

Chapter 1

Introduction: Risk-Return Trade-offs

Imagine you are the Chief Risk Officer (CRO) of a major corporation. The Chief Executive Officer (CEO) wants your views on a major new venture. You have been inundated with reports showing that the new venture has a positive net present value and will enhance shareholder value. What sort of analysis and ideas is the CEO looking for from you?

As CRO it is your job to consider how the new venture fits into the company's portfolio. What is the correlation of the performance of the new venture with the rest of the company's business? When the rest of the business is experiencing difficulties, will the new venture also provide poor returns, or will it have the effect of dampening the ups and downs in the rest of the business?

Companies must take risks if they are to survive and prosper. The risk management function's primary responsibility is to understand the portfolio of risks that the company is currently taking and the risks it plans to take in the future. It must decide whether the risks are acceptable and, if they are not acceptable, what action should be taken.

Most of this book is concerned with the ways risks are managed by banks and other financial institutions, but many of the ideas and approaches we will discuss are equally applicable to nonfinancial corporations. Risk management has become progressively more important for all corporations in the last few decades. Financial institutions in particular are finding they have to increase the resources they devote to risk management. In spite of this, financial institutions still make mistakes. In March 2021, Archegos,

a hedge fund managing the assets of a rich private individual named Bill Hwang, went bankrupt as a result of highly levered risky investments. Several banks that the hedge fund had been dealing with lost huge sums. It is reported that Credit Suisse lost over \$5 billion and Nomura lost almost \$3 billion as a result of their dealings with Archegos.

Risk management is not about minimizing risks. It is about ensuring that the risks taken are manageable and that the expected returns are commensurate with the risks being taken. This chapter sets the scene. It starts by reviewing the classical arguments concerning the risk–return trade–offs faced by a well–diversified investor who is choosing a portfolio of stocks and bonds. It then considers whether the same arguments can be used by a company in choosing new projects and managing its risk exposure. The chapter concludes that there are reasons why all companies, including financial institutions, should be concerned with the total risk they face, not just with the risk from the viewpoint of a well–diversified shareholder.

1.1 Risk vs. Return for Investors

We now examine the theoretical trade–off between risk and return when money is invested. The first point to note is that the trade–off is actually between risk and *expected return*, not between risk and actual return. The term “expected return” sometimes causes confusion. In everyday language an outcome that is “expected” is considered highly likely to occur. However, statisticians define the expected value of a variable as its average (or mean) value. Expected return is therefore a weighted average of the possible returns, where the weight applied to a particular return equals the probability of that return occurring. The possible returns and their probabilities can be either estimated from historical data or assessed subjectively.

Suppose, for example, that you have \$100,000 to invest for one year. Suppose further that Treasury bills yield 5%. One alternative is to buy Treasury bills. There is then no risk and the expected return is 5%. Another alternative is to invest the \$100,000 in a stock. To simplify things, we suppose that the possible outcomes from this investment are as shown in Table 1.1. There is a 0.05 probability that the return will be +50%; there is a 0.25 probability that the return will be +30%; and so on. Expressing the returns in decimal form, the expected return per year is:

$$0.05 \times 0.50 + 0.25 \times 0.30 + 0.40 \times 0.10 + 0.25 \times (-0.10) + 0.05 \times (-0.30) = 0.10$$

This shows that, in return for taking some risk, you are able to increase your expected return per annum from the 5% offered by Treasury bills to 10%. If things work out well, your return per annum could be as high as 50%. But the worst–case outcome is a –30% return or a loss of \$30,000.

One of the first attempts to understand the trade–off between risk and expected return was by Markowitz (1952). Later, Sharpe (1964) and others carried the Markowitz

Table 1.1 One-Year Return from Investing \$100,000 in a Stock

Probability	Return
0.05	+50%
0.25	+30%
0.40	+10%
0.25	-10%
0.05	-30%

analysis a stage further by developing what is known as the capital asset pricing model. This is a relationship between expected return and what is termed *systematic risk*, which will be explained later in this chapter. In 1976, Ross developed arbitrage pricing theory, which can be viewed as an extension of the capital asset pricing model to the situation where there are several sources of systematic risk. The key insights of these researchers have had a profound effect on the way portfolio managers think about and analyze the risk-return trade-offs they face. We will now review these insights.

1.1.1 Quantifying Risk

How do you quantify the risk you are taking when you choose an investment? A convenient measure that is often used is the standard deviation of the return over one year. This is

$$\sqrt{E(R^2) - [E(R)]^2}$$

where R is the return per annum. The symbol E denotes expected value so that $E(R)$ is the expected return per annum. In Table 1.1, as we have shown, $E(R) = 0.10$. To calculate $E(R^2)$ we must weight the alternative squared returns by their probabilities:

$$\begin{aligned} E(R^2) &= 0.05 \times 0.50^2 + 0.25 \times 0.30^2 + 0.40 \times 0.10^2 + 0.25 \times (-0.10)^2 \\ &\quad + 0.05 \times (-0.30)^2 = 0.046 \end{aligned}$$

The standard deviation of the annual return is therefore $\sqrt{0.046 - 0.1^2} = 0.1897$ or 18.97%.

1.1.2 Investment Opportunities

Suppose we choose to characterize every investment opportunity by its expected return and standard deviation of return. We can plot available risky investments on a chart such as Figure 1.1, where the horizontal axis is the standard deviation of the return and the vertical axis is the expected return.




Figure 1.1 Alternative Risky Investments

Once we have identified the expected return and the standard deviation of the return for individual investments, it is natural to think about what happens when we combine investments to form a portfolio. Consider two investments with returns R_1 and R_2 . The return from putting a proportion w_1 of our money in the first investment and a proportion $w_2 = 1 - w_1$ in the second investment is

$$w_1 R_1 + w_2 R_2$$

The portfolio expected return is

$$\mu_p = w_1 \mu_1 + w_2 \mu_2 \quad (1.1)$$

where μ_1 is the expected return from the first investment and μ_2 is the expected return from the second investment. The standard deviation of the portfolio return is given by

$$\sigma_p = \sqrt{w_1^2 \sigma_1^2 + w_2^2 \sigma_2^2 + 2\rho w_1 w_2 \sigma_1 \sigma_2} \quad (1.2)$$

where σ_1 and σ_2 are the standard deviations of R_1 and R_2 and ρ is the coefficient of correlation between the two.

Suppose that μ_1 is 10% per annum and σ_1 is 16% per annum, while μ_2 is 15% per annum and σ_2 is 24% per annum. Suppose also that the coefficient of correlation, ρ , between the returns is 0.2 or 20%. Table 1.2 shows the values of μ_p and σ_p for a number of different values of w_1 and w_2 . The calculations show that by putting part of your money in the first investment and part in the second investment a wide range of risk-return combinations can be achieved. These are plotted in Figure 1.2.




Figure 1.2 Alternative Risk-Return Combinations from Two Investments (as Calculated in Table 1.2)

Table 1.2 Expected Return, μ_p , and Standard Deviation of Return, σ_p , from a Portfolio Consisting of Two Investments

w_1	w_2	μ_p	σ_p
0.0	1.0	15%	24.00%
0.2	0.8	14%	20.09%
0.4	0.6	13%	16.89%
0.6	0.4	12%	14.87%
0.8	0.2	11%	14.54%
1.0	0.0	10%	16.00%

The expected returns from the investments are 10% and 15%; the standard deviations of the returns are 16% and 24%; and the correlation between returns is 0.2.

Most investors are risk-averse. They want to increase expected return while reducing the standard deviation of return. This means that they want to move as far as they can in a “northwest” direction in Figures 1.1 and 1.2. Figure 1.2 shows that forming a portfolio of the two investments we have been considering helps them do this. For example, by putting 60% in the first investment and 40% in the second, a portfolio with an expected return of 12% and a standard deviation of return equal to 14.87% is obtained. This is an improvement over the risk-return trade-off from putting 100% in the first investment. (The expected return is 2% higher and the standard deviation of the return is 1.13% lower.) Assuming that the investor is only concerned with the mean and standard deviation of returns and is risk averse, we can say that putting all funds in the first investment is suboptimal. It is dominated by other opportunities.

1.2 The Efficient Frontier

Let us now bring a third investment into our analysis. The third investment can be combined with any combination of the first two investments to produce new risk–return combinations. This enables us to move further in the northwest direction. We can then add a fourth investment. This can be combined with any combination of the first three investments to produce yet more investment opportunities. As we continue this process, considering every possible portfolio of the available risky investments, we obtain what is known as an *efficient frontier*. This represents the limit of how far we can move in a northwest direction and is illustrated in Figure 1.3. There is no investment that dominates a point on the efficient frontier in the sense that it has both a higher expected return and a lower standard deviation of return. The area southeast of the efficient frontier represents the set of all investments that are possible. For any point in this area that is not on the efficient frontier, there is a point on the efficient frontier that has a higher expected return and lower standard deviation of return.

In Figure 1.3 we have considered only risky investments. What does the efficient frontier of all possible investments look like? Specifically, what happens when we include the risk–free investment? Suppose that the risk–free investment yields a return of R_F . In Figure 1.4 we have denoted the risk–free investment by point F and drawn a tangent from point F to the efficient frontier of risky investments that was developed in Figure 1.3. M is the point of tangency. As we will now show, the line FM is our new efficient frontier.

Consider what happens when we form an investment I by putting β_I of the funds we have available for investment in the risky portfolio, M , and $1 - \beta_I$ in the risk–free investment, F ($0 < \beta_I < 1$). From equation (1.1) the expected return from the investment, $E(R_I)$, is given by

$$E(R_I) = (1 - \beta_I)R_F + \beta_I E(R_M)$$




Figure 1.3 Efficient Frontier Obtainable from Risky Investments




Figure 1.4 The Efficient Frontier of All Investments

Point I is achieved by investing a percentage β_I of available funds in portfolio M and the rest in a risk-free investment. Point J is achieved by borrowing $\beta_J - 1$ of available funds at the risk-free rate and investing everything in portfolio M .

and from equation (1.2), because the risk-free investment has zero standard deviation, the return R_I has standard deviation

$$\beta_I \sigma_M$$

where σ_M is the standard deviation of return for portfolio M . This risk-return combination corresponds to the point labeled I in Figure 1.4. From the perspective of both expected return and standard deviation of return, point I is β_I of the way from F to M .

All points on the line FM can be obtained by choosing a suitable combination of the investment represented by point F and the investment represented by point M . The points on this line dominate all the points on the previous efficient frontier because they give a better risk-return combination. The straight line FM is therefore part of the new efficient frontier.

If we make the simplifying assumption that we can borrow at the risk-free rate of R_F as well as invest at that rate, we can create investments that are on the continuation of FM beyond M . Suppose, for example, that we want to create the investment represented by the point J in Figure 1.4 where the distance of J from F is β_J times the distance of M from F ($\beta_J > 1$). We borrow $\beta_J - 1$ of the amount that we have available for investment at rate R_F and then invest everything (the original funds and the borrowed funds) in

the investment represented by point M . After allowing for the interest paid, the new investment has an expected return, $E(R_J)$, given by

$$E(R_J) = \beta_J E(R_M) - (\beta_J - 1)R_F = (1 - \beta_J)R_F + \beta_J E(R_M)$$

and the standard deviation of the return is

$$\beta_J \sigma_M$$

This shows that the risk and expected return combination corresponds to point J . (Note that the formulas for the expected return and standard deviation of return in terms of beta are the same whether beta is greater than or less than 1.)

The idealized argument that we have presented shows that, when the risk-free investment is considered, the efficient frontier must be a straight line. To put this another way, there should be a linear trade-off between the expected return and the standard deviation of returns, as indicated in Figure 1.4. All investors should choose the same portfolio of risky assets. This is the portfolio represented by M . They should then reflect their appetite for risk by combining this risky investment with borrowing or lending at the risk-free rate.

It is a short step from here to argue that the portfolio of risky investments represented by M must be the portfolio of all risky investments. Suppose a particular investment is not in the portfolio. No investors would hold it and its price would have to go down so that its expected return increased and it became part of portfolio M . In fact, we can go further than this. To ensure a balance between the supply and demand for each investment, the price of each risky investment must adjust so that the amount of that investment in portfolio M is proportional to the amount of that investment available in the economy. The investment represented by point M is therefore usually referred to as the *market portfolio*.

1.3 The Capital Asset Pricing Model

How do investors decide on the expected returns they require for individual investments? Based on the analysis we have presented, the market portfolio should play a key role. The expected return required on an investment should reflect the extent to which the investment contributes to the risks of the market portfolio.

A common procedure is to use historical data and regression analysis to determine a best-fit linear relationship between returns from an investment and returns from the market portfolio. This relationship has the form:

$$R = a + \beta R_M + \epsilon \tag{1.3}$$

where R is the return from the investment, R_M is the return from the market portfolio, a and β are constants, and ϵ is a random variable equal to the regression error.

Equation (1.3) shows that there are two uncertain components to the risk in the investment's return:

1. A component βR_M , which is a multiple of the return from the market portfolio.
2. A component ϵ , which is unrelated to the return from the market portfolio.

The first component is referred to as *systematic risk*. The second component is referred to as *nonsystematic risk*.

Consider first the nonsystematic risk. If we assume that the ϵ variables for different investments are independent of each other, the nonsystematic risk is almost completely diversified away in a large portfolio. An investor should not therefore be concerned about nonsystematic risk and should not require an extra return above the risk-free rate for bearing nonsystematic risk.

The systematic risk component is what should matter to a well-diversified investor. When a large well-diversified portfolio is held, the systematic risk represented by βR_M does not disappear. An investor should require an expected return to compensate for this systematic risk.

We know how investors trade off systematic risk and expected return from Figure 1.4. When $\beta = 0$ there is no systematic risk and the expected return is R_F . When $\beta = 1$, we have the same systematic risk as the market portfolio, which is represented by point M , and the expected return should be $E(R_M)$. In general

$$E(R) = R_F + \beta[E(R_M) - R_F] \quad (1.4)$$

This is the *capital asset pricing model*. The excess expected return over the risk-free rate required on the investment is β times the excess expected return on the market portfolio. This relationship is plotted in Figure 1.5. The parameter β is the *beta* of the investment.




Figure 1.5 The Capital Asset Pricing Model

Example 1.1

Suppose that the risk-free rate is 5% and the return on the market portfolio is 10%. An investment with a beta of 0 should have an expected return of 5%. This is because all of the risk in the investment can be diversified away. An investment with a beta of 0.5 should have an expected return of

$$0.05 + 0.5 \times (0.1 - 0.05) = 0.075$$

or 7.5%. An investment with a beta of 1.2 should have an expected return of

$$0.05 + 1.2 \times (0.1 - 0.05) = 0.11$$

or 11%.

The parameter β is equal to $\rho\sigma/\sigma_M$, where ρ is the correlation between the return on the investment and the return on the market portfolio, σ is the standard deviation of the return on the investment, and σ_M is the standard deviation of the return on the market portfolio. Beta measures the sensitivity of the return on the investment to the return on the market portfolio. We can define the beta of any investment portfolio as in equation (1.3) by regressing its returns against the returns on the market portfolio. The capital asset pricing model in equation (1.4) should then apply with the return R defined as the return on the portfolio. In Figure 1.4 the market portfolio represented by M has a beta of 1.0 and the riskless portfolio represented by F has a beta of zero. The portfolios represented by the points I and J have betas equal to β_I and β_J , respectively.

1.3.1 Assumptions

The idealized analysis we have presented leads to the surprising conclusion that all investors want to hold the same portfolios of assets (the portfolio represented by M in Figure 1.4). This is clearly not true. Indeed, if it were true, markets would not function at all well because investors would not want to trade with each other! In practice, different investors have different views on the attractiveness of stocks and other risky investment opportunities. This is what causes them to trade with each other and it is this trading that leads to the formation of prices in markets.

The reason why the analysis leads to conclusions that do not correspond with the realities of markets is that, in presenting the arguments, we implicitly made a number of assumptions. In particular:

1. We assumed that investors care only about the expected return and the standard deviation of return of their portfolio. Another way of saying this is that investors look only at the first two moments of the return distribution. If returns are normally

distributed, it is reasonable for investors to do this. However, the returns from many assets are non-normal. They have *skewness* and *excess kurtosis*. Skewness is related to the third moment of the distribution and excess kurtosis is related to the fourth moment. In the case of positive skewness, very high returns are more likely and very low returns are less likely than the normal distribution would predict; in the case of negative skewness, very low returns are more likely and very high returns are less likely than the normal distribution would predict. Excess kurtosis leads to a distribution where both very high and very low returns are more likely than the normal distribution would predict. Most investors are concerned about the possibility of extreme negative outcomes. They are likely to want a higher expected return from investments with negative skewness or excess kurtosis.

2. We assumed that the ϵ variables for different investments in equation (1.3) are independent. Equivalently we assumed the returns from investments are correlated with each other only because of their correlation with the market portfolio. This is clearly not true. Ford and General Motors are both in the automotive sector. There is likely to be some correlation between their returns that does not arise from their correlation with the overall stock market. This means that the ϵ for Ford and the ϵ for General Motors are not likely to be independent of each other.
3. We assumed that investors focus on returns over just one period and the length of this period is the same for all investors. This is also clearly not true. Some investors such as pension funds have very long time horizons. Others such as day traders have very short time horizons.
4. We assumed that investors can borrow and lend at the same risk-free rate. This is approximately true in normal market conditions for a large financial institution that has a good credit rating. But it is not true for small investors.
5. We did not consider tax. In some jurisdictions, capital gains are taxed differently from dividends and other sources of income. Some investments get special tax treatment and not all investors are subject to the same tax rate. In practice, tax considerations have a part to play in the decisions of an investor. An investment that is appropriate for a pension fund that pays no tax might be inappropriate for a high-marginal-rate taxpayer living in New York, and vice versa.
6. Finally, and perhaps most importantly, we assumed that all investors make the same estimates of expected returns, standard deviations of returns, and correlations between returns for available investments. To put this another way, we assumed that investors have *homogeneous expectations*. This is clearly not true. Indeed, if we lived in a world of homogeneous expectations there would be no trading.

In spite of all this, the capital asset pricing model has proved to be a useful tool for portfolio managers. Estimates of the betas of stocks are readily available and the expected return on a portfolio estimated by the capital asset pricing model is a commonly used benchmark for assessing the performance of the portfolio manager, as we will now explain.

1.3.2 Alpha

When we observe a return of R_M on the market, what do we expect the return on a portfolio with a beta of β to be? The capital asset pricing model relates the expected return on a portfolio to the expected return on the market. But it can also be used to relate the expected return on a portfolio to the actual return on the market:

$$E(R_p) = R_F + \beta(R_M - R_F)$$

where R_F is the risk-free rate and R_p is the return on the portfolio.

Example 1.2

Consider a portfolio with a beta of 0.6 when the risk-free interest rate is 4%. When the return on the market is 20%, the expected return on the portfolio is

$$0.04 + 0.6 \times (0.2 - 0.04) = 0.136$$

or 13.6%. When the return on the market is 10%, the expected return from the portfolio is

$$0.04 + 0.6 \times (0.1 - 0.04) = 0.076$$

or 7.6%. When the return from the market is -10% , the expected return from the portfolio is

$$0.04 + 0.6 \times (-0.1 - 0.04) = -0.044$$

or -4.4% . The relationship between the expected return on the portfolio and the return on the market is shown in Figure 1.6.




Figure 1.6 Relationship between Expected Return on Portfolio and the Actual Return on the Market When Portfolio Beta Is 0.6 and Risk-Free Rate Is 4%

Suppose that the actual return on the portfolio is greater than the expected return:

$$R_P > R_F + \beta(R_M - R_F)$$

The portfolio manager has produced a superior return for the amount of systematic risk being taken. The extra return is

$$\alpha = R_P - R_F - \beta(R_M - R_F)$$

This is commonly referred to as the *alpha* created by the portfolio manager.¹

Example 1.3

A portfolio manager has a portfolio with a beta of 0.8. The one-year risk-free rate of interest is 5%, the return on the market during the year is 7%, and the portfolio manager's return is 9%. The manager's alpha is

$$\alpha = 0.09 - 0.05 - 0.8 \times (0.07 - 0.05) = 0.024$$

or 2.4%.

Portfolio managers are continually searching for ways of producing a positive alpha. One way is by trying to pick stocks that outperform the market. Another is by *market timing*. This involves trying to anticipate movements in the market as a whole and moving funds from safe investments such as Treasury bills to the stock market when an upturn is anticipated and in the other direction when a downturn is anticipated. Chapter 4 explains other strategies used by hedge funds to try to create positive alpha.

Although the capital asset pricing model is unrealistic in some respects, the alpha and beta parameters that come out of the model are widely used to characterize investments. Beta describes the amount of systematic risk. The higher the value of beta, the greater the systematic risk being taken and the greater the extent to which returns are dependent on the performance of the market. Alpha represents the extra return made from superior portfolio management (or perhaps just good luck). An investor can make a positive alpha only at the expense of other investors who are making a negative alpha. The weighted average alpha of all investors must be zero.

1.4 Arbitrage Pricing Theory

Arbitrage pricing theory can be viewed as an extension of the capital asset pricing model. In the capital asset pricing model, an asset's return depends on just one factor. In arbitrage pricing theory, the return depends on several factors. (These factors might involve

¹It is sometimes referred to as Jensen's alpha because it was first used by Michael Jensen in evaluating mutual fund performance. See Section 4.3.

variables such as the gross national product, the domestic interest rate, and the inflation rate.) By exploring ways in which investors can form portfolios that eliminate exposure to the factors, arbitrage pricing theory shows that the expected return from an investment is linearly dependent on the factors.

The assumption that the ϵ variables for different investments are independent in equation (1.3) ensures that there is just one factor driving expected returns (and therefore one source of systematic risk) in the capital asset pricing model. This is the return on the market portfolio. In arbitrage pricing theory there are several factors affecting investment returns. Each factor is a separate source of systematic risk. Unsystematic (i.e., diversifiable) risk in arbitrage pricing theory is the risk that is unrelated to all the factors.

1.5 Risk vs. Return for Companies

We now move on to consider the trade-offs between risk and return made by a company. How should a company decide whether the expected return on a new investment project is sufficient compensation for its risks?

The ultimate owners of a company are its shareholders and a company should be managed in the best interests of its shareholders. It is therefore natural to argue that a new project undertaken by the company should be viewed as an addition to its shareholders' portfolio. The company should calculate the beta of the investment project and its expected return. If the expected return is greater than that given by the capital asset pricing model, it is a good deal for shareholders and the investment should be accepted. Otherwise it should be rejected.

The argument just given suggests that nonsystematic risks should not be considered when accept/reject decisions on new projects are taken. In practice, companies are concerned about nonsystematic as well as systematic risks. For example, most companies insure themselves against the risk of their buildings burning down—even though this risk is entirely nonsystematic and can be diversified away by their shareholders. They try to avoid taking high risks and often hedge their exposures to exchange rates, interest rates, commodity prices, and other market variables.

Earnings stability and the survival of the company are often important managerial objectives. Companies do try to ensure that their expected returns on new ventures are consistent with the risk–return trade-offs of their shareholders. But there is an overriding constraint that the total risks taken should not be allowed to get too large.

Many investors are also concerned about the overall risk of the companies they invest in. They do not like surprises and prefer to invest in companies that show solid growth and meet earnings forecasts. They like companies to manage risks carefully and limit the overall amount of risk—both systematic and nonsystematic—they are taking.

The theoretical arguments we presented in Sections 1.1 to 1.4 suggest that investors should not behave in this way. They should hold a well-diversified portfolio and encourage the companies they invest in to make high-risk investments when the combination

of expected return and systematic risk is favorable. Some of the companies in a shareholder's portfolio will go bankrupt, but others will do very well. The result should be an overall return to the shareholder that is satisfactory.

Are investors behaving suboptimally? Would their interests be better served if companies took more nonsystematic risks? There is an important argument to suggest that this is not necessarily the case. This argument is usually referred to as the “bankruptcy costs” argument. It is often used to explain why a company should restrict the amount of debt it takes on, but it can be extended to apply to a wider range of risk management decisions than this.

1.5.1 Bankruptcy Costs

In a perfect world, bankruptcy would be a fast affair where the company's assets (tangible and intangible) are sold at their fair market value and the proceeds are distributed to the company's creditors using well-defined rules. If we lived in such a perfect world, the bankruptcy process itself would not destroy value for stakeholders. Unfortunately, the real world is far from perfect. By the time a company reaches the point of bankruptcy, it is likely that its assets have lost some value. The bankruptcy process itself invariably reduces the value of its assets further. This further reduction in value is referred to as *bankruptcy costs*.

What is the nature of bankruptcy costs? Once a bankruptcy has happened, customers and suppliers become less inclined to do business with the company; assets sometimes have to be sold quickly at prices well below those that would be realized in an orderly sale; the value of important intangible assets, such as the company's brand name and its reputation in the market, are often destroyed; the company is no longer run in the best interests of shareholders; large fees are often paid to accountants and lawyers; and so on. The story in Business Snapshot 1.1 is representative of what often happens in practice. It illustrates how, when a high-risk decision works out badly, there can be disastrous bankruptcy costs.

The largest bankruptcy in U.S. history was that of Lehman Brothers on September 15, 2008. The bankruptcy proceedings continued for many years. Denison, Fleming and Sarkar at the Federal Reserve Bank of New York estimated in 2019 that the direct and indirect costs of the bankruptcy totaled between \$46 billion and \$63 billion, or 15% to 21% of Lehman's pre-bankruptcy consolidated assets.²

We mentioned earlier that corporate survival is an important managerial objective and that shareholders like companies to avoid excessive risks. We now understand one reason why this is so. Bankruptcy laws vary widely from country to country, but they all have the effect of destroying value as lenders and other creditors vie with each other to get

²See: <https://libertystreeteconomics.newyorkfed.org/2019/01/how-much-value-was-destroyed-by-the-lehman-bankruptcy/>.

paid. If a company chooses projects with very high risks (but sufficiently high expected returns to be above the efficient frontier in Figure 1.4), the probability of bankruptcy will be quite high. When expected bankruptcy costs are taken into account, projects that have a high total (systematic plus nonsystematic) risk are liable to be rejected as unacceptable. This explains why investors like companies to limit the overall amount of risk they take and reward companies that manage risks so that they meet earnings forecasts.

BUSINESS SNAPSHOT 1.1

The Hidden Costs of Bankruptcy: A Hypothetical Scenario

Several years ago, a company had a market capitalization of \$2 billion and \$500 million of debt. The CEO decided to acquire a company in a related industry for \$1 billion in cash. The cash was raised using a mixture of bank debt and bond issues. The price paid for the company was justified on the basis of potential synergies, but key threats to the profitability of the company were overlooked.

Many of the anticipated synergies were not realized. Furthermore, the company that was acquired was not profitable and proved to be a cash drain on the parent company. After three years, the CEO resigned. The new CEO sold the acquisition for \$100 million (10% of the price paid) and announced that the company would focus on its original core business. However, by then the company was highly leveraged. A temporary economic downturn made it impossible for the company to service its debt and it declared bankruptcy.

The offices of the company were soon filled with accountants and lawyers representing the interests of the various parties (banks, different categories of bondholders, equity holders, the company, and the board of directors). These people directly or indirectly billed the company about \$10 million per month in fees. The company lost sales that it would normally have made because nobody wants to do business with a bankrupt company. Intangible assets, such as the relationships built up over many years with customers, were lost. Key senior executives, who had many years of valuable experience, left.

After two years and three reorganization attempts, an agreement was reached among the various parties, and a new company with a market capitalization of \$700,000 was incorporated to continue the remaining profitable parts of the business. The shares in the new company were entirely owned by the banks and the bondholders. The shareholders got nothing.

When a major new investment is being contemplated, it is important to consider how well it fits in with other risks taken by the company. Relatively small investments can often have the effect of reducing the overall risks taken because of their diversification benefits. However, a large investment can dramatically increase these risks. Many spectacular corporate failures (such as the fictitious one in Business Snapshot 1.1) can be traced to CEOs who made large acquisitions (often highly leveraged) that did not work out.

1.5.2 *Financial Institutions*

One can argue about how important bankruptcy costs are for the decision making of a non-financial company, but there can be no question that it is crucially important for a financial institution such as a bank to keep its probability of bankruptcy very low. Large banks rely on wholesale deposits and instruments such as commercial paper for their funding. Confidence is the key to their survival. If the risk of default is perceived by the market to be other than very low, there will be a lack of confidence and sources of funding will dry up. The bank will then be forced into liquidation—even if it is solvent in the sense of having positive equity. Lehman Brothers was the largest bankruptcy in U.S. history. Northern Rock was a large failure of a financial institution in the United Kingdom. In both cases, the failure was because there was a lack of confidence and traditional sources of funding dried up.

1.5.3 *Regulation*

Even if, in spite of the arguments we have just given, the managers of a bank wanted to take huge risks, they would not be allowed to do so. Unlike other companies, many financial institutions are heavily regulated. Governments throughout the world want a stable financial sector. It is important that companies and private individuals have confidence in banks and insurance companies when they transact business. The regulations are designed to ensure that the probability of a large bank or an insurance company experiencing severe financial difficulties is low. The bailouts of financial institutions in 2008 during the Global Financial Crisis illustrate the reluctance of governments to let large financial institutions fail. Regulated financial institutions are forced to consider total risks (systematic plus nonsystematic).

Bankruptcy generally arises from losses being incurred. Regulators try to ensure that the capital held by a bank is sufficient to provide a cushion to absorb the losses with a high probability. Suppose, for example, that there is considered to be only a 0.1% probability that a financial institution will experience a loss of \$2 billion or more in a year. Regulators might require the bank to hold equity capital equal to \$2 billion. This would ensure that there is a 99.9% probability that the equity capital is sufficient to absorb the losses. The models used by regulators are discussed in more detail in later chapters.

The key point here is that regulators are concerned with total risks, not just systematic risks. Their goal is to make bankruptcy a highly unlikely event.

1.6 Risk Management by Financial Institutions

There are two broad risk management strategies open to a financial institution (or any other organization). One approach is to identify risks one by one and handle each one separately. This is sometimes referred to as *risk decomposition*. The other is to reduce risks by being well diversified. This is sometimes referred to as *risk aggregation*. Both approaches are typically used by financial institutions.

Consider, for example, the market risks incurred by the trading room of a bank. These risks depend on the future movements in a multitude of market variables (exchange rates, interest rates, stock prices, and so on). To implement the risk decomposition approach, the trading room is organized so that a trader is responsible for trades related to just one market variable (or perhaps a small group of market variables). For example, there could be one trader who is responsible for all trades involving the dollar-yen exchange rate. At the end of each day, the trader is required to ensure that certain risk measures are kept within limits specified by the bank. If the end of the day is approached and it looks as though one or more of the risk measures will be outside the specified limits, the trader must either get special permission to maintain the position or execute new hedging trades so that the limits are adhered to. (The risk measures and the way they are used are discussed in later chapters.)

The risk managers, working in what is termed the middle office of a bank, implement the risk aggregation approach for the market risks being taken. This involves calculating at the end of each day the total risk faced by the bank from movements in all market variables. Hopefully, the bank is well diversified so that its overall exposure to market movements is fairly small. If risks are unacceptably high, then the reasons must be determined and corrective action taken.

Risk aggregation is a key tool for insurance companies. Consider automobile insurance. The insurance company's payout on one particular automobile insurance policy is quite uncertain. However, the payout on 100,000 similar insurance policies can be predicted with reasonable accuracy.

Credit risks are also traditionally managed using risk aggregation. It is important for a financial institution to be well diversified. If, for example, a bank lends 40% of its available funds to a single borrower, it is not well diversified and likely to be subject to unacceptable risks. If the borrower runs into financial difficulties and is unable to make interest and principal payments, the bank could become insolvent.

If the bank adopts a more diversified strategy of lending 0.01% of its available funds to each of 10,000 different borrowers, it is in a much safer position. Suppose that the probability of any one borrower defaulting is 1%. We can expect that close to 100 borrowers will default in the year and the losses on these borrowers will be more than offset by the profits earned on the 99% of loans that perform well. To maximize the benefits of diversification, borrowers should be in different geographical regions and different industries. A large international bank with different types of borrowers all over the world is likely to be much better diversified than a small bank in Texas that lends entirely to oil companies.

But, however well diversified a bank is, it is still exposed to systematic risk, which creates variations in the probability of default for all borrowers from year to year. Suppose that the probability of default for borrowers in an average year is 1%. When the economy is doing well, the probability of default is less than this and when there is an economic downturn it is liable to be considerably more than this. Models for capturing this exposure are discussed in later chapters.

Since the late 1990s, we have seen the emergence of an active market for credit derivatives. Credit derivatives allow banks to handle credit risks one by one (risk decomposition) rather than relying solely on risk diversification. They also allow banks to buy protection against the overall level of defaults in the economy. However, for every buyer of credit protection there must be a seller. Many sellers of credit protection, whether on individual companies or on portfolios of companies, took huge losses during the Global Financial Crisis that started in 2007. (This crisis is discussed further in Chapter 7.)

1.7 Credit Ratings

Credit rating agencies provide information that is widely used by financial market participants for the management of credit risks. A credit rating is a measure of the credit quality of a debt instrument such as a bond. However, the rating of a corporate or sovereign bond is often assumed to be an attribute of the bond issuer rather than of the bond itself. Thus, if the bonds issued by a company have a rating of AAA, the company is often referred to as having a rating of AAA.

The three major credit rating agencies are Moody's, S&P, and Fitch. The best rating assigned by Moody's is Aaa. Bonds with this rating are considered to have almost no chance of defaulting. The next best rating is Aa. Following that come A, Baa, Ba, B, Caa, Ca, and C. The S&P ratings corresponding to Moody's Aaa, Aa, A, Baa, Ba, B, Caa, Ca, and C are AAA, AA, A, BBB, BB, B, CCC, CC, and C, respectively. To create finer rating measures Moody's divides the Aa rating category into Aa1, Aa2, and Aa3; it divides A into A1, A2, and A3; and so on. Similarly S&P divides its AA rating category into AA+, AA, and AA-; it divides its A rating category into A+, A, and A-; and so on. Moody's Aaa rating category and S&P's AAA rating are not subdivided, nor usually are the two lowest rating categories. Fitch's rating categories are similar to those of S&P.

There is usually assumed to be an equivalence between the meanings of the ratings assigned by the different agencies. For example, a BBB+ rating from S&P is considered equivalent to a Baa1 rating from Moody's. Instruments with ratings of BBB- (Baa3) or above are considered to be *investment grade*. Those with ratings below BBB- (Baa3) are termed *noninvestment grade* or *speculative grade* or *junk bonds*.

We will learn a lot more about credit ratings in later chapters of this book. For example, Chapter 7 discusses the role of ratings in the Global Financial Crisis. Chapters 25 and 26 provide information on how ratings are used in regulation. Chapter 17 provides statistics on the default rates of companies with different credit ratings. Chapter 19 examines the extent to which the credit ratings of companies change through time.

Summary

An important general principle in finance is that there is a trade-off between risk and return. Higher expected returns can usually be achieved only by taking higher risks. In theory, shareholders should not be concerned with risks they can diversify away. The expected return they require should reflect only the amount of systematic (i.e., non-diversifiable) risk they are bearing.

Companies, although sensitive to the risk–return trade-offs of their shareholders, are concerned about total risks when they do risk management. They do not ignore the unsystematic risk that their shareholders can diversify away. One valid reason for this is the existence of bankruptcy costs, which are the costs to shareholders resulting from the bankruptcy process itself.

For financial institutions such as banks and insurance companies there is another important reason: regulation. The regulators of financial institutions are primarily concerned with minimizing the probability that the institutions they regulate will fail. The probability of failure depends on the total risks being taken, not just the risks that cannot be diversified away by shareholders. As we will see later in this book, regulators aim to ensure that financial institutions keep enough capital for the total risks they are taking.

Two general approaches to risk management are risk decomposition and risk aggregation. Risk decomposition involves managing risks one by one. Risk aggregation involves relying on the power of diversification to reduce risks. Banks use both approaches to manage market risks. Credit risks have traditionally been managed using risk aggregation, but with the advent of credit derivatives the risk decomposition approach can be used.

Further Reading

Markowitz, H. “Portfolio Selection.” *Journal of Finance* 7, no. 1 (March 1952): 77–91.

Ross, S. “The Arbitrage Theory of Capital Asset Pricing.” *Journal of Economic Theory* 13, no. 3 (December 1976): 341–360.

Sharpe, W. “Capital Asset Prices: A Theory of Market Equilibrium.” *Journal of Finance* 19, no. 3 (September 1964): 425–442.

Smith, C. W., and R. M. Stulz. “The Determinants of a Firm’s Hedging Policy.” *Journal of Financial and Quantitative Analysis* 20 (1985): 391–406.

Stulz, R. M. *Risk Management and Derivatives*. Mason, OH: South-Western, 2003.

Practice Questions and Problems (Answers at End of Book)

- 1.1 An investment has probabilities 0.1, 0.2, 0.35, 0.25, and 0.1 of giving returns equal to 40%, 30%, 15%, –5%, and –15%. What are the expected returns and the standard deviations of returns?

- 1.2 Suppose that there are two investments with the same probability distribution of returns as in Problem 1.1. The correlation between the returns is 0.15. What is the expected return and standard deviation of return from a portfolio where money is divided equally between the investments?
- 1.3 For the two investments considered in Figure 1.2 and Table 1.2, what are the alternative risk-return combinations if the correlation is (a) 0.3, (b) 1.0, and (c) -1.0 ?
- 1.4 What is the difference between systematic and nonsystematic risk? Which is more important to an equity investor? Which can lead to the bankruptcy of a corporation?
- 1.5 Outline the arguments leading to the conclusion that all investors should choose the same portfolio of risky investments. What are the key assumptions?
- 1.6 The expected return on the market portfolio is 12% and the risk-free rate is 6%. What is the expected return on an investment with a beta of (a) 0.2, (b) 0.5, and (c) 1.4?
- 1.7 “Arbitrage pricing theory is an extension of the capital asset pricing model.” Explain this statement.
- 1.8 “The capital structure decision of a company is a trade-off between bankruptcy costs and the tax advantages of debt.” Explain this statement.
- 1.9 What is meant by risk aggregation and risk decomposition? Which requires an in-depth understanding of individual risks? Which requires a detailed knowledge of the correlations between risks?
- 1.10 A bank’s operational risk includes the risk of very large losses because of employee fraud, natural disasters, litigation, and so on. Do you think operational risk is best handled by risk decomposition or risk aggregation? (Operational risk will be discussed in Chapter 20.)
- 1.11 A bank’s profit next year will be normally distributed with a mean of 0.6% of assets and a standard deviation of 1.5% of assets. The bank’s equity is 4% of assets. What is the probability that the bank will have a positive equity at the end of the year? Ignore taxes.
- 1.12 Why do you think that banks are regulated to ensure that they do not take too much risk but most other companies (for example, those in manufacturing and retailing) are not?
- 1.13 List the bankruptcy costs incurred by the company in Business Snapshot 1.1.
- 1.14 The return from the market last year was 10% and the risk-free rate was 5%. A hedge fund manager with a beta of 0.6 has an alpha of 4%. What return did the hedge fund manager earn?

Further Questions

- 1.15 Suppose that one investment has a mean return of 8% and a standard deviation of return of 14%. Another investment has a mean return of 12% and a standard

deviation of return of 20%. The correlation between the returns is 0.3. Produce a chart similar to Figure 1.2 showing alternative risk–return combinations from the two investments.

- 1.16 The expected return on the market is 12% and the risk–free rate is 7%. The standard deviation of the return on the market is 15%. One investor creates a portfolio on the efficient frontier with an expected return of 10%. Another creates a portfolio on the efficient frontier with an expected return of 20%. What is the standard deviation of the return on each of the two portfolios?
- 1.17 A bank estimates that its profit next year is normally distributed with a mean of 0.8% of assets and the standard deviation of 2% of assets. How much equity (as a percentage of assets) does the company need to be (a) 99% sure that it will have a positive equity at the end of the year and (b) 99.9% sure that it will have positive equity at the end of the year? Ignore taxes.
- 1.18 A portfolio manager has maintained an actively managed portfolio with a beta of 0.2. During the last year, the risk–free rate was 5% and major equity indices performed very badly, providing returns of about -30% . The portfolio manager produced a return of -10% and claims that in the circumstances it was good. Discuss this claim.