GARP
DIGITAL LIBRARY

"Investors and Risk Management"

Chapter 2 of

Risk Management and Derivatives

By René Stulz

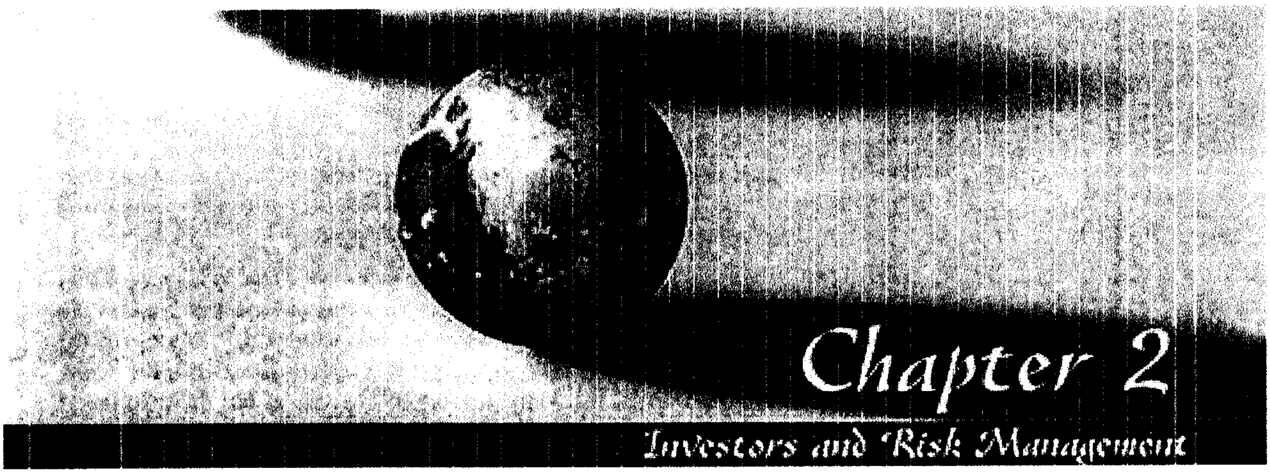

Chapter 2
Investors and Risk Management

Chapter 2 Objectives

At the end of this chapter, you will:

1. Understand expected return and volatility for a security and a portfolio.
2. Know how to use the normal distribution to obtain the probability of ranges of returns for securities.
3. Be able to evaluate the risk of a security in a portfolio.
4. Know how the capital asset pricing model is used to obtain the expected return of a security and to compute the present value of cash flows.
5. Know how hedging affects firm value in perfect financial markets.
6. Understand how investors evaluate risk management policies of firms in perfect financial markets.

How do investors evaluate the risk management policies of firms in which they invest? More specifically, when do investors want a firm in which they hold shares to spend money to reduce the volatility of its stock price?

First, we have to know how investors decide to invest their money, and how the risk management policies of firms affect the riskiness of their investments. To do this, we describe some of the tools used to evaluate the distribution of the returns of securities and portfolios. Ignoring derivatives for the time being, investors have two powerful risk management tools that enable them to invest their wealth with a level of risk that is optimal for them. The first tool is **asset allocation**, which specifies how wealth is allocated across types of securities or asset classes. The second tool is **diversification**. A portfolio's degree of diversification is the extent to which the funds invested are distributed across securities to lessen the dependence of the portfolio's return on the return of individual securities.

This chapter shows that with these risk management tools investors do not need a firm to manage risk to help them achieve their optimal risk-return trade-off. Consequently, they benefit from a firm's risk management policy only if that policy increases the present value of the cash flows the firm expects to generate. The next chapter will show how a firm can use risk management to increase that present value.

2.1. Evaluating the risk and the return of individual securities and portfolios

Suppose an investor named John Smith has wealth of $100,000 that he wants to invest in equities for one year. His broker recommends two companies, IBM and XYZ. John knows about IBM, but has never heard of XYZ. He decides that first he wants to understand what his wealth would amount to after putting all his wealth in IBM shares for one year.

The return of a stock per dollar invested over a period of time is the total gain from holding the stock divided by the stock price at the beginning of the period. If the stock price is $100 at the beginning of the year, the dividend payments are $5, and the stock price appreciates by $20 during the year, the return per dollar invested or decimal return is (20 + 5)/100, or 0.25. Alternatively, we can express the return in percentage, so that a decimal return of 0.25 is a return of 25 percent. Unless we mention otherwise, returns are decimal returns.

For each dollar invested in the stock, John has one dollar plus the return of the stock at the end of the year. Since he puts all his wealth in IBM, his wealth at the end of the year is his initial wealth times one plus the return of IBM, or (Initial wealth)(1 + Return of IBM). In this example, John's wealth at the end of the year is $100,000(1 + 0.25), or $125,000. We first discuss how to figure out how likely various return outcomes are for a stock, and then do the same for a portfolio. Most readers may be familiar with these materials, but we include them because the concepts are basic to an understanding of derivatives and risk management.

Throughout the analysis in this chapter, we assume that the frictions that affect financial markets are unimportant. More specifically, we assume that there are no taxes, no transaction costs, no costs to writing and enforcing contracts, no restrictions on investments in securities, no differences in information across investors, and that investors take prices as given because they are too small to affect prices. Financial economists call markets that satisfy these assumptions **perfect financial markets**. Real-world financial markets are not perfect financial markets, but we make this assumption because it allows us to avoid distractions in discussing important concepts and to clarify the conditions under which financial risk management can increase firm value. Later on, we take into account financial markets imperfections and build on our understanding of perfect financial markets.

2.1.1. The distribution of the return of IBM

We first describe the concepts of return distribution, expected return, and return variance. We then show how to use the distribution of the return to infer the probability of various return outcomes for a stock. Finally, we address the implications of past returns for future returns.

2.1.1.A. Return distribution, expected return, and return variance

Because stock returns are uncertain, John has to figure out which outcomes are likely to occur and which are not. To do this, he uses basic statistical tools. The return of IBM is a random variable—we do not know what its value will be until that value is realized. A probability distribution provides a quantitative measure of the likelihood of the possible outcomes or realizations of a random variable by assigning probabilities to these outcomes. The statistical tool used to measure the likelihood of various returns for a stock is called the stock's **return probability distribution**. The most common probability distribution is the **normal distribution**. There is substantial empirical evidence that, for many purposes, the normal distribution provides a good but not perfect approximation of the true, unknown, distribution of stock returns.

With the normal distribution, all there is to know about the distribution of a stock's return is given by the expected return of the stock and by its variance. The **expected value** of a random variable is a probability-weighted average of all the possible distinct outcomes of that variable. Each distinct outcome has a probability, and all probabilities add up to one. For example, if the probability distribution of a stock specifies that it can have only one of two returns, 0.1 with probability 0.4 and 0.15 with probability 0.6, its expected return is $0.4 \times 0.1 + 0.6 \times 0.15$, or 0.13 in decimal form. IBM's **expected return** is the average return John would earn if next year were repeated over and over, each time yielding a different return drawn from the return distribution of IBM. Everything else equal, the higher the expected return, the better off the investor. If y is a random variable, we denote its expected value by $E(y)$.

The **variance** of a random variable is a quantitative measure of how the realizations of the random variable are distributed around their expected value; it provides a measure of risk. More precisely, it is the expected value of the square of the difference between the realizations of a random variable and its expected value, $E[y - E(y)]^2$. Using our example of a return of 0.10 with probability 0.4 and a return of 0.15 with probability 0.6, the decimal variance of the return is $0.4(0.10 - 0.13)^2 + 0.6(0.15 - 0.13)^2$ or 0.0006. For returns, the units of the variance are returns squared. The square root of the variance, however, is in the same units as the returns and is called the **standard deviation**. In finance, the standard deviation of returns is generally called the **volatility** of returns. We write Var(y) and Vol(y), respectively, for the variance and the volatility of random variable y. In our example, the square root of 0.0006 is 0.0245. Since the volatility is in the same units as the returns, we can use a volatility expressed as 2.45 percent. As returns are spread farther from the expected return, volatility increases. For example, if instead of having returns of 0.10 and 0.15 we have returns of 0.025 and 0.20, the expected return is unaffected but the volatility becomes 8.57 percent instead of 2.45 percent. Similarly, if IBM's return volatility is low, its return is likely to be close to its expected value, so that a return substantially greater or less than the expected return would be surprising. As IBM's volatility increases, a return close to the expected return becomes less likely.

Since investors prefer more to less, an increase in their expected wealth and hence in the expected return of their investments is good for them. However, investors are typically risk-averse, so, keeping the expected return on their wealth constant, they would prefer the volatility of the return on their wealth to be lower.

2.1.1.B. Using the return distribution to infer the likelihood of various return outcomes The **cumulative distribution function** of a random variable y specifies, for any number Y, the probability that the realization of the random variable will be no greater than Y. We denote the probability that the random variable y has a realization no greater than Y as prob($y \leq Y$). For IBM, a reasonable estimate of the stock return volatility is 30 percent. With an expected return of 13 percent and a volatility of 30 percent, we can draw the cumulative distribution function for the return of IBM as plotted in Figure 2.1. For a given return, the function specifies the probability that the return of IBM will not exceed that return.

To use the cumulative distribution function, we choose a value on the horizontal axis, say 0 percent. The corresponding value on the vertical axis tells us the probability that IBM will earn less than 0 percent is 0.33. In other words, there is a 33 percent chance that over one year, IBM will have a negative return.

The easiest way to compute a probability is to use a spreadsheet program such as Excel. Box 2.1, Computing a probability using Excel, shows how to do this. Suppose John is worried about making losses. Using the normal distribution, we can tell him that there is a 33 percent chance he will lose money. This probabili-

Box 2.1 **Computing a probability using Excel**

The NORMDIST function of Excel is used to obtain probabilities for a normal distribution. Suppose we want to know how likely it is that IBM will earn less than 10 percent over one year. To get the probability that the return will be less than 10 percent, we choose x = 0.10. The mean is 0.13 and the standard deviation is 0.30. We finally write TRUE in the last line to obtain the cumulative distribution function. The result is 0.46. This number means that there is a 46 percent chance that the return of IBM will be less than 10 percent over a year.

Computing a probability using Excel *Figure 2.1*

The expected return of IBM is 13 percent and its volatility is 30 percent. The horizontal line corresponds to a probability of 0.05. The cumulative probability function of IBM crosses that line at a return almost twice as high as the cumulative probability function of the riskier stock. There is a 5 percent chance that IBM will have a lower return than the one corresponding to the intersection of the IBM cumulative distribution function and the horizontal line, which is a return of -36 percent. There is a 5 percent chance that the stock with twice the volatility of IBM will have a return lower than -0.66 percent.

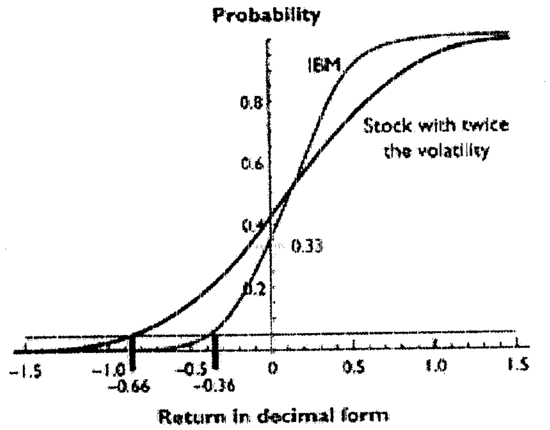

ty depends on the expected return. As the expected return of IBM increases, the probability of making a loss falls.

One concern John could have is that his wealth might not be sufficient to pay for living expenses. Suppose he needs to have $50,000 to live on at the end of the year. By putting all his wealth in a stock, he knows that he takes the risk that he will have less than that amount at the end of the year, but he wants the probability of that outcome to be less than 0.05. With the cumulative normal distribution with an expected return of 13 percent and a volatility of 30 percent, the probability of a 50 percent loss is 0.018. John can therefore invest in IBM given his objective of making sure that there is a 95 percent chance that he will have at least $50,000 at the end of the year.

Suppose John wants to understand how likely it is that his portfolio will have a value between $50,000 and $100,000 at the end of the year. We know that the probability that the portfolio will be worth less than $50,000 is 0.018 and the probability that the portfolio will be worth less than $100,000 is 0.33. The probability that the portfolio will be worth less than $100,000 is the sum of two probabilities: the probability that the portfolio is worth less than $50,000 and the probability that the portfolio is worth more than $50,000 but less than $100,000. The sum of these two probabilities is 0.33. Subtracting from 0.33 the probability that the portfolio will be worth less than $50,000, we get the probability that the portfolio will be worth more than $50,000 but less than $100,000:

0.33 − 0.018, or 0.312. The probability of 0.312 is the sum of the probability of all the possible values the portfolio could take between $50,000 and $100,000.

The **probability density function** tells us what the probabilities of these various portfolio values are. If a random variable takes discrete values, the probability density function tells us the probability of each of the values that the random variable can take. With the normal distribution, the random variable is continuous—there are many possible values over any range of numbers. In this case, the probability density function tells us the probability that the random variable will take a value within an infinitesimally small range of its possible values--it gives us the increase in prob($x \leq X$) as X increases by an infinitesimal amount.

In the case of IBM, we see that the cumulative distribution function first increases slowly, then more sharply, and finally again slowly. This explains why the probability density function of IBM shown in Figure 2.2 first has a value close to zero, increases to reach a peak, and then falls again. This bell-shaped probability density function is characteristic of the normal distribution. Note that this bell-shaped function is symmetric around the expected value of the distribution. For comparison, the figure also shows the distribution of the return of a security that has twice the volatility of IBM but the same expected return. The distribution of the more volatile security has more weight in the tails and less around the mean than IBM, implying that outcomes substantially away from the mean are more likely.

The distribution of the more volatile security shows a limitation of the normal distribution for simple returns: It has returns worse than −100 percent. Because

Figure 2.2 Normal density function for IBM assuming an expected return of 13 percent and a volatility of 30 percent and of a stock with the same expected return but twice the volatility

This figure shows the probability density function of the one-year return of IBM assuming an expected return of 13 percent and a volatility of 30 percent. It also shows the probability density function of the one-year return of a stock that has the same expected return but twice the volatility of return of IBM.

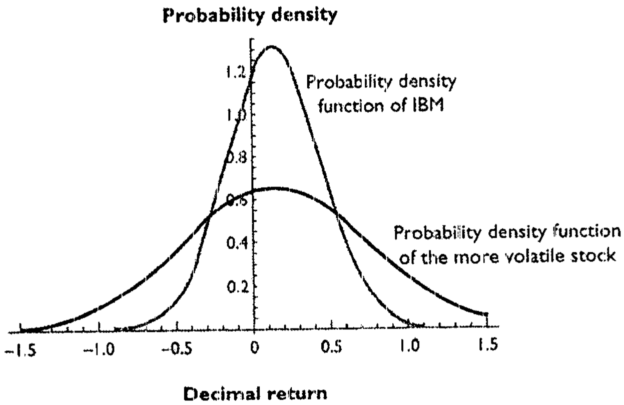

stocks have limited liability, the most one can lose owning a stock is what one paid for it, corresponding to a simple return of –100 percent. In general, this limitation is not important in that the probability of such a return is very small.

2.1.2. The distribution of the return of a portfolio

To be thorough, John wants to consider XYZ. He first wants to know if he would be better off investing \$100,000 in XYZ rather than in IBM. He finds out that the expected return of XYZ is 26 percent and the return volatility is 60 percent, so that XYZ has twice the expected return and twice the volatility of IBM. Using volatility as a summary risk measure, XYZ is riskier than IBM. The probability that the price of XYZ will fall by 50 percent is 0.102. Consequently, John cannot invest all his wealth in XYZ if he wants his probability of losing \$50,000 to be at most 0.05.

Since XYZ has a high expected return compared to IBM, though, John wants to consider investing something in XYZ, forming a portfolio of the two stocks. Section 2.1.2.A presents the computation of the return and the expected return of the portfolio, while section 2.1.2.B shows how to compute and use the return volatility of the portfolio.

2.1.2.A. The return and expected return of a portfolio The return of a portfolio is the weighted average of the return of the securities in the portfolio, where the weight for a security is the fraction of the portfolio invested in that security. The fraction of the portfolio invested in a security is called the **portfolio share** (or portfolio weight) of that security. Suppose John puts \$75,000 in IBM and \$25,000 in XYZ. The portfolio share of IBM is \$75,000/\$100,000, or 0.75. Portfolio shares sum to one since the entire portfolio must be invested. A negative portfolio share corresponds to a short sale. With a **short sale,** an investor borrows shares from a third party and sells them. With our assumption of perfect financial markets, the investor can then use the proceeds from the sale fully. To close the short-sale position, the investor must buy shares and deliver them to the lender. If the share price increases, the investor loses because he has to pay more for the shares he delivers than he received for the shares he sold.

Using w_i for the portfolio share of security i in a portfolio with N securities and R_i for the return on security i, the portfolio return is:

$$\sum_{i=1}^{N} w_i R_i = \text{Portfolio return} \tag{2.1}$$

If the realized return on IBM is 20 percent and the realized return on XYZ is –10 percent, applying formula (2.1) the decimal return of the investor's portfolio is:

$$0.75(0.20) + 0.25(-0.10) = 0.125 \tag{2.2}$$

With this return, the wealth of the investor at the end of the year is 100,000 x (1 + 0.125), or \$112,500.

At the start of the year, John wants to compute the expected return of the portfolio and the return volatility of the portfolio for different choices of portfolio shares to help allocate his wealth between IBM and XYZ shares. The portfolio weights are taken as given and therefore are treated as constants. The expected return of a portfolio is therefore the portfolio share weighted average of the expected return of the securities in the portfolio:[1]

$$\sum_{i=1}^{N} w_i E(R_i) = E(R_p) = \textbf{Porfolio expected return} \tag{2.3}$$

Applying this formula using an expected return for IBM of 13 percent and an expected return for XYZ of 26 percent, the expected decimal return of the portfolio with a portfolio share of 0.75 in IBM and 0.25 in XYZ is:

$$0.75 \times 0.13 + 0.25 \times 0.26 = 0.1625 \tag{2.4}$$

The expected wealth of the investor at the end of the year is E[(Initial wealth) x (1 + Portfolio return)]. Since (see footnote 1) the expectation of the product of a constant with a random variable is equal to the constant times the expectation of the random variable, the expectation of the investor's wealth at the end of the year is Initial wealth x [1 + E(Portfolio return)]. John therefore expects his wealth to be 100,000 x (1 + 0.1625), or $116,250, at the end of the year.

2.1.2.B. The volatility of the return of a portfolio An investor naturally wants to be able to compare the volatility of the stock portfolio to the volatility the portfolio would have if the entire amount were invested in one stock. To compute the volatility of a portfolio, it is best to compute the variance of the portfolio return first and then take its square root to get the volatility to avoid cumbersome square roots. To compute the variance of the portfolio return, we first need to review two properties of the variance.

The first property is that the variance of a constant times a random variable is the constant squared times the variance of the random variable. This implies that $\mathrm{Var}(w_i R_i) = w_i^2 \mathrm{Var}(R_i)$. The portfolio return is a weighted sum of returns. Consequently, to compute the variance of the portfolio return, we have to compute the variance of a sum. If a and b are random variables, to obtain the variance of $a + b$ we have to compute $E[a + b - E(a + b)]^2$. Remember that the square of a sum of two terms is the sum of each term squared plus two times the cross-product of the two numbers (the square of 5 + 4 is $5^2 + 4^2 + 2 \times 5 \times 4$, or 81). Consequently:

$$\begin{aligned} \mathrm{Var}(a+b) &= E[a + b - E(a + b)]^2 \\ &= E[a - E(a) + b - E(b)]^2 \\ &= E[a - E(a)]^2 + E[b - E(b)]^2 + 2E[a - E(a)][b - E(b)] \\ &= \mathrm{Var}(a) + \mathrm{Var}(b) + 2\mathrm{Cov}(a, b) \end{aligned} \tag{2.5}$$

The bold term is the covariance between a and b, denoted by Cov(a, b), which measures how a and b move together. The covariance is the expected value of the

[1] To compute the portfolio's expected return, we use two properties of expectations. First, the expected value of the product of a random variable and a constant is equal to the constant times the expected value of the random variable. If $E(w_iR_i)$ is the expected return on security i times its portfolio share, this property of expectations implies that $E(w_iR_i) = w_iE(R_i)$. Second, the expected value of a sum of random variables is simply the sum of the expected values of the random variables. This second property implies that if the portfolio has only securities 1 and 2, $E(w_1R_1 + w_2R_2) = E(w_1R_1) + E(w_2R_2)$, which is equal to $w_1E(R_1) + w_2E(R_2)$ because of the first property.

product of the deviations of two random variables from their mean: $E[a - E(a)][b - E(b)]$. The covariance can take negative as well as positive values. Its value increases as a and b are more likely to exceed their expected values together. If the covariance is zero, the fact that a exceeds its expected value provides no information about whether b exceeds its expected value also. Equation (2.5) shows the second important property of variances for the case of two random variables: The variance of a sum of random variables is the sum of the variances of the random variables plus twice the covariance of each pair of random variables.

The covariance is closely related to the correlation coefficient. The correlation coefficient takes values between –1 and +1. If a and b have a correlation coefficient of 1, they move in lockstep in the same direction. If the correlation coefficient is –1, they move in lockstep in the opposite direction. Finally, if the correlation coefficient is zero, a and b are independent if they are normally distributed. Denote by Corr(a, b) the correlation between a and b. If one knows the correlation coefficient, one can obtain the covariance by using the formula:

$$\mathbf{Cov}(a, b) = \mathbf{Corr}(a, b) \times \mathbf{Vol}(a) \times \mathbf{Vol}(b) \qquad (2.6)$$

The variance of $a + b$ increases with the covariance of a and b since an increase in the covariance makes it less likely that an unexpectedly low value of a is offset by an unexpectedly high value of b. In the special case where a and b have the same volatility, $a + b$ has no risk if the correlation coefficient is –1 because a high realization of one of the random variables is always exactly offset by a low realization of the other. Note that if a and b are the same random variables, they have a correlation coefficient of +1, so that Cov(a, b) is Cov(a, a) = 1 x Vol(a) x Vol(a), which is Var(a) since the square of the volatility of a is its variance, so that the covariance of a random variable with itself is its variance.

From what we have just seen, John does not have enough information to compute the variance of the portfolio if he knows just the variance and the portfolio weights of the securities in the portfolio. He must also know how the securities in the portfolio covary. More generally, therefore, the formula for the variance of the return of a portfolio is:[2]

$$\sum_{i=1}^{N} w_i^2 Var(R_i) + \sum_{i=1}^{N}\sum_{j\neq i}^{N} w_i w_j Cov(R_i, R_j)$$
$$= \textbf{Variance of portfolio return} \qquad (2.7)$$

Applying equation (2.7) to the portfolio of IBM and XYZ, we need to know the covariance between the returns of the two securities. If the correlation coefficient between the two securities is 0.5, the covariance is 0.5(0.30)(0.60), or 0.09, and the variance is:

$$0.75^2(0.3^2) + 0.25^2(0.6^2) + 2(0.25)(0.75)(0.5)(0.3)(0.6) = 0.11 \qquad (2.8)$$

[2] This formula is obtained in the following way in the case of a portfolio with two assets, assets 1 and 2. Using the formula for the variance of a sum, the variance of the portfolio, Var($w_1R_1 + w_2R_2$) is equal to Var(w_1R_1) + Var(w_2R_2) + 2Cov(w_1R_1, w_2R_2). Since we saw that if k is a constant and a a random variable, the variance of ka is k^2Var(a), Var(w_1R_1) + Var(w_2R_2) + 2Cov(w_1R_1, w_2R_2) = w_1^2Var(R_1) + w_2^2 Var(R_2) + 2 w_1w_2Cov(R_1, R_2).

The volatility of the portfolio is the square root of 0.11, which is 0.3317. By investing less in IBM and more in a stock that has twice the volatility of IBM, John can increase his expected return from 13 to 16.25 percent, but in doing so he increases the volatility of his portfolio from 30 to 33.17 percent. Since the portfolio has a higher expected return than IBM but also a higher volatility, we cannot determine a priori which of the three possible investments the investor prefers (investing in IBM, XYZ, or the portfolio). We know that John would prefer the portfolio if it had a higher expected return than IBM and less volatility, but this is not the case. An investor who is risk-averse is willing to give up some expected return in exchange for less risk. If John dislikes risk sufficiently, he prefers IBM to the portfolio because IBM has less risk even though it has less expected return. By altering portfolio shares, the investor can create many different portfolios that differ in their risk and expected return.

2.2. Diversification, asset allocation, and expected returns

In section 2.2.1, we determine how diversification affects the distribution of the return of a portfolio. In section 2.2.2, we examine asset allocation when there is a risk-free asset. In section 2.2.3, we show how investors measure risk when they hold a diversified portfolio. Finally, in section 2.2.4, we explain how required expected returns are determined when investors care only about the expected return and the volatility of their portfolio. Once again, the reader may be familiar with the materials presented in this section, but a review is necessary to understand when risk management creates value.

2.2.1. Diversification and the return of a portfolio

The impact on the volatility of the investor's portfolio return of investing in XYZ depends on the correlation coefficient between XYZ and IBM. Figure 2.3 shows that the volatility of the portfolio with portfolio share of 0.25 in XYZ and 0.75 in IBM increases directly with the correlation coefficient between XYZ and IBM. Equation (2.1), however, shows that the expected return of a portfolio does not depend on the covariances and variances of the securities that constitute the portfolio. Consequently, it follows from Figure 2.3 that if the return correlation coefficient is zero, the volatility of the investor's portfolio falls from 30 to 26 percent without a change in expected return as she invests 25 percent of her wealth in XYZ instead of all of it in IBM, making her unambiguously better off.

That John should want to invest in XYZ despite its high volatility is made clear in Figure 2.4. The graph represents all the combinations of expected return and volatility that can be obtained by investing in IBM and XYZ when the correlation is zero. Such a graph drawn for any correlation coefficient between IBM and XYZ would have a similar shape. The upward-sloping part of the curve drawn in Figure 2.4 is called the efficient frontier. The investor wants to choose portfolios on the efficient frontier because, for each volatility, there is a portfolio on the efficient frontier that has a higher expected return than any other portfolio with the same volatility.

By choosing portfolios on the efficient frontier instead of holding only shares of IBM, the investor benefits from diversification. Diversification is such a good risk management tool that sometimes it can eliminate risk completely. To see this,

Portfolio volatility and correlation *Figure 2.3*

This figure shows the volatility of a portfolio with a portfolio share of 0.75 in IBM and 0.25 in XYZ when IBM has a volatility of 30 percent and XYZ has a volatility of 60 percent as a function of the correlation coefficient between IBM and XYZ.

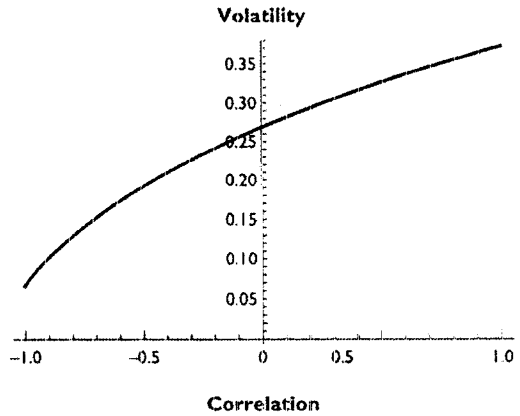

Efficient frontier without a riskless asset *Figure 2.4*

The function represented in the figure gives all the combinations of expected return and volatility that can be obtained with investments in IBM and XYZ. The point where volatility is the smallest has an expected return of 15.6 percent and a standard deviation of 26.83 percent. The upward-sloping part of the curve is the efficient frontier. The portfolio on the efficient frontier that has the same volatility as a portfolio wholly invested in IBM has an expected return of 18.2 percent.

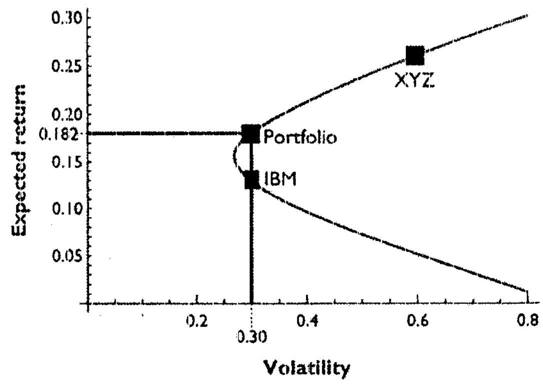

suppose that an investor can choose to invest among many uncorrelated securities that all have the same volatility and the same expected return as IBM: a volatility of 30 percent and an expected return of 13 percent. Dividing one's wealth among all these uncorrelated securities has no impact on the expected return, because all the securities have the same expected return. Using the formula for the variance of a portfolio, equation (2.7), however, we find that the volatility of the portfolio is:

$$\text{Volatility of portfolio} = \left[\sum_{i=1}^{N} (1/N)^2 (0.3)^2\right]^{0.5} = \frac{0.3}{\sqrt{N}} \qquad (2.9)$$

Applying this result, we find that for $N = 10$, the volatility is 9 percent; for $N = 100$ it is 3 percent; and for $N = 1000$ it is less than 1 percent. As N is increased further, the volatility becomes infinitesimal.

In other words, by holding uncorrelated securities, one can eliminate portfolio volatility if one holds sufficiently many of these securities. Risk that disappears in a well-diversified portfolio is called **diversifiable risk**. In our example, all of the risk of each security becomes diversifiable as N increases.

In the real world, securities tend to be positively correlated because changes in aggregate economic activity affect most firms. News of the onset of a recession, for instance, is generally bad news for almost all firms. As a result, we cannot eliminate risk through diversification but we can reduce it. The risk that cannot be eliminated through diversification, the risk that remains, is often called **systematic risk**.

2.2.2. Asset allocation when there is a risk-free asset

So far, we have assumed that the investor forms the portfolio holding shares of two companies, IBM and XYZ. We now consider portfolio choice when there is also an asset that has no risk over the investment horizon of the investor. An example of such an asset is a Treasury bill (T-bill). T-bills are **zero-coupon bonds**. For zero-coupon bonds, the interest payment comes in the form of capital appreciation of the bond. Since T-bills have no default risk, they have a sure return if held to maturity. Box 2.2, Treasury bills, shows how they are quoted and how one can use a quote to obtain a yield.

John can decrease the volatility of year-end wealth by investing some fraction in risk-free bonds, perhaps half in risk-free bonds and the other half in the portfolio of risky assets with the lowest volatility. Assume that the bonds earn 5 percent over the year. This minimum-volatility portfolio of risky assets has an expected return of 15.6 percent and a standard deviation of 26.83 percent. With this asset allocation, John's portfolio would have a volatility of 13.42 percent and an expected return of 10.3 percent.

The efficient frontier in Figure 2.4 that was formed using only risky stocks is called the efficient frontier of risky assets. All combinations of the minimum-volatility portfolio and the risk-free asset lie on a straight line that intersects this efficient frontier of risky assets at the minimum-volatility portfolio. Figure 2.5 shows this straight line. Portfolios on the straight line to the left of the minimum-volatility portfolio have positive investments in the risk-free asset. In contrast,

Treasury Bills *Box 2.2*

T-bills are securities issued by the U.S. government that mature in one year or less. They pay no coupon, so that the investor's dollar return is the difference between the price paid on sale or at maturity and the price paid on purchase. Suppose that T-bills maturing in one year sell for \$95 per \$100 of face value. This means that the holding period return computed annually is 5.26 percent (100 x 5/95) because each investment of \$95 returns \$5 of interest after one year.

T-bills are quoted on a bank discount basis. The price of a T-bill is quoted as a discount:

$$D_t(t + n/365) = (360/n)[100 - P_t(t + n/365)]$$

where $D_t(t + n/365)$ is the discount for a T-bill that matures in n days and $P_t(t + n/365)$ is the price of the same T-bill. The bank discount method uses 360 days for the year. Suppose that the price of a 90-day T-bill, $P_t(t + 90/365)$, is \$98.5. In this case, the T-bill would be quoted at a discount of 6.00. From the discount rate, one can recover the price using the formula:

$$P_t(t + n/365) = 100 - (n/360)D_t(t + 365/n)]$$

For our example, we have 100 – (90/360)6.00 = 98.5

Efficient frontier with a riskless asset *Figure 2.5*

The function giving the expected returns and volatilities of all combinations of holdings in IBM and XYZ is reproduced here. The risk-free asset has a return of 5 percent. By combining the risk-free asset and a portfolio on the frontier, the investor can obtain all the expected return and volatility combinations on the straight line that meets the frontier at the portfolio of risky assets chosen to form these combinations. The figure shows two such lines. The line with the steeper slope is tangent to the efficient frontier at the portfolio *m*. The investor cannot form combinations of the risk-free asset and a risky portfolio that dominate combinations formed from the risk-free asset and portfolio *m*.

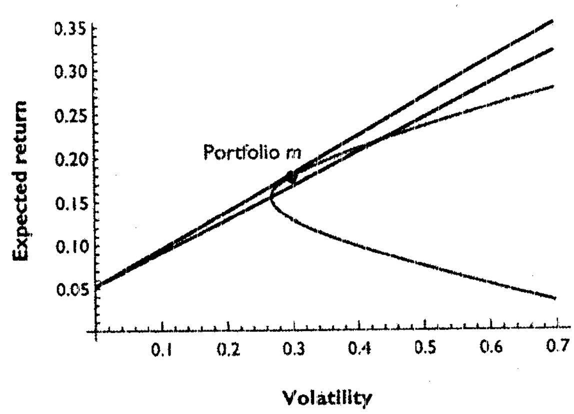

portfolios to the right of the minimum-volatility portfolio require borrowing in the risk-free asset. Figure 2.5 suggests that the investor could do better by combining the risk-free asset with a portfolio more to the right on the efficient frontier of risky assets than the minimum-volatility portfolio because all possible combinations would have a higher return. There is no way that the investor can do better than combining the risk-free asset with portfolio *m*, however, because in that case the straight line is tangent to the efficient frontier of risky assets at *m*. The portfolio *m* is therefore called the **tangency portfolio**. There is no straight line starting at the risk-free rate that touches the efficient frontier of risky assets at least at one point and that has a steeper slope than the line tangent to *m*.

Whenever investors face the same universe of securities and agree on the expected returns, volatilities, and covariances of securities, they end up looking at the same efficient frontier of risky assets and they all want to invest in portfolio *m*. This can be possible only if portfolio m is the market portfolio.

The **market portfolio** is the portfolio of all securities available. Security *i*'s portfolio share in the market portfolio is the ratio (market value of the outstanding supply of security *i*)/ (market value of the outstanding supply of all securities). The value of securities held by all investors together must be the market value of the outstanding supply of all securities. If a security's portfolio share in portfolio *m* were greater than that security's portfolio share in the market portfolio, investors would want to hold more of that security than its outstanding supply. This cannot be an equilibrium. As investors want to hold too much of a security, its expected return has to fall so that investors want to hold less of it. In equilibrium, the expected return of each security must be such that its outstanding supply and no more than its outstanding supply is held by investors.

If all investors have the same views on expected returns, volatilities, and covariances of securities, all of them hold the same portfolio of risky securities, portfolio *m*, the market portfolio. To achieve the right volatility for their invested wealth, they allocate their wealth to the market portfolio and to the risk-free asset. Investors who have little aversion to risk borrow to invest more than their wealth in the market portfolio. The most risk-averse investors put most or all of their wealth in the risk-free asset.

2.2.3. The risk of a security in a diversified portfolio

The extra expected return of a security (or of a portfolio) over the risk-free rate is called the **risk premium** of the security (or portfolio). The risk premium is the reward the investor expects to receive for bearing the risk associated with that security or portfolio. If R_m is the return of portfolio *m*, and R_F is the risk-free rate, $E(R_m) - R_F$ is the risk premium on the market portfolio.

For John to hold the market portfolio, the risk premium on any security has to be just sufficient. Any change in the portfolio's expected return resulting from a very small increase in the investor's holdings of the security must just compensate for the change in the portfolio's risk. If this is not the case, the investor will want to hold a different portfolio, and will no longer hold the market portfolio.

The variance of the return of the market portfolio is the return covariance of the market portfolio with itself. Denote by w_i^m the portfolio share or weight of security *i* in the market portfolio. Since the covariance of a sum of random variables with R_m is equal to the sum of the covariances of the random variables with R_m, it follows that Var(R_m) is equal to a portfolio share weighted sum of the return covariances of the securities with the market portfolio return:

$$Var(R_m) = Cov(R_m, R_m) = Cov\left(\sum_{i=1}^{N} w_i^m R_i, R_m\right) = \sum_{i=1}^{N} w_i^m Cov(R_i, R_m) \tag{2.10}$$

Equation (2.10) shows that a portfolio is risky to the extent that the returns of its securities covary with the return of the market portfolio. The part of the return of a security that covaries with the market portfolio is systematic risk that cannot be eliminated through diversification, since it is part of the risk of the market portfolio and the market portfolio has to be held. The part of the return of security i that does not covary with the market portfolio is diversified away when the investor holds the market portfolio—it is diversifiable risk that the investor does not know is there because it does not affect the volatility of the portfolio.

2.2.4. The capital asset pricing model

With our assumptions, investors care only about the systematic risk of securities and not about their diversifiable risk because the risk of the market portfolio depends only on the systematic risk of securities. Consequently, investors will require a risk premium to bear systematic risk but will not be compensated for bearing diversifiable risk because any investor can get rid of such risk costlessly. A security's systematic risk is proportional to the covariance of its return with the return of the market portfolio. In equation (2.10), if the return of security k has twice the covariance with the return of the market portfolio than the return of security q, it contributes twice to the variance of the market return as security q. Investors should therefore receive twice the reward for holding security k than for holding security q. Otherwise, they would not hold the market portfolio. To see this, suppose that q and k have the same risk premium. An investor could create a portfolio with the same return variance as the market portfolio but with a greater expected return by holding more of security q and less of security k than in the market portfolio.

The result that the risk premium on a security is proportional to its systematic risk is the key insight of the **capital asset pricing model (CAPM)**. The CAPM equation is:

$$E(R_i) - R_F = \beta_i[E(R_m) - R_F]$$

$$\beta_i = \frac{Cov(R_i, R_m)}{Var(R_m)} \tag{2.11}$$

The CAPM tells us that the expected excess return of a risky security is equal to the systematic risk of that security measured by its beta times the market's risk premium. A security's beta (β) is the covariance of the return of the security with the return of the market portfolio divided by the variance of the return of the market portfolio. With the CAPM, if the return of security k has a covariance with the market return that is twice the covariance with the market return of security q, it has twice the **beta** of security q and twice the risk premium. In equation (2.11), the covariance of the return of a security with the market return is divided by the variance of the market return so that a security that has the same systematic risk as the market earns the same risk premium as the market. The covariance of the return of that security with the return of the market is equal to the variance of the market, so that it has a beta of one.

The relation between expected return and beta that results from the CAPM is shown in Figure 2.6. The relation is called the **security market line**. Any portfolio for which the CAPM does not hold is one that investors will want to go long or short in. Consequently, the CAPM must apply to any portfolio. The beta of a portfolio is the portfolio share weighted average of the betas of the securities in the portfolio.

We can apply the CAPM to IBM. Suppose that the risk-free rate is 5 percent, the market risk premium is 6 percent, and the beta of IBM is 1.33. In this case, IBM's expected return is:

Expected return of IBM = 5% + 1.33[6%] = 13%

This is the expected return we used earlier for IBM. Box 2.3, The CAPM in practice, shows how we produce these numbers.

2.3. Diversification and risk management

Once we understand how an investor values a firm using the CAPM, we can find out when risk management increases firm value. For simplicity, let's start with a gold mining firm, Pure Gold Inc. Financial markets are assumed to be perfect as

Figure 2.6 The CAPM

The straight line titled the security market line gives the expected return of a security for a given beta. This line intersects the vertical axis at the risk-free rate and has a value equal to the expected return of the market portfolio for a beta of one.

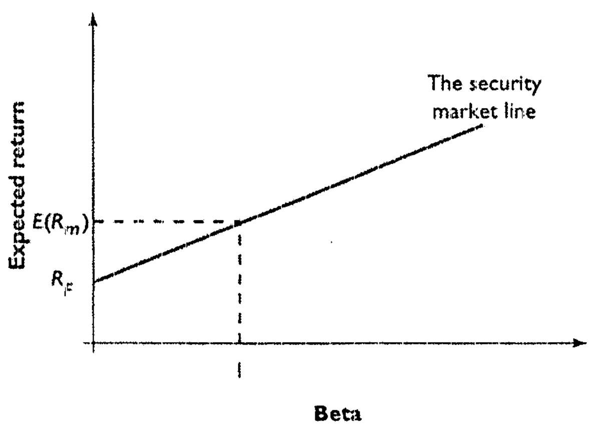

The CAPM in Practice *Box 2.3*

The CAPM provides a formula for the expected return on a security required by capital markets in equilibrium. To implement the CAPM to obtain the expected return on a security, we need to:

Step 1. Identify a proxy for the market portfolio.

Step 2. Identify the appropriate risk-free rate.

Step 3. Estimate the risk premium on the market portfolio.

Step 4. Estimate the beta of the security.

If we are trying to find the expected return of a security over the next month, the next year, or further in the future, all steps involve forecasts except for the first two steps. Using zero-coupon bonds of the appropriate maturity, we can always find the risk-free rate of return for the next month, the next year, or longer in the future, but we always have to forecast the risk premium on the market portfolio and the beta of the security.

What is the appropriate proxy for the market portfolio? Remember that the market portfolio represents how the wealth of investors is invested when the assumptions of the CAPM hold. We cannot observe the market portfolio directly, so we have to use a proxy for it. Most applications of the CAPM in the United States involve the use of some broad U.S. index, such as the S&P 500 as a proxy for the market portfolio. Estimates of the risk premium have decreased substantially over recent years. Some academics and practitioners argue that it should be 4 percent or even less. Using historical data, the estimate is much higher. We compromise with an estimate of 6 percent.

How do we get beta? Consider a security that has traded for a number of years. Suppose that the relation between the return of that security, $R_c(t)$, and the return of the proxy for the market portfolio, $R_m(t)$, is expected to be the same in the future as it was in the past. In this case, one can estimate beta over the past and apply it to the future. To estimate beta, we use linear regression and define a sample period. Five or six years of monthly returns are typical. Having defined the sample period, one estimates an equation over the sample period using regression analysis:

$$R_c(t) = c + bR_m(t) + e(t)$$

In this equation, $e(t)$ is residual risk. It has zero mean and is uncorrelated with the return on the market portfolio. Hence, it corresponds to unsystematic risk. The estimate for b will then be used as the beta of the stock.

Let's look at an example using data for IBM. We have monthly data from January 1992 through the end of September 1997, 69 observations in total. We use as the market portfolio the S&P 500 index. Using Excel, we can get the beta estimate for IBM using the regression program in data analysis under tools. We use the return of IBM in decimal form as the dependent or Y variable and the return on the S&P 500 as the independent or X variable. We

(continued)

Box 2.3 (continued)

use a constant in the regression. The Excel output is reproduced below. The estimates are:

Return on IBM = −0.00123 + 1.371152 (Return on S&P 500)

The standard error associated with the beta estimate is 0.333592. The difference between the beta estimate and the unknown true beta is a normally distributed random variable with zero mean. Using our knowledge about probabilities and the normal distribution, we can find that there is a 95 percent chance that the true beta of IBM is between 0.7053 and 2.037004.

The t-statistic is the ratio of the estimate to its standard error. It provides a test here of the hypothesis that the beta of IBM is different from zero. A t-statistic greater than 1.65 means that we can reject this hypothesis in that there is only a 10 percent chance or less that zero is in the confidence interval constructed around the estimate of beta. Here, the t-statistic is 4.110271. This means that zero is 4.110271 standard errors from 1.371152. The probability that the true beta would be that many standard errors below the mean is the p-value 0.00011.

As the standard error falls, the confidence interval around the coefficient estimates narrows. Consequently, we can make stronger statements about the true beta. The R-square coefficient of 0.201376 means that the return of the S&P 500 explains a fraction 0.201376 of the volatility of the IBM return. As this R-square increases, the independent variable explains more of the variation in the dependent variable.

As one adds independent variables in a regression, the R-square increases. The adjusted R-square takes this effect into account and hence is a more useful guide of the explanatory power of the independent variables when comparing regressions that have different numbers of independent variables.

Summary Output

Regression Statistics

Multiple R	0.44875
R Square	0.201376
Adjusted R-square	0.189457
Standard Error	0.076619
Observations	69

	Coefficients	*Standard Error*	*t-Stat*	*P-value*
Intercept	−0.00123	0.010118	−0.12157	0.903604
X Variable	1.371152	0.333592	4.110271	0.00011

before. The firm will produce one million ounces of gold this year, but after that it will no longer produce gold and it liquidates. For simplicity, the firm has no costs. At the end of the year, the firm has a cash flow of C corresponding to the market value of one million ounces of gold. The firm then pays that cash flow to equity as a liquidating dividend.

Viewed from today, the cash flow is random. The value of the firm today is the present value of receiving the cash flow in one year. We denote this value by V. If the firm is riskless, its value V is its cash flow discounted at the risk-free rate, $C/(1 + R_F)$. If the gold price is fixed at \$350 an ounce and the risk-free rate is 5 percent, the value of the firm is \$350 million/(1 + 0.05), or \$333.33 million.

Now suppose the cash flow is random because the gold price is random. In this case, the random liquidating cash flow C is the market value of the firm at the end of the year. The gain on holding the shares of the firm is therefore $C - V$, where V is the value of the firm at the beginning of the year. The return on shares is therefore $(C - V)/V$. Since C is equal to a quantity of gold times the gold price, the return is perfectly correlated with the gold price, so the firm must have the same beta as gold. Shareholders receive the cash flow in one year for an investment today equal to the value of the firm. This means that the cash flow is equal to the value of the firm times one plus the rate of return of the firm. We know that the expected return of the firm has to be given by the CAPM. Consequently, firm value must be such that:

$$E(C) = V(1 + R_F + \beta[E(R_m) - R_F]) \tag{2.12}$$

If we know the distribution of the cash flow C, the risk-free rate R_F, the β of the firm's shares, and the risk premium on the market $E(R_m) - R_F$, we can compute V because it is the only variable in the equation that we do not know. Solving for V, we get:

$$\frac{E(C)}{1 + R_F + \beta[E(R_m) - R_F]} = V \tag{2.13}$$

The value of the firm is therefore the expected cash flow discounted at the appropriate discount rate from the CAPM. Using this formula, we can value Pure Gold. Let's say that the expected gold price is \$350. In this case, the expected payoff to shareholders is \$350 million, which is one million ounces times the expected price of one ounce. As before, we use a risk-free rate of 5 percent and a risk premium on the market portfolio of 6 percent. We assume a beta of 0.5. Consequently:

$$\frac{E(C)}{1 + R_F + \beta[E(R_m) - R_F]} = \frac{\$350 \text{ million}}{1 + 0.05 + 0.5(0.06)} = \$324.074 \text{ million} \tag{2.14}$$

We can extend this approach to firms expected to remain in existence more than one year. The value is again the present value of the cash flows to shareholders. Nothing else affects the value of the firm for its shareholders—they care only about the present value of cash the firm generates over time for them. We can therefore value a firm's equity in general by computing the sum of the present values of all future cash flows to shareholders using the same approach we used to value one year's future cash flow. For a levered firm, we often consider the value of the firm to be the sum of debt and equity: the present value of the cash flows to the debt and equity holders.

Cash flow to shareholders is computed as net income plus depreciation and other noncash charges minus investment. To get cash flow to the debt and equity holders, one adds to cash flow to equity the payments made to debt holders. The cash flow to shareholders does not necessarily correspond each year to the payouts to equity because firms smooth dividends. A firm may have a positive cash flow to equity holders in excess of its planned dividend; it keeps the excess cash flow in liquid assets and pays it to shareholders later. All cash generated by the firm after debt payments belongs to the shareholders, however, and hence contributes to firm value whether it is paid out in a year or not.

2.3.1. Risk management and shareholder wealth

Would shareholders want a firm to spend cash to reduce the volatility of its cash flow when the only benefit of risk management is to decrease share return volatility? To answer this question, let's assume that the shareholders of the firm are investors who care only about the expected return and the volatility of their wealth invested in securities. These investors hold a diversified portfolio of risky assets, the market portfolio or a portfolio not too different from it, and choose the risk of their end-of-period wealth by allocating their wealth between the risk-free asset and their diversified portfolio of risky assets.

To reduce its volatility, a firm must reduce either its diversifiable risk or its systematic risk. We consider these two approaches to reducing volatility in turn. The firm can reduce risk either through financial transactions or through changes in its operations.

2.3.1.A. Financial risk management policy to reduce the firm's diversifiable risk Assume Markowitz Inc. has a market value of $1 billion and that its management can transfer the diversifiable risk of the firm's shares to an investment bank by paying $50 million. We can think of such a transaction as a hedge offered by the investment bank that exactly offsets the firm's diversifiable risk. Would shareholders ever want the firm to make such a payment when the only benefit to them is to eliminate the diversifiable risk of their shares? We already know that firm value does not depend on diversifiable risk when expected cash flow is given.

Consider then a risk management policy eliminating diversifiable risk that reduces expected cash flow by its cost, but has no other impact on expected cash flow. Since the value of the firm is the expected cash flow discounted at the rate determined by the systematic risk of the firm, this risk management policy does not affect the rate at which cash flow is discounted. In terms of our valuation equation, this policy reduces the numerator of the valuation equation without a change in the denominator, so that firm value is reduced.

Shareholders are diversified; they have no reason to care about diversifiable risks. Therefore, they are not willing to discount expected cash flow at a lower rate if the firm makes cash flow less risky by eliminating diversifiable risk. This means that if shareholders could vote on a proposal to implement risk management to decrease the firm's diversifiable risk at a cost, they would vote no and refuse to incur the cost as long as the only effect of risk management on expected cash flow is to reduce expected cash flow by the cost of risk management.

Managers, therefore, will never be rewarded by shareholders for decreasing the firm's diversifiable risk at a cost because shareholders can eliminate the firm's diversifiable risk through their own diversification at zero cost. For shareholders to value a reduction in diversifiable risk, it has to increase their wealth and hence the share price.

2.3.1.B. Financial risk management policy to reduce the firm's systematic risk Is it worthwhile for management to incur costs to reduce the firm's systematic risk through financial transactions? Suppose IBM decides to reduce its beta because it believes this will make its shares more attractive to investors. It can easily do this by taking a short position in the market, since such a position has a negative beta. The proceeds of the short position can be invested in the risk-free asset.

In our discussion of IBM, we saw that the beta of IBM is 1.33. This means that a dollar invested in IBM has the same systematic risk as $1.33 invested in the market portfolio. Consequently, if IBM were an all-equity firm, the management of IBM could make IBM a zero-beta firm by selling short $1.33 of the market portfolio per dollar of shareholder equity and investing the proceeds in the risk-free asset.

Would investors be willing to pay for IBM management to do this? The answer is no because this action creates no value for the shareholders. In perfect financial markets, shareholders could eliminate the systematic risk of their IBM shares on their own by following the strategy we used when we showed why the CAPM must hold. They would not be willing to pay for the management of IBM to do something that they could do at zero cost on their own if they want to.

The reduction in systematic risk, however, decreases the denominator of the present value formula for shares, since it decreases the discount rate. Why is it that this does not increase the value of the shares? The reason is that reducing systematic risk has a cost, in that it reduces expected cash flow. To get rid of its systematic risk, IBM has to sell the market short. Selling the market short earns a negative risk premium since holding the market long has a positive risk premium. Hence, the expected cash flow of IBM has to fall by the risk premium of the short sale.

The impact of the short sale on firm value is therefore the sum of two effects. The first effect is the reduction in expected cash flow and the second is the drop in the discount rate. The two effects cancel out. Going short in the market is equivalent to getting perfect insurance against market fluctuations. In perfect markets, insurance is priced at its fair value. This means that the risk premium IBM would earn by not changing its systematic risk has to be paid to an entity that will now bear this systematic risk.

Hence, financial risk management in this case simply determines who bears the systematic risk—but IBM's shareholders charge the same price for market risk as anybody else, since that price is determined by the CAPM. Consequently, IBM management cannot create value by selling market risk to other investors at the price that shareholders would require to bear that risk.

2.3.1.C. Does using operations to reduce risk make a difference? What if the firm changes its systematic or its unsystematic risk by changing its operations? The same reasoning applies in this case also, but with a twist. Let's first look at unsystematic risk. Reducing unsystematic risk does not make shareholders better off if the only benefit of doing so is to reduce share return volatility. It does not matter, therefore, whether the decrease in share volatility is due to financial transactions or to operating changes. If the firm can change its operations costlessly to reduce its beta without changing its expected cash flow, however, firm value increases because expected cash flow is discounted at a lower rate. Hence, decreasing cash flow beta through operating changes is worth it if firm value increases as a result.

In financial markets, every investor charges the same for systematic risk. This means that nobody can make money from selling systematic risk to one group of investors instead of another. The ability to change an investment's beta through operating changes depends on technology and strategy. A firm can become more flexible so that it has lower fixed costs in cyclical downturns. This greater flexibility translates into a lower beta. If flexibility has low cost but a great impact on beta, the firm's shareholders are better off if the firm improves its flexibility. If greater flexibility has a high cost, though, shareholders will not want it because this will decrease share value.

2.3.2. Risk management and shareholder clienteles

One gold firm, Homestake, had for a long time a policy of not hedging at all. Homestake justified this policy in its 1990 annual report (p. 12):

> So that its shareholders might capture the full benefit of increases in the price of gold, Homestake does not hedge its gold production. As a result of this policy, Homestake's earnings are more volatile than those of many other gold producers. The Company believes that its shareholders will achieve maximum benefit from such a policy over the long-term.

The rationale for this policy is that some investors want to benefit from gold price movements, and that giving them this benefit increases firm value because they are willing to pay for it. These investors form a clientele the firm caters to. Our analysis so far has not accounted for the possible existence of clienteles such as investors wanting to bear gold price risks. In the world of the CAPM, investors care only about their portfolio's expected return and volatility, not about its sensitivity to other variables, such as gold prices.

The CAPM has limitations in explaining the returns of securities. Small firms, for instance, earn more on average than predicted by the CAPM. It is also possible that investors require a risk premium to bear some risks other than the CAPM's systematic risk; they might, for instance, want a risk premium to bear inflation risk. The presence of such risk premiums could explain why small firms earn more on average than the CAPM predicts.

It could be the case, then, that investors value gold price risk. To see the impact of additional risk premiums besides the market risk premium on our reasoning about the benefits of hedging, let's suppose that Homestake is right, and see what this implies for our analysis of the implications of hedging for the value of Pure Gold Inc., the gold mining firm we valued earlier.

Suppose first that Pure Gold Inc. hedges its gold price risk with a forward contract on gold. It produces one million ounces, so that it wants to sell one million ounces of gold forward. There is no uncertainty about Pure Gold's production, so that we can focus on its value per ounce of gold produced. The price of gold in one year is S and the current forward price is F.

Let's start by assuming there is no clientele effect to establish a benchmark. In this case, the CAPM applies. Empirically, gold has a beta close to zero, and we assume that the gold beta is actually zero. In this case, all the risk of Pure Gold is diversifiable. Let's verify that hedging does not affect firm value in this case. Pure Gold eliminates gold price risk by selling gold forward. The cash flow per ounce of gold to shareholders when the gold is sold forward is F, which is known today. Firm value today per ounce of gold produced is F discounted at the risk-free rate. If the firm does not hedge, the expected cash flow to shareholders per ounce is $E(S)$, which is known today also. In this case, firm value per ounce is obtained by discounting $E(S)$ at the risk-free rate since there is no systematic risk.

The difference between the hedged value of the firm and its unhedged value per ounce is $[F - E(S)]/(1 + R_F)$. The hedged firm is worth more than the unhedged firm if the forward price exceeds the expected spot price, which is true if $F - E(S)$ is positive. Remember that with a short forward position the firm receives F for delivering gold worth S per ounce. $F - E(S)$ is therefore equal to the expected payoff from selling one ounce of gold forward at the price F.

If this expected payoff is positive, an investor can expect to make a profit from entering a short forward position without using any of his resources, since no cash changes hands when a forward position is entered. In equilibrium, the only way the investor can expect to make money without investing any money is if the expected payoff is a reward for bearing risk. Yet we have assumed that the risk associated with the gold price is diversifiable, so that $F - S$ represents diversifiable risk. The expected value of $F - S$ has to be zero, since diversifiable risk does not earn a risk premium. Consequently, $F = E(S)$, and hedging does not affect the firm's value.

Suppose now the case where gold has a positive beta. By taking a long forward position in gold that pays $S - F$, an investor takes on systematic risk. The only way investors would enter such a position is if they are rewarded with a risk premium, which means that they expect to make money on the long forward position. Hence, if gold has systematic risk, it must be that $E(S) > F$, so that the expected payoff to shareholders is lower if the firm hedges than if it does not. However, since a forward contract must have zero value for both parties to be willing to enter the contract, the present value of receiving S in one year and paying F must be zero.

If the firm is hedged, the cash flow has no systematic risk, and the expected cash flow is discounted at the risk-free rate. If the firm is not hedged, the cash flow has systematic risk, so that the higher expected cash flow of the unhedged firm is discounted at a higher discount rate than the lower expected cash flow of the hedged firm. The lower discount rate used for the hedged firm just offsets the decrease in expected cash flow resulting from hedging, so that the present value of expected cash flow is the same whether the firm hedges or not.

We can extend this argument to the case where some investors value exposure to gold for its own sake—perhaps because they feel that it is a hedge against systemic threats. These investors are willing to pay for gold exposure, and as a result the risk premium attached to gold price risk is lower than predicted by the CAPM. Consequently, the forward price has to be higher than otherwise because investors require less compensation to bear gold price risk. If the firm does not hedge, its share price reflects the benefit from exposure to gold that the market values because its discount rate is lower. If the firm hedges, its shareholders are no longer exposed to gold. To hedge, however, the firm sells gold at a higher forward price than if no investors value exposure to gold for its own sake, so in this case the firm's expected cash flow is higher.

Shareholders can earn the premium for gold exposure either because the unhedged firm has a lower discount rate—because of its exposure to gold—or because the hedged firm sells gold forward at a higher price—it has a greater cash flow. The firm has a natural exposure to gold; it is just a matter of which investors bear it. By our reasoning, it does not matter whether the firm's shareholders themselves value gold exposure or not. If the firm's shareholders value gold exposure, they will get it one way or another, but gold exposure will always be priced so that the expected return of investors is not affected by where they get that exposure. If the firm gets rid of its gold exposure, the firm's shareholders can buy it on their own. If the firm's shareholders do not want the gold exposure, they can sell it on their own. No matter how investors who value gold exposure get this exposure, they will have to pay the same price for it, or otherwise the same good—gold exposure—would have different prices on the capital markets, making it possible for investors to profit from these price differences.

An important lesson of this analysis is that the value of the firm is the same whether the firm hedges or not, and however the forward price is determined. If hedging were to create value, there must be opportunities for riskless profits, called **arbitrage profits**, with our assumptions. Suppose that the value of the gold-producing firm is higher if it is hedged than if it is not. Let's assume that each share of the unhedged gold-producing firm pays the value of one ounce of gold. In this case, an investor can create a share of a hedged gold-producing firm on his own by buying a share of the unhedged firm and selling one ounce of gold forward at the price F to hedge. The investor's cash cost today is the price of a share since the forward position has no cash cost today. Having created a share of the hedged firm through homemade hedging, the investor can then sell the share hedged through **homemade hedging** at the price of the share of the hedged firm. There is no difference between the two shares, and hence they should sell for the same price. Through this transaction, the investor makes a profit equal to the difference between the share price of the hedged firm and the share price of the unhedged firm. This profit has no risk attached to it. Consequently, firm value must be the same whether the firm hedges or not with our assumptions.

Let's apply what we have learned to Homestake's clientele argument for not hedging. With perfect financial markets, anybody can get exposure to gold without Homestake's help by taking a long forward position. Whenever investors can do what the firm does on their own and at the same cost—in other words, whenever homemade hedging is possible—the firm cannot possibly create value through hedging. In 2001, Homestake was acquired by American Barrick. American Barrick historically had a policy of protecting itself fully against price declines that could affect its revenue from its anticipated production over the next three years. From American Barrick's perspective, this policy created value for its shareholders—but for reasons we will discuss in the next chapter.

We have shown that Homestake's hedging policy cannot benefit its shareholders even if there is a clientele of investors who value exposure to gold, but we did not address the plausibility of Homestake's claim that investors who value exposure to gold would want to obtain this exposure by buying Homestake shares. That claim has to be questioned. Suppose you are an investor who wants to benefit from increases in the gold price over the coming year. You face the following choice: You can obtain gold exposure by buying Homestake shares and holding them for one year or you can buy a financial security, a gold-indexed zero-coupon bond, that pays you the dollar value of 100 ounces of gold in one year. The gold-indexed zero-coupon bond would be a much better way to obtain gold exposure than Homestake shares. The reason is that the price of Homestake shares might fall over the year even though the gold price increased—a mine could flood, for instance. Hence, using Homestake shares to bet on an increase in the gold price might fail. Provided that the default risk of the gold-indexed zero-coupon bond is trivial, betting on the gold price increase by buying the gold-indexed zero-coupon bond would always be successful if the gold price increases. Gold-indexed bonds can be purchased and there are other securities you could buy that have a payoff indexed to the gold price.

2.3.3. The risk management irrelevance proposition

The major lesson is that a firm cannot create value by hedging risks when it costs the same for the firm to bear these risks directly than to pay the capital markets to bear them. For our purposes in this chapter, the only cost of bearing risks within the firm is the risk premium the capital markets attach to these risks when they value the firm. The same risk premium is required by the capital markets for bearing these risks outside the firm. Consequently, shareholders can alter the firm's risk on their own through homemade hedging at the same terms as the firm, and the firm has nothing to contribute to the shareholders' welfare through risk management. Let's confirm that this is the case by looking at the types of risk the firm faces:

1. **Diversifiable risk.** Diversifiable risk does not affect the share price, and investors do not care about it because it gets diversified within their own portfolios. Hence, eliminating it does not affect firm value.
2. **Systematic risk.** Shareholders require the same risk premium for systematic risk as all investors. Hence, eliminating it for the shareholder just means that the investors who take it on bear it at the same cost. Again, this cannot create value.
3. **Risks valued by investors differently from what the CAPM would predict.** Again, shareholders and other investors charge the same price for bearing such risks.

The bottom line can be summarized in the **hedging irrelevance proposition**: Hedging a risk does not increase firm value when the cost of bearing the risk is the same whether the risk is borne within the firm or outside the firm by the capital markets.

2.4. Summary

In this chapter, we first examined how investors evaluate the risk of securities. We saw how we can use a distribution function to evaluate the probability of various outcomes for the return of a security. The ability to specify the probability that a security will experience a return lower than some pre-specified benchmark will be of crucial importance throughout this book. We then saw how investors can diversify, and that the ability to diversify affects how an investor evaluates the riskiness of a security. A security's contribution to the risk of the market portfolio is its systematic risk; its diversifiable risk does not affect the riskiness of the portfolio. This fundamental result allowed us to present the capital asset pricing model, which states that a security's risk premium is given by its beta times the risk premium on the market portfolio. The CAPM allows us to compute the value of future cash flows. We saw that only the systematic risk of cash flows affects the rate at which investors discount expected future cash flows.

We then showed that, in perfect financial markets, hedging does not affect firm value, whether hedging systematic or unsystematic risks through financial instruments. Further, we demonstrated that even if investors have preferences for some types of risks, like gold price risks, hedging is still irrelevant in perfect financial markets. If it costs the same for a firm to bear a risk as it does for the firm to pay somebody else to bear it, hedging cannot increase firm value.

Key Concepts

Review Questions

1. Assume a stock return follows the normal distribution. What do you need to know to compute the probability that the stock's return will be less than 10 percent during the coming year?

2. What does the variance of the return of a portfolio depend on?
3. What does diversification of a portfolio do to the distribution of the portfolio's return?
4. What is beta?
5. When does beta measure risk?
6. For a given expected cash flow, how does the beta of the cash flow affect its current value?
7. How does hedging affect firm value if financial markets are perfect?
8. Why can hedging affect a firm's expected cash flow when it does not affect its value?
9. Why does the fact that some investors have a preference for gold exposure have no bearing on whether firms should hedge gold exposure?
10. What is the risk management irrelevance proposition?

Questions and Exercises

1. The typical level of the monthly volatility of the S&P 500 index is about 4 percent. Using a risk premium of 6 percent and a risk-free rate of 5 percent per year, what is the probability that a portfolio of $100,000 invested in the S&P 500 will lose $5,000 or more during the next month? How would your answer change if you used current interest rates from T-bills?
2. During 1997, the monthly volatility on the S&P 500 increased to about 4.5 percent from its typical value of 4.0 percent. Using the current risk-free rate, construct a portfolio worth $100,000 invested in the S&P 500 and the risk-free asset that has the same probability of losing $5,000 or more in a month when the S&P 500 volatility is 4.5 percent as a portfolio of $100,000 invested in the S&P 500 when its volatility is 4 percent.
3. Compute the expected return and the volatility of return of a portfolio that has a portfolio share of 0.9 in the S&P 500 and 0.1 in an emerging market index. The S&P 500 has a volatility of return of 15 percent and an expected return of 12 percent. The emerging market has a return volatility of 30 percent and an expected return of 10 percent. The correlation between the emerging market index return and the S&P 500 is 0.1.
4. If the S&P 500 is a good proxy for the market portfolio in the CAPM, and the CAPM applies to the emerging market index, use the information in question 3 to compute the beta and risk premium for the emerging market index.
5. Compute the beta of the portfolio described in question 4 with respect to the S&P 500.
6. A firm has an expected cash flow of $500 million in one year. The beta of the common stock of the firm is 0.8 and this cash flow has the same risk as the firm as a whole. Using a risk-free rate of 5 percent and a risk premium on the market portfolio of 6 percent, what is the present value of the cash

flow? If the beta of the firm doubles, what happens to the present value of the cash flow?

7. Using the data in the previous question, consider the impact on the firm of hedging the cash flow against systematic risk. If management wants to eliminate the systematic risk of the cash flow completely, how could it do so? How much would the firm have to pay investors to bear the systematic risk of the cash flow?

8. Consider the situation you analyzed in question 6. To hedge the firm's systematic risk, management has to pay investors to bear this risk. Why is it that the value of the firm for shareholders does not fall when the firm pays other investors to bear the cash flow's systematic risk?

9. The management of a gold-producing firm agrees with the hedging irrelevance result and has concluded that it applies to the firm. However, the CEO wants to hedge because the price of gold has fallen over the last month. He asks for your advice. What do you tell him?

10. Consider again an investment in the emerging market portfolio of question 3. You consider investing $100,000 in that portfolio because you think it is a good investment. You decide that you are going to ignore the benefits from diversification, in that all your wealth will be invested in that portfolio. Your broker nevertheless presents you with an investment in a default-free bond in the currency of the emerging country, which matures in one year. The expected return on the foreign currency bond is 5 percent in dollars, its volatility is 10 percent, and the correlation of its return with the dollar return of the emerging market portfolio is 1. Compute the expected return of a portfolio with $100,000 in the emerging market portfolio, –$50,000 in the foreign currency bond, and $50,000 in the domestic risk-free asset that earns 5 percent per year. How does this portfolio differ from the portfolio that has only an investment in the emerging market portfolio? Which one would you choose and why? Could you create a portfolio with investments in the emerging market portfolio, in the emerging market currency risk-free bond, and in the risk-free asset that has the same mean return but a lower volatility?

Literature Note

Much research examines the appropriateness of the assumption of the normal distribution for security returns. Fama (1965) argues that monthly stock returns are well-described by the normal distribution and that these returns are independent across time. We will see later that, while the normal distribution is a good starting point, it is sometimes necessary to make different distributional assumptions.

The fundamental research on diversification and the CAPM is mainly the work, respectively, of Markowitz (1952) and Sharpe (1964), who were awarded a share of the Nobel Memorial Prize in Economics in 1990. Textbooks on investments cover this material in much greater detail. Elton, Gruber, Brown, and Goetzmann (2002) provide an extensive presentation of portfolio theory. Valuation theory using the CAPM is discussed in corporate finance textbooks or in textbooks specialized in valuation. For corporate finance textbooks, see Brealey and Myers (2002) or Jordan, Ross, and Westerfield (2002). A book devoted to valuation is Copeland, Koller, and Murrin (1996).

The hedging irrelevance result is discussed in Smith and Stulz (1985). This result is a natural extension of the leverage irrelevance result of Modigliani and Miller (1958). Modigljani and Miller (1958) argue that in perfect markets leverage cannot increase firm value. Their result led to the award of a Nobel Memorial Prize in Economics in 1990 for Miller. Modigliani received such a Prize earlier for a different contribution.