1

Introduction to Enterprise Risk Management

1.1 Summary

In this chapter, we present an overview of enterprise risk management (ERM). We begin by discussing the concepts of risk and uncertainty. We then review some of the more important historical developments in different areas of risk management, propose a definition for ERM, and show how ERM has its origins in all of the individual areas of risk management that came before.

We discuss how ERM can be implemented as an ongoing process, which is optimally built into the operations of an organization from the top down through a risk governance framework. Stages of the ERM cycle include risk identification and analysis, risk evaluation, and risk treatment. Each of these stages is introduced in this chapter and then developed in more detail in subsequent chapters.

1.2 Risk and Uncertainty

Suppose you are planning an outdoor party. You check the weather forecast and see that there is a 10% probability that it will rain during the event. That means there is a risk that your party will not be successful, unless you provide some shelter for your guests. It is uncertain whether or not the shelter will be needed, but you have very good information about the uncertainty – there is a 90% chance that the shelter will not be needed and a 10% chance that it will save your party from being washed out. More problematic in planning your party are the things about which you have no information. Perhaps, without your knowledge, two of the guests hate each other; if they meet at your party, a big fight will ensue, ruining the event. That is a risk out of the blue – you did not know that there was any risk of a fight, and, even if you had, it would have been hard to assign probabilities to the events that unfolded. Both the rain and the fight represent risks to the success of your party, but the risks are different in nature. The risk of rain is much easier to handle – you can decide whether the risk justifies the cost of renting a temporary shelter for the day. The risk of a fight is a more difficult problem – you cannot prepare for it, or avoid it, if you are unaware that the risk exists.

Some writers use the term **risk** specifically for events which are identifiable and quantifiable, like the risk of rain at the party, and **uncertainty**, or **Knightian uncertainty**[1] for events which are not foreseeable, or are impossible to quantify, like the fight risk. In this text, we will use the term 'risk' in the broadest sense to encompass all uncertainty, but we will also refer to Knightian uncertainty when it is useful to make the distinction.

Risk is an essential feature of all business. A manufacturer must invest substantial amounts in research and development, equipment, and raw materials to establish a new product line, in the expectation that their product will prove profitable but allowing for the possibility that it does not. A bank that wishes to avoid all risk would never lend money to its customers. An insurer pursues risk because the purpose of insurance is to profit from managing the risks transferred to it by its customers.

In the business context, taking a risk can have a positive or negative outcome. In common usage, 'risk' is often used to refer only to the possibility of an unwelcome outcome; we refer to this type of risk as **downside risk**.

For many years, business education sought ways to incentivize higher risk-taking by managers in order to benefit from the extra returns attainable. The reasoning was based on the theory that, although managers tended to be risk-averse, businesses should be risk-neutral – meaning (loosely) that any risk that was expected to generate profit was worth pursuing, even if the potential downside could be catastrophic to the firm.[2]

More recently, it has become better recognized that, while some risk-taking is necessary for all business, it is also appropriate to manage those risks, avoiding some and perhaps mitigating others. Modern organizations aim to operate in a risk corridor: not too little, or the business cannot thrive, but not too much, in case the business is destroyed by foreseeable adverse events. The **risk tolerance** of the business sets the upper and lower limits of the risk corridor, and the **risk appetite** describes the target level of risk within the corridor under normal circumstances (although it should be noted that the precise definitions

[1] After Frank Knight, who distinguished risk and uncertainty in his 1921 book *Risk, Uncertainty and Profit*.
[2] For more in-depth explanation of these arguments, see, e.g., Grossman and Hart (1982) and Mayers and Smith (1982).

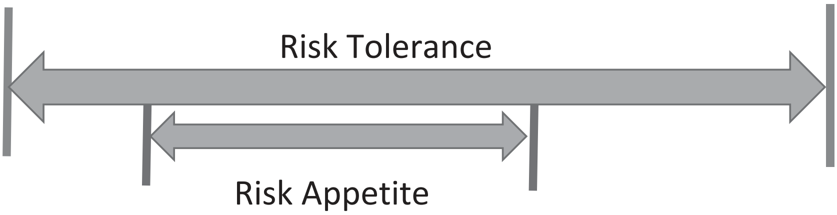

Figure 1.1 Risk appetite and risk tolerance

of risk appetite and risk tolerance vary across the literature) (see Figure 1.1). Risk appetite and risk tolerance are discussed in more detail in Section 1.12.1.

The objective of risk management is closely tied to the management of the **capital** of a firm. Broadly, capital refers to the funds available to absorb unexpected losses. Typically, a firm's capital includes all of its equity (excess of assets over liabilities) and some, or all, of its long-term debt. Following common usage, we will use the term 'capital' in this very loose sense. Where more precision is required, the calculation depends on the context. In some cases, long-term debt will be excluded, as it is less flexible than equity. In other cases, long-term debt is included, as it can be applied to absorb losses in the short term.

The term '**risk management**' is used by accountants, actuaries, project managers, quants (quantitative finance specialists), business leaders, engineers, and professional risk managers, and each of these groups has adopted a slightly different interpretation of the term. In the following sections, we describe the meaning and origins of a range of types of risk management, and we then discuss how the various strands of risk management converge in enterprise risk management.

1.3 Insurance Risk Management

The field of **risk management and insurance** (RMI) emerged in the 1960s as the study of the impact of insurable risk, or **pure risk**, on firms. Pure risks are all downside, such as the loss from windstorms or cyberattacks. Many pure risks are insurable, which means that firms may transfer the risk to an insurer, but at a cost. The RMI discipline was essentially a subfield of corporate finance, as it developed methods for determining the optimal financial approach to pure risk in reference to its impact on share values. That means that the firm might seek opportunities to reduce risks, but only if the analysis

Cambridge University Press & Assessment
978-1-009-09846-5 — Quantitative Enterprise Risk Management
Mary R. Hardy , David Saunders
Excerpt
More Information

showed the cost of risk reduction to be less than the benefit. For example, a firm's risk of losses from fire damage may be reduced through insurance or through better-quality building materials or fire prevention strategies; the calculation the firm should make is to weigh up the costs and risks to decide on the best approach.

The choices available for managing pure risk are:

- Transfer – typically by purchasing insurance.
- Retain – self-insure.
- Avoid – by not engaging in the risky activity.
- Reduce – take actions to reduce the frequency and/or severity of loss.

These are not exclusive, as we shall see in later chapters. The firm might insure against severe losses but self-insure minor losses.

The analysis of pure risk, including the analysis of risk reduction or elimination (through insurance or otherwise), is now a major element of ERM. The options for managing risk in an ERM framework include the choices for pure risk, but ERM addresses both risk and risk management more comprehensively. In Section 1.12.4, we describe an expanded list of management strategies that recognizes potential rewards as well as losses.

1.4 Financial Risk Management

Two strands of modern finance deal, in different ways, with risk analysis and risk management. The first, is the quantitative assessment of the risk-return trade-off through modern portfolio theory. The second is pricing and hedging financial options.

1.4.1 Portfolio Risk

Portfolio analysis traditionally measures the risk associated with a portfolio of assets using the standard deviation of the portfolio return. The risk-reward trade-off is captured in the calculation of the efficient frontier, where the expected return on a portfolio of risky assets is plotted against the risk, as represented by the standard deviation. An example is given in Figure 1.2.

The efficient frontier represents the set of optimal portfolios, where optimal means that no other combination offers a higher expected return for any given standard deviation. When a risk-free asset is included in the portfolio mix, the efficient frontier is a straight line with y-intercept equal to the risk-free rate of return, denoted here by r_f. This is illustrated by the grey line in Figure 1.2.

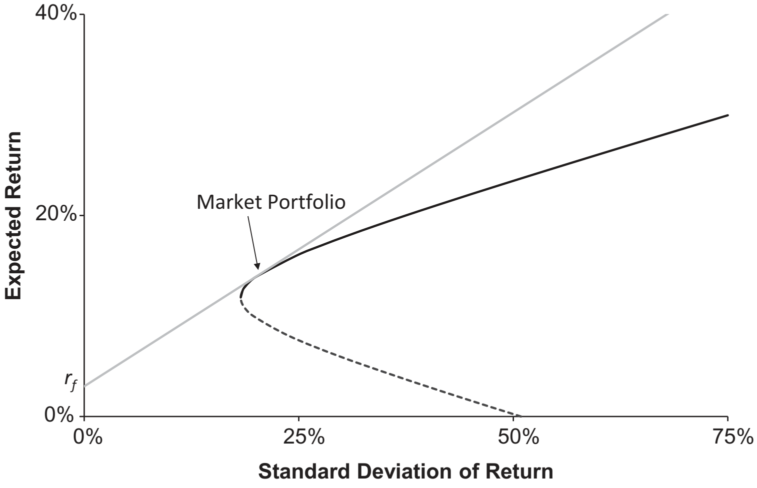

Figure 1.2 Illustration of a portfolio frontier – dark line is the efficient frontier for the risky asset portfolios; the straight line includes the risk-free asset.

There is a unique point on the efficient frontier that is achieved with zero investment in the risk-free asset. This is illustrated in Figure 1.2 by the point where the straight line is tangent to the original frontier. This point represents the mean and standard deviation of the return on the market portfolio, denoted by R_M; the market portfolio is a portfolio of all the investments represented in the market, each held in proportion to their total market capitalization. Any point on the efficient frontier can be achieved with a combination of an investment in the risk-free asset (possibly a negative, or short position), and an investment in the market portfolio. This implies that, under the model assumptions, there can be no reason to invest in anything other than the market portfolio, in respect of the risky asset allocation.

Now consider a single asset, for example, an investment in shares of XYZ company. The capital asset pricing model (CAPM) shows that the random return on shares of XYZ, denoted by R_X, must be related to the return on the market portfolio, denoted by R_M, as:

$$E[R_X] = r_f + \beta_X \left(E[R_M] - r_f \right). \tag{1.1}$$

The beta term is defined as follows: let $\rho_{X,M}$ denote the correlation between the return on the market portfolio and the return on the XYZ shares, and let σ_X and σ_M represent the standard deviation of returns on shares of XYZ and on the market portfolio, respectively. Then

$$\beta_X = \frac{\rho_{X,M}\,\sigma_X}{\sigma_M}. \tag{1.2}$$

Hence, the beta of the asset is a measure of the dependence of the returns of the XYZ shares on the market overall return; it is sometimes referred to as the **systematic risk**.

Non-systematic risk is the variation in the asset value that is independent of the market as a whole. In theory, this risk can be eliminated for investors through diversification, so the systematic risk of a portfolio, or of any investment, is considered the key risk.

There are some important considerations missing from this analysis. The first is that standard deviation, in general, is not a very good measure of risk. One reason is that it counts both upside and downside deviations equivalently; in addition, it may not capture rare events that could have significant impact. Furthermore, the efficient frontier analysis and equation (1.1) assume that the expected returns, standard deviations, and correlations of the individual investments are known by all and are (reasonably) static over time. In reality, each of these parameters can only be estimated, with a lot of associated uncertainty, and the parameters change significantly over time.

1.4.2 Options and Derivatives

The second major development in financial risk grew out of the seminal work of Black, Scholes, and Merton (BSM) in the field of options and other derivatives.

A derivative is an asset with a payoff that is dependent on another (underlying) security. A European call option with strike price K and term T years, written on an underlying security with price S_t at t, pays out the greater of $S_T - K$ and 0 at time T. A European put option with strike price K, term T years, and underlying asset price process S_t, pays out the greater of $K - S_T$ and 0 at T. Options are described and explored in depth in Chapter 9. The BSM work combined a scientific approach to pricing options with a consistent method for hedging the resulting cash flows. That is, as well as measuring risk, the approach offers important insight into the management of risk.

The BSM analysis of option pricing was significant in risk management for several reasons. The first is that options and other derivatives can be used to hedge risk, and so the development of scientific pricing methods expanded the options market and, therefore, the availability of pseudo-insurance within capital markets. For example, an investor might hold 10 shares in Gaggle.com with a current value \$100 per share. Suppose she wishes to ensure that, in six months her portfolio does not fall in value below \$850. She could purchase six-month put options with a strike price of \$85 per share at a cost of around \$1.50 per share or \$15 in total. In the (unlikely) event that Gaggle.com shares fall below \$85 at the maturity date, the put option would pay the shortfall, maintaining the portfolio value. If the share price is higher than \$85 per share at the end of the six-month period, then the options expire without value.

More exotic derivatives can be used to hedge more exotic risks. Derivatives can be used to offset currency risk for businesses working in different markets or to offset default risk using credit derivatives. More simple forward and futures contracts, which set a fixed price for a commodity some time in advance of a transaction, can be used by businesses to lock in prices – for example, a farmer might lock-in the price of corn or pork bellies that they plan to sell, or an airline might lock-in the price of fuel that it plans to buy. Unlike options, under futures contracts the agreed price is paid at maturity even if the market price is more favourable. It is important to note that derivatives can be used for speculation as well as for hedging, unlike traditional insurance.

The second major contribution of the BSM analysis was to provide a systematic framework for hedging derivatives that is integrated with the pricing methodology. For example, consider the Black–Scholes formula for the price of a T-year put option, on a non-dividend paying stock with price S_t at time t, and with strike price K. The price at time 0 is

$$p_0 = Ke^{-rT}\Phi(-d_2(0,T)) - S_0\Phi(-d_1(0,T)),$$

where r is the continuously compounded risk-free rate of interest, Φ is the distribution function of the standard normal (N(0, 1)) distribution, and, for any t_1, t_2 where $0 \leq t_1 < t_2$

$$d_1(t_1,t_2) = \frac{\log(S_{t_1}/K) + (r + \sigma^2/2)\,(t_2 - t_1)}{\sigma\sqrt{t_2 - t_1}}, \tag{1.3}$$

$$d_2(t_1,t_2) = d_1(t_1,t_2) - \sigma\sqrt{t_2 - t_1}. \tag{1.4}$$

Here, $\sigma > 0$ is a measure of the volatility of the returns on the underlying stock. This formula[3] is also used to find the price at intermediate dates $0 < t < T$,

$$p_t = Ke^{-r(T-t)}\Phi(-d_2(t,T)) - S_t\,\Phi(-d_1(t,T)). \tag{1.5}$$

Now, although the formula is written in terms of the N(0, 1) distribution function, the development is based on the cost of replicating the option payoff,

[3] Note that the t_1 and t_2 arguments are usually omitted from the d_1, d_2 notation and that in finance texts N is generally used in place of Φ.

not on evaluating the expected costs under real-world probabilities. That is, under the BSM assumptions, the option seller can eliminate their risk by holding the replicating portfolio at time t, which is then costlessly rebalanced throughout the rest of the term. At time t, the replicating portfolio comprises $\Phi(-d_2(t,T))$ units of a $(T-t)$ year risk-free zero-coupon bond, with total face value K, together with $-\Phi(-d_1(t,T))$ units (the minus sign indicating that this is a short position) of the underlying stock at t, where each unit of stock has price S_t. The total value of this replicating portfolio is then exactly p_t at t, as given in equation (1.5). The $-\Phi(-d_1(t,T))$ term represents the derivative of the option price with respect to the stock price at time t; this is the option **delta**.

This is the real beauty of the Black–Scholes formula, that it simultaneously offers a valuation framework and a risk management strategy. The fundamental principle underlying the formula, which applies far more widely than vanilla European options, is the no-arbitrage principle, which requires that two portfolios with identical payoffs must have identical prices. There are some strong assumptions involved in the BSM framework, but years of research and experience have shown that the approach can be useful in practice, even where the underlying assumptions are breached.

1.5 Actuarial Risk Management

Actuaries typically work on quantitative analysis for insurance companies and pension plans. The traditional actuarial role in insurance includes determining premiums, setting reserves, and contributing to investment and risk management strategies.

Actuarial risks involve a combination of financial and non-financial uncertainty. For example, a term life insurance policy involves financial uncertainty with respect to the return earned by investing the premiums, and non-financial uncertainty in the benefit payout, which is contingent on the death of the insured life. Similarly, an auto insurance policy involves financial uncertainty in the value of accumulated premiums, non-financial uncertainty in the claim frequency (that is, the probability of a claim), and some financial and non-financial uncertainty in the claim severity because, when a claim arises, the amount will depend not only on the nature of the accident but also, potentially, on the inflation rates applying in the time between the loss event and the claim settlement. Most non-life insurance is quite short-term and involves little financial risk, but some life insurance liabilities are much more sensitive to financial risk than demographic risk and the risk management techniques for those are very similar to the techniques used by investment banks.

Actuarial risk management integrates RMI and financial risk management, but from a different perspective. The objective of RMI is to use insurance as a tool in risk management with the ultimate goal of maximizing share value of the firm. Financial risk management generally takes the perspective of pure investment enterprises, such as banks or hedge funds, where the objective is to create investment return with mitigation of downside risk. Actuarial risk management takes the perspective of the insurer but adds a broader public responsibility beyond maximizing profitability. For example, the actuary is responsible for the security of the policyholders' contingent benefits, and for ensuring that there is broad equity of treatment of different groups of policyholders. Some actuaries work for pension plans, where profit is not an objective. In this case, actuarial risk management consists of understanding the risks involved and advising plan managers as to the best strategy for balancing cost and benefit security without excessive risk-taking.

In the 1990s, actuaries developed **dynamic financial analysis** (DFA) techniques as an extension of their traditional risk management toolkit. DFA used integrated projections of assets and liabilities to assess the net effect of different investment and risk scenarios on the liabilities. The key insight in DFA was that the uncertainty in asset values should not be analysed separately from the uncertainty in liability values. Furthermore, different liability classes could be aggregated, allowing simultaneous modelling of risk from all sources with respect to a portfolio of liabilities and supporting assets. The scenarios analysed comprised paths for all key factors influencing cash flows (such as interest rates, claim frequency, and inflation rates), allowing the insurer to stress test their financial strength through a range of adverse scenarios. Additionally, the cash flows could be projected through a large number of stochastically simulated scenarios, thereby offering a probabilistic analysis of the portfolio solvency over a specified time horizon.

A major contribution to ERM was the recognition that jointly modelling assets and liabilities is essential where they are connected, for example, through common dependencies on capital market factors or on other factors influencing claim frequency and severity. The idea of projecting cash flows with dynamic control mechanisms, which means that the projection in each future year depends to some extent on the experience up to that year, is still a widely used tool in ERM.

1.6 Asset-Liability Management

Traditionally, financial risk management has focussed on asset values, and insurance risk management on liabilities. Asset liability management (ALM) considers risks arising from mismatching of assets and liabilities. It is an integral part of dynamic financial analysis and, hence, of actuarial risk management.

The main focus of ALM is interest rate risk. Often, the assets supporting an uncertain liability have different sensitivity to interest rate movements than the liability itself. This can create a situation where the assets appear to be sufficient to meet liabilities, but after a shift in interest rates the asset values fall below the liability values. We illustrate this with the following example.

Example 1.1 A company has a debt of 1,000 due in 10 years ('the liability'). It makes provision for the debt by purchasing a zero-coupon bond with term 8 years and face value 920 ('the asset').

The company demonstrates that the asset is sufficient to meet the liability based on market values. Market values are determined by discounting the cash flows at the risk-free rate of interest. Assume that the risk-free rate of interest (annually compounded) is 3.25% for an 8-year term and 3.5% for a 10-year term.

(a) Calculate the market value of the surplus of assets over liabilities.
(b) Assume that one year later interest rates have fallen, such that the 9-year risk-free rate of interest (annually compounded) is 3.0% and the 7-year risk-free rate is 2.8%.
    Calculate the change in value from time 0 to time 1 in: (i) the asset and (ii) the liability.
(c) Discuss the impact of the interest rate movement on the sufficiency of the assets with respect to meeting the liabilities.

Solution 1.1

(a) At time 0, the market value of the liability is

$$V_L(0) = 1{,}000(1.035)^{-10} = 708.92.$$

The market value of the asset is $V_A(0) = 920(1.0325)^{-8} = 712.31$.
The surplus is, therefore, $V_A(0) - V_L(0) = 3.39$.

(b) (i) At time 1, the market value of the asset is $V_A(1) = 920(1.028)^{-7} = 758.29$. The asset value has increased by 6.455% over the year.
(ii) The market value of the liability is $V_L(1) = 1{,}000(1.030)^{-9} = 766.42$. The liability value has increased by 8.11% over the year.