GARP
DIGITAL LIBRARY

"Investors and Risk Management"

Chapter 2 of

Risk Management and Derivatives

By René Stulz

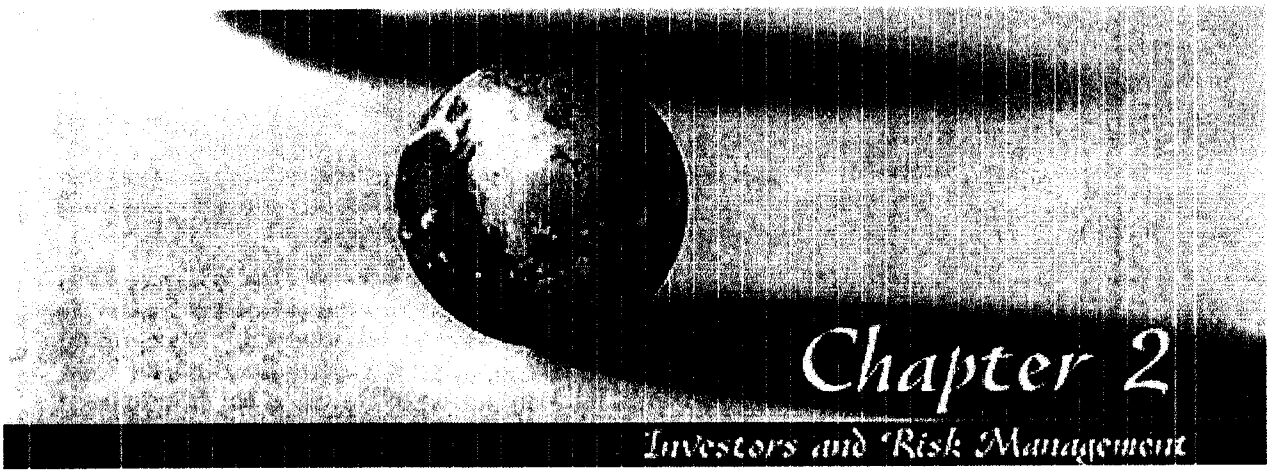

Chapter 2

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.

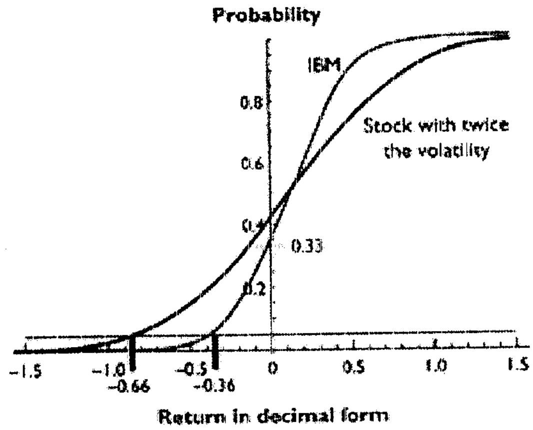

ty depends on the expected return. As the expected return of IBM increases, the probability of making a loss falls.

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.

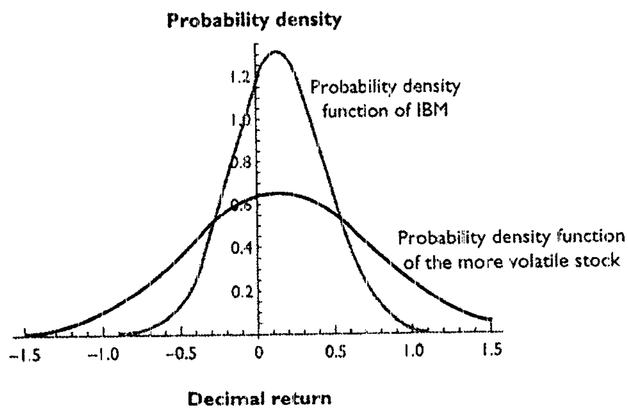

stocks have limited liability, the most one can lose owning a stock is what one paid for it, corresponding to a simple return of –100 percent. In general, this limitation is not important in that the probability of such a return is very small.

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