely, the expected value of your leveraged investment is $\frac{1}{2}(\$1{,}800) + \frac{1}{2}(\$400) = \$1{,}100$. Your expected gain—the difference between your investment of $1,000 and its expected value of $1,100—is now $100 and your expected return is 10 percent. That's double the expected return from your investment of $1,000 without any borrowing—double what it would be without leverage. So we have part of the answer to our question: *Leverage increases the expected return.*

$$\text{St. dev.} = \sqrt{\frac{1}{2}(1{,}800-1{,}100)^2 + \frac{1}{2}(400-1{,}100)^2} = \$700$$

But what about risk? To figure it out, let's calculate the standard deviation of your leveraged investment.

The standard deviation has doubled too—*twice the expected return at twice the risk!*

We can repeat these calculations for any amount of leverage we want. For example, homebuyers commonly pay 20 percent of the price of a house with their savings and borrow the remaining 80 percent. Since mortgage lenders expect to be repaid, changes in the price of the house become gains or losses to the owner. Say you buy a $100,000 house by borrowing $80,000 and paying $20,000 from your savings, often called your *equity*. A 10 percent increase in your home's value would raise the price to $110,000. Subtracting the $80,000 you borrowed, your $20,000 down payment would become $30,000, a 50 percent increase. On the other hand, if your home's value fell by 10 percent, you would *lose* half your $20,000 down payment. *Leverage magnifies the effect of price changes* (see Figure 5.1).

*Corporate borrowing through the issuance of bonds is another form of leverage. As we will see in Chapter 8, this type of leverage affects the risk of owning a firm's equity or stock.

which we can measure the risk of alternative investments. Given a choice between two investments with the same expected payoff, most people would choose the one with the lower standard deviation. A higher-risk investment is less desirable.

Let's compare this two-payoff investment with the one that has four possible payoffs (Table 5.3). We already concluded that the second investment is riskier, since the payoffs are more spread out. But how much riskier is it? To answer this question, let's compute the standard deviation:

$$\text{St. dev.} = \sqrt{0.1(100-1{,}050)^2 + 0.4(1{,}400-1{,}050)^2 + 0.4(700-1{,}050)^2 + 0.1(2{,}000-1{,}050)^2}$$
$$= \sqrt{0.1(950)^2 + 0.4(350)^2 + 0.4(350)^2 + 0.1(950)^2}$$
$$= \$528$$

We can use these examples to develop a formula for the impact of leverage on the expected return and standard deviation of an investment. If you borrow to purchase an asset, you increase both the expected return and the standard deviation by a leverage ratio of

$$\text{Leverage ratio} = \frac{\text{Cost of investment}}{\text{Owner's contribution to the purchase}}$$

where the "Owner's contribution to the purchase" in the denominator is just the cost of investment minus the amount borrowed. If the expected return and standard deviation of the unleveraged investment are 5 percent and 35 percent (as in our first example), then borrowing half and contributing half means that for each dollar invested, the buyer is contributing 50 cents. The formula tells us that the leverage ratio is 1/0.5, which equals 2. Thus the investment's expected return is 2 × 5 percent = 10 percent, and its standard deviation is 2 × 35 percent = 70 percent. And if the homeowner borrows 80 percent of the purchase price of the house, his or her contribution is 20 percent, so the leverage ratio is $\frac{1}{(1-80/100)} = \frac{1}{0.2} = 5$ times what it would be for someone who could buy the house outright, with no mortgage.

We have focused on the impact of leverage on risk, but leverage has at least as big an impact on value at risk. Note that for the $1,000 investment without leverage in Table 5.2, the worst case was a loss of $300, or 30 percent, half the time. If an investor borrowed 90 percent of the funds needed to make the investment, half the time the investor would lose not only the entire $100 invested but an additional $200 of borrowed funds as well. *Leverage compounds the worst possible outcome.*

Figure 5.1 The Effect of Leverage on Risk and Return

To understand leverage, picture a set of two gears, one large and one small. The movement in the price of the leveraged investment is measured by the number of revolutions in the big gear. The investor's risk and return are measured by the number of revolutions in the small gear. The bigger the big gear, the more times the small gear goes around with each revolution of the big gear. That's leverage.

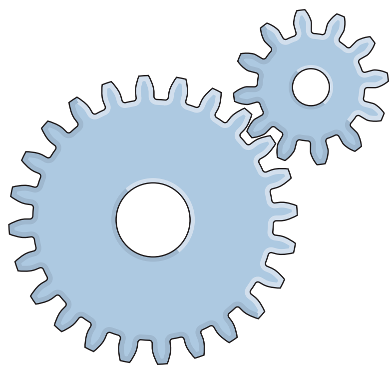

This result is much higher than the $350 standard deviation of the first investment, which has only two possible payoffs. Since the two investments have the same expected value, the vast majority of people would prefer the first. *The greater the standard deviation is, the higher the risk.*

Value at Risk Standard deviation is the most common measure of financial risk, and for most purposes it is adequate. But in some circumstances we need to take a different approach. Sometimes we are less concerned with the spread of possible outcomes than with the value of the worst outcome. For example, no one wants the local bank to close its doors. Neither the depositors nor the government regulators care how well or badly the bank's shareholders fare as long as they do well enough to keep the doors open. The concept used to assess this sort of risk is called **value at risk (VaR)**.

Let's assume you are considering buying a house. In going through your finances, you conclude that you can afford a monthly mortgage payment of $750 per month and not a dollar more. You find a nice house and a mortgage lender who is willing to lend you $100,000 to buy it. You expect to complete the purchase transaction and move in within six months.

Now assume that the current mortgage interest rate is 7 percent, yielding a $651 monthly payment that is well within your budget. But over the next six months, interest rates could go up or down. A decline in interest rates reduces your monthly payment, but a rise in rates increases it. This creates a risk for you. If interest rates rise high enough, the required monthly payment might exceed your budget.

Realizing your predicament, the mortgage company offers you a sort of insurance policy. It will commit to financing your mortgage at a fixed rate of 8 percent, or one percentage point above the current market rate. The alternative is to wait and take the interest rate that prevails in six months when you actually get the loan. What should you do?

To decide, start by listing the possibilities and their associated probabilities—that is, the possible interest rates and the likelihood that each one will occur. Let's assume that the rate could stay the same, rise to 10 percent, or fall to 6 percent. There is a 0.50 probability that it will stay at 7 percent, a 0.25 probability that it will rise to 10 percent, and a 0.25 probability that it will fall to 6 percent. Each of these possibilities implies a monthly payment on a 30-year, $100,000 fixed loan (see Table 5.4). The 10 percent loan means a payment of $846 (too high) and the 8 percent loan requires a monthly payment of $714.[10] Should you accept the risk of interest-rate fluctuations or take the insurance and lock in the interest rate at 8%?

The fixed rate of 8 percent is insurance against a risk, but the risk of what? We could compute the standard deviation of the monthly payment under both cases (and you should do that as an exercise). But without picking up a calculator, we know that the standard deviation of the payment associated with the 8 percent rate is zero and the standard deviation of the fluctuating rate is greater than zero. With a calculator we can figure out that the expected value of the interest rate is 7.5 percent. But what does that tell us?

In this case, the computation of the expected value and standard deviation does not seem to get at the heart of the problem. The reason is that it doesn't take proper account of the worst case, when the interest rate rises to 10 percent and your required monthly payments climb to $846 per month—well beyond your $750 budget. If that happens, you don't just lose the $132 per month difference between the guaranteed 8 percent payment and the 10 percent payment. You lose the house!

[10]As a rule, on a long-term mortgage, the monthly payments will approximate the interest alone. For instance, if we take a 10 percent, 30-year mortgage and compute the monthly interest rate, we get 0.797 percent. Paying the monthly interest costs $797, just $49 less than the full payment of $846.

Table 5.4 Monthly Payments on a $100,000 Mortgage

Interest Rate	Approximate Monthly Payment	Probability
10%	$846	0.25
8%	$714	—
7%	$651	0.50
6%	$589	0.25

Monthly payments on a 30-year, $100,000 fixed-rate mortgage at various interest rates.

APPLYING THE CONCEPT
GOVERNMENT-RUN LOTTERIES

Governments use lotteries to finance a range of activities, including public schools and the arts. But for lotteries to remain profitable, the people who run them must keep a large percentage of the revenue. State-run lotteries commonly pay out only 60 percent of the revenue they receive. That is, for each $1 bet on the lottery, the government pays out 60 cents. The expected value of a $1 lottery ticket is 60 cents.

Lotteries, then, are a risky investment. And since people generally don't like risk, you would think that the government would have to pay a premium to get people to play. Instead, the opposite is true. Millions of people pay good money for a very small chance to win big. As the jackpots grow larger, the lines of those waiting to buy tickets grow longer and longer. How can we explain this puzzle?

One answer is that playing the lottery is a form of entertainment, like going to the movies. But that doesn't really seem to explain it. We can use the concept of value at risk to provide a more coherent explanation. Compare paying $1 for a chance to win $1 million with paying $10,000 for a chance to win $10 billion. We see people spending $1 but not $10,000 for lottery tickets. The reason is that the risk of losing $1 is inconsequential, but the potential gain of $1 million is significant. The risk of losing $10,000, however, would loom large even compared to a payoff of $10 billion. Value-at-risk calculation tells us that the $1 lottery isn't very risky, and that's why people play.*

*Before you run out and buy a lottery ticket, think about the lottery's advertising. The people who run lotteries are allowed to advertise that they are paying a jackpot of, say, $10 million, when they are really promising to pay $500,000 per year for 20 years. At a 6 percent interest rate, the present value of 20 payments of $500,000 per year is about $6 million, not $10 million. If private companies tried to advertise that way, they would probably get into trouble.

This example highlights the fact that sometimes risk should be measured by the value of the worst case rather than by expected value and standard deviation. Value at Risk, which measures risk as the maximum potential loss, is more appropriate in a case like this. VaR is the answer to the question: How much will I lose if the worst possible disaster occurs? In the $1,000 investment summarized in Table 5.2, the worst case was a loss of $300. In the more complex $1,000 investment, summarized in Table 5.3, the value at risk was $900, the most you could possibly lose. In the mortgage example (Table 5.4), the value at risk is the house: If the payment increases to more than $750 a month, you can't make the payments on your loan. If that happens, you will be forced to sell the house and move.

A more sophisticated value-at-risk analysis would include a time horizon and probabilities. In fact, the formal definition of value at risk is *the worst possible loss over a specific time horizon, at a given probability*. For the mortgage example, the time horizon is the six months over which the interest rate can move, and the probability that the worst case will actually occur is 0.25. VaR is a measure of risk that we will find very useful in discussing the management and regulation of financial institutions. Bank managers and regulators work hard to ensure that financial collapse is an extremely remote possibility, and to do it they employ the concept of value at risk.

Risk Aversion, the Risk Premium, and the Risk-Return Trade-off

The implication of our discussion so far is that most people don't like risk and will pay to avoid it. While some people enjoy risky activities like skydiving and car racing, most of us are more careful. And while some people gamble large sums, most of us don't because we can't sustain large losses comfortably. In fact, the reason we buy insurance is that we want someone else to take the risk. Insurance is an interesting case; remember, for an insurance company to make a profit, it must charge more than it expects to pay out. Thus, insurance premiums are higher than the expected value of the policyholder's losses. We pay to avoid risks because most of us are *risk averse*.

To understand risk aversion, imagine that you are offered a single chance to play a game in which a fair coin will be tossed. If it comes up heads you will win $1,000; if it comes up tails, you will get nothing. How much would you be willing to pay to play the game just once? The expected value of the game is $500—that is, on average, the game yields $500—but you may play only one time. Would you pay $500 to play the game? If so, you are **risk neutral**. Most people would not play the game at $500, though they would at less than that amount. These people are **risk averse**.

Because the coin toss is similar to an investment, we can apply the same logic to investor behavior and conclude that *a risk-averse investor will always prefer an investment with a certain return to one with the same expected return but any amount of uncertainty*. (A risk-neutral person wouldn't care as long as the expected return is the same.) A risk-free investment with a guaranteed return is clearly preferable to a risky investment with the same expected return but an uncertain outcome. In the case of the coin toss, most people would take $500 with certainty rather than risk tossing the coin and getting double or nothing.

"I'm sorry, but you don't get frequent flyer miles for regularly investing in high risk securities."

One result of this desire to avoid risk is that investors require compensation for taking risk. That's the flip side of buying insurance. When we buy insurance, we pay someone else to take our risks, so it makes sense that if someone wants us to take on a risk, we need to be paid to do it. A risky investment, then, must have an expected return that is higher than the return on a risk-free asset. In economic terms, it must offer a **risk premium**. In general, *the riskier an investment (the higher the compensation investors require for holding it), the higher the risk premium* (see Figure 5.2).

By extension, if riskier investments have higher risk premiums, they must have higher expected returns. Thus, there is a trade-off between risk and expected return; you can't get a high return without taking considerable risk. So if someone tells you he or she made a big return on an investment, you should suspect that it was a very risky investment. No risk, no reward!

Sources of Risk: Idiosyncratic and Systematic Risk

Risk is everywhere. It comes in many forms and from almost every imaginable place. In most circumstances the sources of risk are obvious. For drivers, it's the risk of an accident; for farmers, the risk of bad weather; for investors, the risk of fluctuating stock prices. Regardless of the source, however, we can classify all risks into one of two groups: (1) those affecting a small number of people but no one else and (2) those

Figure 5.2 The Trade-off between Risk and Expected Return

The higher the risk, the higher the expected return. The risk premium equals the expected return on the risky investment minus the risk-free return.

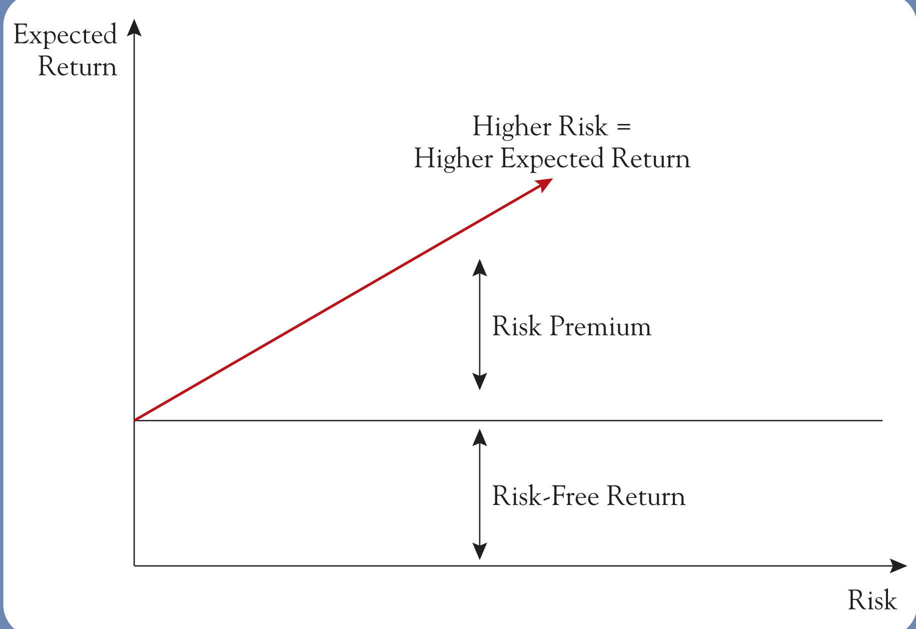

YOUR FINANCIAL WORLD

Your Risk Tolerance

How much risk should you tolerate? Figuring that out isn't easy, but there are a few ways to get some sense of the right level of risk for you. First, there are risk quizzes, short sets of questions financial advisers give their clients to determine the level of risk they can live with. For instance, "What would you do if a month after you invest in the stock market, the value of your stocks suddenly falls by 20 percent?" Answers might include "Sell right away," "Nothing," and "Buy more." Taking such a quiz is a useful first step, so you might want to try the one in Appendix 5A.

But don't stop there. Even if you are willing to take risks, that doesn't mean you should. You may not have time to make back the losses you might suffer. Think about the difference between a 25-year-old and a 60-year-old both saving for their retirement. Which one of these people can afford to suddenly lose a quarter of her savings? Obviously, it is the 25-year-old. If a 60-year-old loses a quarter of her retirement savings, it's a disaster! Likewise, if you're saving to buy a car or a home, the sooner you are planning to make the purchase, the less you can afford to lose what you have. Always ask yourself: How much can I stand to lose? The longer your time horizon (and the wealthier you are), the more risk you can tolerate.

affecting everyone. We'll call the first of these **idiosyncratic** or *unique* **risks** and the second **systematic** or *economywide* **risks**.[11]

To understand the difference between idiosyncratic and systematic risk, think about the risks facing General Motors stockholders. Why would the value of GM's stock go up or down? There are two reasons. First, there is always the risk that GM will lose its position as the largest producer of cars and trucks in the world. The fact is, DaimlerChrysler, Ford, and Toyota are working every day to take away some of GM's business. So GM may do poorly compared to its competition, and its market share may shrink (see Figure 5.3). This risk is unique to GM because if GM does poorly, someone else must be doing better. Idiosyncratic risk affects specific firms, not everyone.

The second risk GM's stockholders face is that the entire auto industry will do poorly (see Figure 5.3). This is systematic, economywide risk. If idiosyncratic risk is a change in the *share* of the auto-market pie, systematic risk is a change in the *size* of the pie. It is the risk that everyone will do poorly at the same time. Sales of cars and trucks could simply collapse for reasons that are completely unrelated to any individual company's performance.

We can apply the concept of idiosyncratic and systematic risk to the entire economy. Surely some events will affect firms like GM and Ford in one way and other firms in another way. An example would be a change in the price of oil. History tells us that when oil prices rise, auto sales fall and the automobile industry suffers. But higher oil prices improve the profits of firms that supply energy, such as ExxonMobil, Shell, and Texaco. An oil price change that is bad for GM is good for the oil companies. Looking at the economy as a whole, this is an idiosyncratic risk.

Not all idiosyncratic risks are balanced by opposing risks to other firms or industries. Some unique risks are specific to one person or company and no one else. The risk that two people will have an automobile accident is unrelated to whether anyone else has one. We include these completely independent risks in the category of idiosyncratic risks.

[11]These are also sometimes referred to as *specific* and *common risks*.

Figure 5.3 Idiosyncratic and Systematic Risk

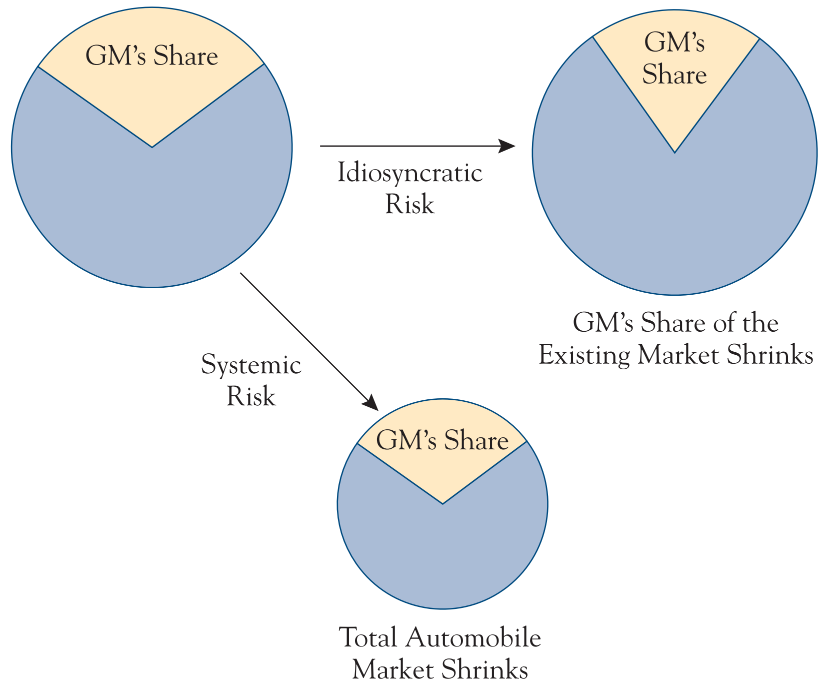

Systematic risks affect *all* firms and individuals in the entire economy. They are economywide risks that come from changes in general economic conditions that have an impact on everyone. Macroeconomic factors, such as swings in consumer and business confidence brought on by changes in the political climate or in the global economic conditions, are the source of systematic risks.

Reducing Risk through Diversification

When George T. Shaheen left his $4 million-a-year job overseeing 65,000 employees of a large management consulting firm to become chief executive of the Webvan Group, he may not have realized how much of a risk he was taking. He thought Webvan would change the way people bought their groceries. Consumers would order their cereal, milk, apples, and ice cream over the Internet, and Webvan would deliver to their door. In November 1999, just a few months after Mr. Shaheen joined the company, his stock in Webvan was worth more than $280 million. But by April 2001, his shares were worth a paltry $150,000 and Mr. Shaheen had left the company. On July 10, 2001, Webvan collapsed and stockholders were left with nothing.

What happened to Webvan and its plan to change the way people shop? Maybe people actually like getting out of the house and going to the grocery store. But this story is about more than shopping; it's also about risk. Shaheen took on so much risk that a single big loss wiped him out. Traders in the financial markets call this experience

YOUR FINANCIAL WORLD

A Guide to Evaluating Risk

Deciding whether a risk is worth taking is extremely difficult, but some simple rules can help. Let's start with the investment described in Table 5.2, where $1,000 yields either $1,400 or $700 with equal probability. If we think about it in terms of gains and losses, this investment offers an equal chance of gaining $400 or losing $300. Should you take the risk? The answer depends on how risk averse you are, but most of us would say no. To see why, let's break the investment down into two parts, the gain and the loss (see Table 5.5).

Taking the gain first, how much would you pay for a 50 percent chance of making $400? Again, the answer depends on your risk aversion, but you surely would pay less than $200, the expected value of such an investment. Let's assume that your answer is $150.

Next, let's turn to the loss. How much would you be willing to pay to avoid a $300 loss altogether? To put it another way, assume that you risk losing $300 and are considering buying insurance against the loss. The insurance company will take the bet for you, losing the $300 in your place if that is the outcome. How much would you be willing to pay an insurance company to avoid taking a 50 percent chance of losing $300? Again, the answer depends on how risk averse you are, but we know that you will pay more than the expected value of the loss, which is $150. (The insurance company would insist on receiving more.) Let's assume you will pay $200 to avoid the loss.

Now we are ready to answer our original question: Is the value of the potential gain sufficient to compensate you for the cost of the potential loss? Subtracting the $200 that you are willing to pay to avoid the $300 loss from the $150 you will pay for the opportunity to gain $400, we get $150 - $200 = -$50, a result less than zero. In short, the potential gain is not big enough to compensate you for the potential loss, so you should not take the risk. In fact, our computation suggests you would be willing to pay $50 *not* to make this investment!

Table 5.5 Evaluating the Risk of a $1,000 Investment

A. The Gain

Payoff	Probability
+$400	$\frac{1}{2}$
$0	$\frac{1}{2}$

B. The Loss

Payoff	Probability
$0	$\frac{1}{2}$
–$300	$\frac{1}{2}$

Deciding if a Risk Is Worth Taking

1. List all the possible outcomes, or payoffs.
2. Assign a probability to each possible payoff.
3. Divide the payoffs into gains and losses.
4. Ask how much you would be willing to pay to *receive* the gain.
5. Ask how much you would be willing to pay to *avoid* the loss.
6. If you are willing to pay more to receive the gain than to avoid the loss, you should take the risk.

"blowing up." Surely Shaheen could have done something to protect at least a portion of his phenomenal wealth from the risk that it would suddenly disappear. But what?

Cervantes answered this question in *Don Quixote* in 1605: "It is the part of a wise man to keep himself today for tomorrow, and *not to venture all his eggs in one basket* [emphasis added]." In today's terminology, risk can be reduced through **diversification**, the principle of holding more than one risk at a time. Though it may seem counterintuitive, holding several different investments can reduce the overall risk an investor bears. A combination of risky investments is often less risky than any one individual investment. There are two ways to diversify your investments. You can *hedge* risks or you can *spread* them among the many investments. Let's discuss hedging first.

Hedging Risk

Hedging is the strategy of reducing overall risk by making two investments with opposing risks. When one does poorly, the other does well, and vice versa. So while the payoff from each investment is volatile, together their payoffs are stable.

Consider the risk an investor faces from a potential change in the price of oil. Increases in the price of oil are bad for most of the economy, but they are good for oil companies. So an investor might buy stock in both General Electric (GE), maker of everything from lightbulbs to dishwashers and jet engines, and Texaco, a large oil company. For the sake of our example, let's assume that oil prices have an equal chance of rising or falling. When they rise, owners of Texaco stock receive a payoff of \$120 for each \$100 they invested. When oil prices fall, Texaco's shareholders just get their \$100 investment back. The reverse is true for GE. When oil prices fall, owners of GE stock get \$120 for each \$100 they invested; when oil prices rise, they get \$100.

Table 5.6 summarizes these relationships.

Table 5.6 Payoffs on Two Separate Investments of \$100

	Payoff from Owning Only		
Possibility	GE	Texaco	Probability
Oil prices rise	\$100	\$120	$\frac{1}{2}$
Oil prices fall	\$120	\$100	$\frac{1}{2}$

Let's compare three strategies for investing \$100, given the relationships shown in the table:

1. Invest \$100 in GE.
2. Invest \$100 in Texaco.
3. Invest half in each company: \$50 in GE and \$50 in Texaco.

Regardless of whether you invest \$100 in GE or Texaco, the expected payoff is $\frac{1}{2}\$120 + \frac{1}{2}\$100 = \$110$; and the

$$\text{Standard deviation of the payoff} = \sqrt{\frac{1}{2}(\$120 - \$110)^2 + \frac{1}{2}(\$100 - \$110)^2} = \$10$$

But what about the third option? What if you split your \$100 and put half in GE and half in Texaco? Since \$50 is half the size of your initial investment, the payoff is half as big as well—a \$50 investment in either stock pays off either \$60 or \$50. But the important point about this strategy is that it reduces your risk (see Table 5.7). When oil prices go up, Texaco does well but GE does badly. When oil prices fall, the reverse happens. Regardless of whether oil prices go up or down, you will get back \$110 on your \$100 investment. Investing \$50 in each stock ensures your payoff. Hedging—splitting your investment between two stocks with different payoff patterns—has eliminated your risk entirely.

Could George Shaheen have hedged the risk of owning so much Webvan stock? To do it, he would have had to find a company whose stock price would rise when Webvan's

Table 5.7 Results of Possible Investment Strategies: Hedging Risk
Initial Investment = $100

Investment Strategy	Expected Payoff	Standard Deviation
GE only	$110	$10
Texaco only	$110	$10
$\frac{1}{2}$ and $\frac{1}{2}$	$110	$0

fell. That would have been difficult, since Webvan's business concept was new and untested. But Shaheen did have another option.

Spreading Risk

Because investments don't always move predictably in opposite directions, you can't always reduce risk through hedging. Fortunately, there is another way. You can simply **spread risk** around—and that's what George Shaheen should have done. To spread your risk, all you need to do is find investments whose payoffs are unrelated. Let's replace Texaco with Microsoft and assume that GE and Microsoft's payoffs are independent of each other. So we toss a coin once to see if GE does well or badly, and then we toss it a second time to see how Microsoft does. As before, a $100 investment in either company pays off either $120 or $100 with equal probability.

Again, we'll consider three investment strategies: (1) GE only, (2) Microsoft only, and (3) half in GE and half in Microsoft. The expected payoff on each of these strategies is the same: $110. For the first two strategies, $100 in either company, the standard deviation is still $10, just as it was before. But for the third strategy, $50 in GE and $50 in Microsoft, the analysis is more complicated. There are four possible outcomes, two for each stock.

To solve the problem, we need to create a table showing all the possibilities, their probabilities, and the associated payoffs (see Table 5.8). We're familiar with possibilities 2 and 3, in which one stock pays off but the other one doesn't, just as in the GE/Texaco example. But possibilities 1 and 4 are new. To assess how risky this investment is, we need to compute its expected payoff and standard deviation of the payoff. The expected payoff is $\frac{1}{4}\$120 + \frac{1}{4}\$110 + \frac{1}{4}\$110 + \frac{1}{4}\$100 = \$110$. Using the information in the table, we can compute the standard deviation of the payoff.

St. dev. of payoff =

$$\sqrt{\frac{1}{4}(\$120-\$110)^2+\frac{1}{4}(\$110-\$110)^2+\frac{1}{4}(\$110-\$110)^2+\frac{1}{4}(\$100-\$110)^2}$$

$$=\sqrt{\frac{1}{4}(\$10)^2+\frac{1}{4}(\$0)^2+\frac{1}{4}(\$0)^2+\frac{1}{4}(\$10)^2}$$

$$=\sqrt{50(\text{dollars})^2}=\$7.1$$

Table 5.8 Payoffs from Investing $50 in Each of Two Stocks
Initial Investment = $100

Possibilities	GE	Microsoft	Total Payoff	Probability
#1	$60	$60	$120	$\frac{1}{4}$
#2	$60	$50	$110	$\frac{1}{4}$
#3	$50	$60	$110	$\frac{1}{4}$
#4	$50	$50	$100	$\frac{1}{4}$

Table 5.9 summarizes the results so that we can compare the three investment strategies. As we have already noted, they all have the same expected payoff, but the strategy of investing in both stocks has a lower standard deviation. By spreading your investment among independently risky investments, you have lowered your overall risk.

Measures of risk other than standard deviation will give us the same result. When you split your investment between the two stocks, 75 percent of the time the payoff is $110 or higher; only 25 percent of the time is the payoff $100. For most people, that prospect is more appealing than a 50 percent probability of getting $100 and a 50 percent probability of getting $120, the odds an investor faces in holding only one stock.

In the real world, there is no reason for an investor to stop diversifying at two stocks. The more independent sources of risk you hold in your portfolio, the lower your overall risk. Using the same numbers in this example—a payoff of either $100 or $120 per $100 investment, with equal probability—we can increase the number of stocks from two to three to four, and the standard deviation of a $100 investment will fall from $7.1 to $5.8 to $5.0. As we add more and more independent sources of risk, the standard deviation becomes negligible. (Appendix 5B explains the algebra behind this statement.)

In summary, spreading the risk is a fundamental investment strategy. As Cervantes put it (and George Shaheen learned), never put all your eggs in one basket. If Shaheen had sold his Webvan stock and invested the proceeds in a portfolio composed of many stocks representative of the stock market as a whole, he would probably still have most of his $280 million.[12] Diversification really works.

Diversification through the spreading of risk is the basis for the insurance business. A large automobile insurer writes millions of policies. It counts on the fact that not everyone will have accidents at the same time, for the risk of any one policyholder

Table 5.9 Results of Possible Investment Strategies: Spreading Risk
Initial Investment = $100

Investment Strategy	Expected Payoff	Standard Deviation
GE only	$110	$10
Microsoft only	$110	$10
$\frac{1}{2}$ and $\frac{1}{2}$	$110	$7.1

[12]Shaheen may not have been able to sell all his stock because of restrictions on sales by company executives. But he surely could have sold some of it and invested the proceeds more prudently.

IN THE NEWS

A Poor Retirement

Financial Times

by Andrew Hill and Elizabeth Wine

December 12, 2001

Enron had thousands of employees, customers, creditors, and shareholders, all of whom have been hit by the energy trader's abrupt collapse into bankruptcy 10 days ago.

But the broadest impact of the company's demise is likely to be felt by millions of working Americans who hold defined-contribution retirement savings accounts, called 401(k)s after the paragraph of the U.S. tax code that brought them into being in the early 1980s.

Under the 401(k) system, companies provide employees with a menu of investment options, generally including several choices of stock and bond mutual funds, as well as money-market funds. Frequently, companies will also offer their own stock for purchase by employees. Workers set aside a percentage of their salary each month to purchase assets for their 401(k) accounts.

Even before the Enron debacle, the benefits of 401(k)s had been called into question. William Bernstein, a principal at an investment management firm, says that high fees, a lack of understanding, and low long-term returns mean that most employees investing through 401(k)s are bound to be disappointed. "I think for the average person, 401(k)s are going to be bad and for a significant minority of people—20 to 30 percent—they are going to be very, very bad," he says.

Enron's plan suffered from particular flaws. The $2.1 billion plan offered employees 18 investment options, including a range of mutual funds and its own stock, but by the end of 2000, 60 percent of the plan was invested in Enron shares—a far higher proportion than most investors would consider prudent diversification. The allocation to Enron rose that high at least partly because Enron matched employee contributions to their own retirement accounts with company stock rather than cash.

Much of the blame for that lack of diversification falls squarely on the shoulders of Enron employees, who were not forced to choose Enron stock. But the problem was compounded by Enron "locking down" its 401(k) plan. Between October 17 and November 19, just as the crisis of confidence in the energy trader became more acute, the company stopped employees from selling Enron shares out of their plan.

Enron shares were priced at $32.20 on October 17, already a fraction of their peak of more than $80 last year. By November 19, the stock had dropped 71 percent to $9.06. With the company in bankruptcy, the shares now trade at less than a dollar.

crashing is independent of the risk of another policyholder crashing. If it writes enough policies, the company can expect a predictable number of accident claims each year. Just as when you toss a coin a million times, you know it will turn up heads 500,000 times and tails 500,000 times, assembling a large enough pool of independent risks isn't risky.

Terms

average, 94
benchmark, 91
diversification, 106
expected return, 95
expected value, 94
hedging, 107
idiosyncratic risk, 104
leverage, 98
mean, 94
payoff, 94

Enron was not alone in offering to match employee contributions with stock instead of cash. About 2,000 U.S. companies do so. That represents only 0.5 percent of 401(k) plans, but the practice affects a greater proportion of workers because the companies are among the biggest and oldest in the U.S. Many of the Fortune 500 companies run their plans in this manner.

Among 140 of the largest companies, the average allocation to company stock in the 401(k) plan is about 35 percent, according to the Committee on the Investment of Employee Benefit Assets, a group of corporate pension plan sponsors. At General Electric, for example, 75 percent of the defined-contribution plan consists of GE stock. About 78 percent of Coca Cola's plan is invested in the company's shares.

[W]hen it comes to taking away or restricting existing investment choices, politicians are cautious. Senator Jeff Bingaman, a Democrat from New Mexico, says retirement policy cannot be "based on the assumption that it is always going to be a sunny day" in the markets but at the same time he strongly believes in "the ability and intelligence of employees to invest their own money" in 401(k)s. David Wray, president of the Profit Sharing/401(k) Council of America (PSCA), a trade association of 1,200 companies that sponsor such plans, concurs: "This is about the long term, it isn't about getting rich in the short term or disappearing because you have a temporary downturn."

That does not convince former Enron employees such as Cregg Lancaster, a 43-year-old father of two, who has lost $100,000 in 401(k) retirement savings as well as his job. "Long-term I'm confident but it's just getting through [to] that long term," he says. Mr Lancaster did not consider selling stock from his 401(k) because it was money he had set aside for retirement. Enron seemed a solid company and a favorite of Wall Street.

It does not yet seem likely that the bankruptcy of Enron will lead to a fundamental restructuring of the U.S. retirement savings system. But whatever the impact, it has taught at least one important lesson to holders of 401(k) plans: retirement nest eggs can go down as well as up.

LESSON OF THE ARTICLE

Don't put all of your eggs in one basket; diversify. Diversification is especially important for retirement savings. But Enron employees lost more than just their investment in Enron stock. They lost their jobs, too. Almost every aspect of their financial well-being was tied up in the same company. The real lesson is that you shouldn't buy stock in the company you work for. If the company goes bankrupt, as Enron did, you're in bad enough shape. You don't want your savings to go down the drain along with your job.

probability, 93
risk, 91
risk averse, 102
risk neutral, 102
risk-free asset, 97
risk-free rate of return, 97
risk premium, 103
spreading risk, 108
standard deviation, 98
systematic risk, 104
value at risk (VaR), 100
variance, 97

Chapter Lessons

1. Risk is a measure of uncertainty about the possible future payoffs of an investment. It is measured over some time horizon, relative to a benchmark.
2. Measuring risk is crucial to understanding the financial system.
   a. To study random future events, start by listing all the possibilities and assign a probability to each. Be sure the probabilities add to one.

b. The expected value is the probability-weighted sum of all possible future outcomes.

c. A risk-free asset is an investment whose future value, or payoff, is known with certainty.

d. Risk increases when the spread (or range) of possible outcomes widens but the expected value stays the same.

e. One measure of risk is the standard deviation of the possible payoffs.

f. A second measure of risk is value at risk, the worst possible loss over a specific time horizon, at a given probability.

3. A risk-averse investor

   a. Always prefers a certain return to an uncertain one with the same expected return.

   b. Requires compensation in the form of a risk premium in order to take risk.

   c. Trades off between risk and expected return: the higher the risk, the higher the expected return risk-averse investors will require for holding an investment.

4. Risk can be divided into idiosyncratic risk, which is specific to a particular business or circumstance, and systematic risk, which is common to everyone.

5. There are two types of diversification:

   a. Hedging, in which investors reduce risk by making investments with offsetting payoff patterns.

   b. Spreading, in which investors reduce risk by making investments with independent payoff patterns.

Problems

1. Consider a game in which a coin will be flipped three times. For each heads you will be paid $100. Assume that the coin has a two-thirds probability of coming up heads.

   a. Construct a table of the possibilities and probabilities in this game.

   b. Compute the expected value of the game

   c. How much would you be willing to pay to play this game?

   d. Consider the effect of a change in the game so that if tails comes up twice in a row, you get nothing. How would your answers to the first three parts of this question change?

2. Begin with the $1,000 investment depicted in Table 5.2.

   a. Recompute the payoff and the payoff times the probability as percentages rather than dollars. Then compute the expected value and standard deviation of the percentage return.

   b. Repeat this exercise using the investment in Table 5.3.

3. An investor is boasting about his ability to obtain a better return than everyone else's. Using the tools in this chapter, how would you explain his performance?

4. You are the founder of IGRO, an Internet firm that delivers groceries.

   a. Describe the idiosyncratic and systematic risks your company faces.

   b. As founder of the company, you own a significant portion of the firm, and your personal wealth is highly concentrated in IGRO shares. What risks do you face, and how should you try to reduce them?

5. Assume that the economy can experience high growth, normal growth, or recession. You expect the following stock-market returns for the coming year under these conditions.

State of the Economy	Probability	Return
High Growth	0.2	+30%
Normal Growth	0.7	+12%
Recession	0.1	–15%

   a. Compute the expected value of a $1,000 investment both in dollars and as a percentage over the coming year.
   b. Compute the standard deviation of the return as a percentage over the coming year.
   c. If the risk-free return is 7 percent, what is the risk premium for a stock-market investment?

6. You can save $5,000 per year from your salary and currently have $15,000 in savings. One year from now you hope to purchase a house for $100,000. To obtain a mortgage you can afford, you will need a down payment equal to 20 percent of the purchase price of the house. You have two possible investments available. The first is a risk-free bond that pays 5 percent; the second is the stock-market investment described in problem 5. How would you decide which investment to make?

7. An investment advisor offers you an opportunity to buy a financial instrument with the following payoffs.

State of the Economy	Probability	Return
High Growth	0.2	–10%
Normal Growth	0.7	+4%
Recession	0.1	+8%

   Assuming that you also have available the opportunity in problem 5, is this investment valuable to you? Why or why not?

8. You are a typical American investor. An insurance broker calls and asks if you would be interested in an investment with a high payoff if the annual Indian monsoons are less damaging than normal. If damage is high, you will lose your investment. On calculating the expected return, you realize that it is roughly the same as that of the stock market. Is this opportunity valuable to you? Why or why not?

9. Among the many consequences of Enron's bankruptcy in 2001 was the loss of savings Enron employees suffered. Roughly 47 percent of employee pension funds

were invested in Enron stock. Why was this investment so risky? How could the risk have been reduced?

10. In Table 5.4, what interest rate is consistent with your $750 threshold? That is, what maximum interest rate will you accept to insure against the possibility of interest rates and payments rising above your means?

11. Car insurance companies eliminate risk (or come close) by selling a large number of policies. Explain how they do this.

12. One morning you are intrigued by an e-mail offering you income insurance. A company has formed to guarantee college students a fixed income for the rest of their lives. As payment, you have to agree to give them your salary. Is this the sort of insurance people are likely to buy? Will the company be able to stay in business?

13. Mortgages increase the risk faced by homeowners.
    a. Explain how.
    b. What happens to the homeowner's risk as the down payment on the house rises from 10 percent to 50 percent?

14. Banks pay substantial amounts to monitor the risks that they take. One of the primary concerns of a bank's risk managers is to compute the value at risk. Why is value at risk so important for a bank (or any financial institution)?

15. How much would you be willing to pay for the investment described in Table 5.3? Explain why it is more or less than $1,050.

Appendix 5A

A Quick Test to Measure Your Risk Tolerance*

The following quiz is adapted from one prepared by the T. Rowe Price group of mutual funds. It can help you figure out how comfortable you are with varying degrees of investment risk. Other things being equal, your risk tolerance is a useful guide for deciding how heavily you should weight your portfolio toward low- or high-risk investments.

1. You are the winner of a TV game show. Which prize would you choose?
   - $2,000 cash (1 point).
   - A 50 percent chance to win $4,000 (3 points).
   - A 20 percent chance to win $10,000 (5 points).
   - A 2 percent chance to win $100,000 (9 points).
2. You are down $500 in a poker game. How much more would you be willing to bet to win the $500 back?
   - $500 (6 points).
   - $250 (4 points).
   - $100 (2 points).
   - Nothing. You'll cut your losses and quit now (1 point).
3. A month after you invest in a stock, it suddenly goes up 15 percent. With no further information, what would you do?
   - Hold it, hoping for further gains (3 points).
   - Sell it and take your gains (1 point).
   - Buy more. It will probably go higher (4 points).
4. Your investment suddenly goes down 15 percent. Its fundamentals still look good. What would you do?
   - Buy more. If it looked good at the original price, it looks even better now (4 points).
   - Hold on and wait for it to come back (3 points).
   - Sell it to avoid losing even more (1 point).
5. You are a key employee in a start-up company. You can choose one of two ways to take your year-end bonus. Which would you pick?
   - $1,500 cash (1 point).
   - Company stock options that could bring you $15,000 next year if the company succeeds but will be worthless if it fails (5 points).

Your score: ______

*From Jack Kapoor, Les Dlabay, Robert J. Hughes, *Personal Finance* (New York: McGraw-Hill, 2004).

Scoring

5–18 points: You are a conservative investor. You prefer to minimize financial risks. The lower your score, the more cautious you are. When you choose investments, look for high credit ratings, well-established records, and an orientation toward stability. In stocks, bonds, and real estate, focus on income.

19–30 points: You are a less conservative investor. You are willing to take more chances in pursuit of greater rewards. The higher your score, the bolder you are and the more risk you are willing to take. You may want to consider bonds with higher yields and lower credit ratings, the stocks of newer companies, and real estate investments that use mortgage debt.

Appendix 5B

The Mathematics of Diversification

With a small amount of mathematics, we can show how diversification reduces risk. Let's begin with two investments in GE and Texaco. We'll label the payoffs to these investments x and y. If x is the payoff to buying GE stock, then it must equal either \$120 or \$100, each with a probability of one-half (see Table 5.6). Then y is the payoff to buying Texaco stock.

Hedging Risk

In the chapter, we considered splitting our investment between GE and Texaco. If x and y are the payoff from holding GE and Texaco, respectively, then the payoff on the investment is

$$\text{Investment payoff} = \tfrac{1}{2}x + \tfrac{1}{2}y \qquad \text{(A1)}$$

What is the variance of this payoff? (Since the standard deviation is the square root of the variance, the two must move together—a lower variance means a lower standard deviation—so we can skip the standard deviations.) In general, the variance of any weighted sum $ax + by$ is

$$\text{Var}(ax + by) = a^2\text{Var}(x) + b^2\text{Var}(y) + 2ab\,\text{Cov}(x,y), \qquad \text{(A2)}$$

where *Var* is the variance and *Cov* is the covariance. While the variance measures the extent to which each payoff moves on its own, the covariance measures the extent to which two risky assets move together. If the two payoffs rise and fall together, then the covariance will be positive. If one payoff rises while the other falls, then the covariance will be negative.

It is useful to express these quantities symbolically. Assume that p_i is the probability associated with a particular outcome x_i. Then the expected value of x is the probability-weighted sum of the possible outcomes.

$$\text{Expected value of } x = E(x) = \bar{x} = \sum_i p_i x_i. \qquad \text{(A3)}$$

As described in the chapter (page 97), the variance of x is the probability-weighted sum of the squared deviations of x from the expected value.

$$\text{Variance of } x = \text{Var}(x) = \sigma_x^2 = \sum_i p_i (x_i - \bar{x})^2. \qquad \text{(A4)}$$

The covariance of x and y is defined analogously as

$$\text{Covariance of } x \text{ and } y = \text{Cov}(x,y) = \sigma_{x,y} = \sum_i p_i\,(x_i - \bar{x})\,(y_i - \bar{y}). \qquad \text{(A5)}$$

In our GE/Texaco examples $a = b = \frac{1}{2}$, so

$$\text{Var(Investment payoff)} = \tfrac{1}{4}\text{Var}(x) + \tfrac{1}{4}\text{Var}(y) + \tfrac{1}{2}\text{Cov}(x,y) \qquad \text{(A6)}$$

We know from Table 5.7 that the expected payoff to GE and Texaco is \$110 each and the standard deviation is \$10 as well. The variance is the standard deviation squared, so it is 100. What about the covariance? We can compute it easily from Table 5.6:

$$\text{Cov(Payoff on GE and Texaco)} = \tfrac{1}{2}(100 - 110)(120 - 110) + \tfrac{1}{2}(120 - 110)(100 - 110) = -100 \qquad \text{(A7)}$$

Substituting this value into the formula for the variance of the investment payoff, we get

$$\text{Var(Investment payoff)} = \tfrac{1}{4}(100) + \tfrac{1}{4}(100) - \tfrac{1}{2}(100) = 0 \qquad \text{(A8)}$$

The fact that the covariance is negative means that the variance, or risk, in a portfolio containing both GE and Texaco stock is lower than the risk in a portfolio containing one or the other. The stocks act as hedges for each other.

Spreading Risk

Showing how spreading reduces risk is a bit more complex. Let's consider spreading our investment between GE and Microsoft. Again, the variance of the investment payoff depends on the variances of the individual stock payoffs and on their covariance. But here we must assume that the covariance between the GE and Microsoft payoffs is zero. That is, they are independent of each other. As before, each stock has a variance of 100, so the variance of a portfolio that is split half and half is

$$\text{Var(Investment payoff)} = \tfrac{1}{4}(100) + \tfrac{1}{4}(100) = 50, \qquad \text{(A9)}$$

and the standard deviation is 7.1 (see Table 5.9).

This result suggests that individual stocks or groups of stocks with independent payoffs are potentially valuable, as they will reduce risk. Let's consider an arbitrary number of independent investments, each with the same individual variance. What is the variance of an equally weighted portfolio of these investments? Assume that the number of investments is n, each with the same expected payoff, $\bar{x}$, and the same variance, σ_x^2. We hold $1/n$ of our portfolio in each stock, so the expected payoff is

$$\text{Expected payoff} = \frac{1}{n}\sum_{i=1}^{n} x_i = \bar{x} \qquad \text{(A10)}$$

Since the payoff on each stock is independent of all the rest, the covariances are all zero. So the variance is

$$\text{Variance of payoff} = \left(\frac{1}{n}\right)^2 \sum_{i=1}^{n} \sigma_x^2 = \frac{\sigma_x^2}{n} \qquad \text{(A11)}$$

That is, the variance of the payoff on a portfolio of n independent stocks is the variance divided by n. Most important, as n increases, the variance declines, so when the value of n is very large, the variance is essentially zero.

In summary, spreading exposure to risk among a wide range of independent risks reduces the overall risk of a portfolio. As we saw in this chapter, this is the strategy that insurance companies use. While the payoff to an individual policyholder is highly uncertain, the payoffs to a large group of policyholders are almost unrelated. By selling millions of insurance policies, the company reduces the payoffs it must make and simply pays the expected value.