Part I
Risk and the ART Market

1 Overview of Risk Management

Risk management is a dynamic and well-established discipline practiced by many companies around the world. Traditional forms of risk management – loss control, loss financing and risk reduction, arranged through mechanisms such as insurance and derivatives – have been actively used by companies for many decades, and are an essential element of most corporate strategies. But newer forms of risk protection – including those from the **alternative risk transfer**[1] (ART) market, the *combined risk management marketplace for innovative insurance and capital market solutions* – often surface as viable, flexible and cost-efficient options. In fact, some firms already use ART mechanisms to supplement their traditional risk management strategies; many others, however, have yet to take advantage of the benefits offered by the marketplace. Regardless of a firm's specific approach to risk management, it should always consider ART-related solutions so that it has complete knowledge of all available options and can make the best, most informed, decisions possible.

Our discussion in this book is on the ART market, its function, participants and products, its advantages and disadvantages, and its future prospects. Before considering the specifics of the marketplace, however, we review some of the essential concepts of risk and risk management; this helps to provide a proper framework for the material that follows. In the remainder of this chapter we explore issues related to risk and return, general risk management processes and techniques, and fundamental risk concepts and measures.

1.1 RISK AND RETURN

Risk is a broad, complex and vitally important topic that touches on virtually all aspects of modern corporate operation. Although we shall consider matters in greater detail as we progress through the text, we begin by defining **risk,** in its most general form, as *uncertainty associated with a future outcome or event*. To apply this more specifically to corporate activities, we can say that risk is the *expected variance in profits, losses, or cash flows* arising from an uncertain event. Other terms commonly associated with risk – such as peril and hazard – are often encountered in the risk management industry (indeed, we shall also use them throughout the text); they are, however, distinct. A **peril**, for instance, is a cause of loss, while a **hazard** is an event that creates, or increases, peril. While both have a bearing on risk, risk itself is a broader concept. Companies are exposed to a wide range of risks that might, at any time, include such things as business interruption, catastrophic and non-catastrophic property damage, product recall/liability, directors and officers liability, credit default/loss, workers compensation, environmental liability, and so on. These risks must be managed if the market value of the company is to be increased – or, at a minimum, if the probability of financial distress is be lowered. Some of the risks can be retained as part of core business operations, while others are best transferred elsewhere – but only when it is cost-effective to do so.

[1] Emboldened items are listed in the glossary.

We shall consider risks in more detail later, but we begin by classifying them broadly as operating risks and financial risks:

- **Operating risk** The risk of loss arising from the daily physical (non-financial) operating activities of a firm.
- **Financial risk** The risk of loss arising from the financial activities of a firm.

Operating and financial risks can be decomposed further. For example, within the general category of operating risks we can consider subclasses such as personal liability and commercial property/casualty liability. Within commercial property/casualty (P&C) liability we might differentiate between losses related to commercial property (direct/indirect), machinery, transportation (inland/marine), crime, commercial liability, commercial auto, workers compensation, and employers' liability. Similar decomposition is possible within the category of financial risks, where we might first divide exposures into credit risk, market risk, liquidity risk, and model risk. A category such as market risk might then be segregated into directional risk, volatility risk, time decay risk, curve risk, basis risk, spread risk, correlation risk, and so on.

We can also categorize financial and operating risks as being pure or speculative.

- **Pure risk** A risk that only has the prospect of downside, i.e., loss.
- **Speculative risk** A risk that has the possibility of upside or downside, i.e., gain or loss.

Regardless of the taxonomy, the central point is that risk comes in many forms, a factor that becomes apparent and important in the risk management process.

A company creates goods and services that it sells to clients in order to generate returns. These returns are used to expand business (e.g., internal funding via retained earnings) and compensate equity investors who have supplied the equity risk capital needed to fund productive assets (e.g., factories, machinery, intellectual property). Investors must be compensated for supplying risk capital. Generally speaking, they require returns related to the inherent riskiness of the company: the riskier the company, the greater the return (or **risk premium**) the investors demand. Whether a company is risky or not, however, investors will always seek the maximum possible return. This means a key corporate goal is the maximization of **enterprise value** (EV), which we define as the sum of a firm's expected future net cash flows (NCFs), discounted back to the present at an appropriate discount rate (e.g., risk-free rate plus relevant risk premium). We summarize this as:

$$\mathrm{EV} = \sum_{t=1}^{n} \frac{\mathrm{NCF}_t}{(1 + r)^t},$$

where NCF(t) is the expected net cash flow at time t, and r is the discount rate, comprising a risk-free rate $r(\mathrm{f})$ and a risk premium $r(\mathrm{p})$.

We shall explore this in more detail in Chapter 2, but note for the moment that expected NCFs can be impacted by the expected size, timing, and variability of cash flows. Risk can also change all three dimensions, meaning that it can alter the value of the firm. In fact, unexpected changes in NCF can be quite damaging to enterprise value, and protecting against such changes surfaces as one of the primary motivations for active risk management.

1.2 ACTIVE RISK MANAGEMENT

Companies need to control their exposure to risk in the normal course of business. While speculative risks can bring gains or losses, pure risks generate only losses. In either case, failure to focus on the potential downside through active risk management means that firms face financial uncertainty – to the possible detriment of shareholders, creditors, and other stakeholders who might be economically impacted if a firm becomes insolvent. Risk management is an important discipline because, unlike the world presented through pure corporate finance theory,[2] shareholders cannot effectively manage a firm's risks by themselves. Investors face information asymmetries, lack access to the same risk transfer mechanisms as a corporate entity (which faces lower friction costs), and cannot influence or control corporate investment policy. Accordingly, active risk management is not only desirable, but also necessary, if corporate value is to be maximized in practice.

There are, of course, many reasons why a company should actively, rather than passively, manage its risks. An active approach to risk management – centered on control, retention, transfer and/or hedging – can help to:

- provide funds when they are most needed, helping to ensure a liquid position and minimizing the possibility of **financial distress** – a state of financial weakness that might include a higher cost of capital, poorer supplier terms, lower liquidity, and departure of key personnel;
- lower cash flow volatility and minimize the disruption of investment plans;
- reduce the possibility of **underinvestment**, or the process of directing capital toward projects with lower returns and risks (to the benefit of creditors rather than equity investors);
- stabilize revenue streams and thus benefit from specific tax treatment (e.g., asymmetrical tax structures where firms with more volatile revenue and profit performance pay greater taxes);
- create more stable earnings, which often helps to generate higher stock price valuations.

It is increasingly common in the corporate world of the twenty-first century for companies to implement a risk management process to control risks. It is important to stress at the outset that the exercise relates to *controlling* risks, not *eliminating* them. This is an important distinction because risk is not inherently bad, and is not a variable that must be removed from corporate operations at any cost. As we shall see in subsequent chapters, there are times when it makes sense for a company to retain, and even increase, its risk exposure, as this helps to increase the value of the firm to shareholders. Instead, the focus is on controlling – that is, understanding and closely managing – risk exposures, so that stakeholders are fully aware of how the firm might be impacted. The essential element of controlling risks is ensuring that no surprises arise. Losses are acceptable if the possibility that they may occur is understood by stakeholders, and if the appropriate economic evaluation occurs. Indeed, risk is a game of chance: speculative risks will produce favorable outcomes and losses, pure risk events only losses. The risk-taking firm must expect both, and if it is controlling its exposures properly it is helping to increase its value. Unexpected losses that occur when the company and its stakeholders have no idea that the firm is exposed to particular types, or amounts, of risk, must be regarded as unacceptable; this essentially means that risk is not being controlled. The development and use of a formalized risk management process must therefore be a central part of overall corporate operations and governance.

[2] For instance, in a Modigliani and Miller framework.

1.2.1 Risk management processes

The standard **risk management process** can be seen as a four-stage process centered on identification, quantification, management, and reporting. Each element is a vital link in the chain and must be implemented correctly in order to be effective.

- **Risk identification** The identification process centers on defining and identifying all of the firm's actual, perceived, or anticipated risks. In a large firm, this might encompass dozens of financial and operating risk drivers, implying a significant degree of complexity. In some cases risks are readily identifiable, at other times they are more difficult to discern. For instance, a firm that produces goods in the US for dollars and sells them in Japan for yen is exposed to changes in the \$/¥ foreign exchange rate, and identifying this risk is relatively simple. Likewise, a company that has a factory located in the path of hurricanes can easily identify potential exposure to catastrophic damage. Alternatively, a firm that has to purchase power in the spot electricity market when temperatures rise above 95 °F is actually exposed to the absolute level of, and correlation between, electricity prices and temperature; in this case the different dimensions of exposure are somewhat more difficult to identify. This stage of the process is vital, of course, as failure to properly identify all financial or operating risks impacting the firm may lead to surprise losses (e.g., those coming from an 'unknown' source).
- **Risk quantification** The quantification process determines the financial impact that risks can have on corporate operations. This is typically done through various quantitative tools. Returning to the \$/¥ example, a company with a foreign exchange exposure will be interested in knowing, as precisely as possible, the impact of the risk on its profit and loss (P&L) account (e.g., a 5% decline in the value of the yen might produce a \$5m loss). The company with a factory in the hurricane path may need to quantify a number of different types of scenarios, including smaller losses from temporary business interruption (e.g., if a hurricane causes damage that forces it to suspend operations for 2 months) to larger losses from total destruction (e.g., the hurricane destroys the facility beyond repair). Specific techniques for measuring the financial impact of risks vary widely, and depend largely on the nature of the underlying exposures. Some, such as credit and market risks, can be measured through financial mathematics based on analytic computation, closed-form pricing models, and simulation methods. Others, such as high-frequency insurance-related risks, can often be estimated by using actuarial techniques; certain low-frequency insurance exposures, such as catastrophic risks, may be modeled through simulation.
- **Risk management** After risks have been identified and quantified, they must be managed. Through the core process of active decision-making, a firm must decide whether it will control, retain, eliminate or expand its exposures. For instance, a firm may decide that it is comfortable retaining a potential loss (or gain) of \$10m on its \$/¥ foreign exchange exposure and will constrain it at that level; alternatively, if it wants to face zero chance of loss, it might eliminate the risk entirely (for a price). Similarly, the potential cost of sustaining partial or complete destruction as a result of a powerful hurricane may be too great for the firm, so it might decide to transfer the exposure entirely. Risk management decisions ultimately depend on several variables, including the financial resources of the firm, the operating philosophy of management, the expectations of shareholders, and the costs and benefits of various risk strategies. We consider these points in the section below.
- **Risk monitoring** Once the firm has decided how it wants to manage its risk profile, it must actively monitor its exposures. This means regularly tracking and reporting both risks

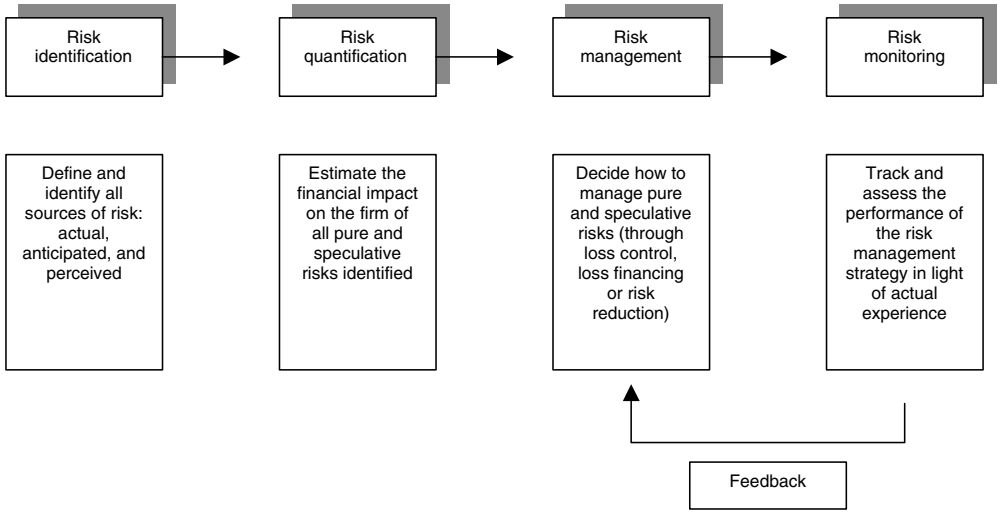

Figure 1.1 The generic risk management process

and risk decision experience, and communicating information internally and externally so that interested parties (e.g., executive management, board directors, regulators, creditors and investors) are aware of any possible upside or downside. Good monitoring is especially important for internal decision-makers, who require feedback in order to assess, and even adjust, their decisions. Thus, the $/¥ exposure that the firm has chosen to retain must be measured and reported regularly (e.g., daily, weekly) so that managers are aware of its size and potential impact as the market moves and the risk position changes. The catastrophic hurricane exposure, which is unlikely to change very often (unless the firm expands or contracts the size of its factories), must still be monitored and reported, but less frequently. An important by-product of the risk-monitoring process is the ability to change how risks are managed; without such visibility, a firm's risk strategies remain static. Monitoring thus feeds back into management.

We shall revisit aspects of this generic risk process (summarized in Figure 1.1) at greater length in the next few chapters, but for the moment let us expand on the third stage of the process below by considering specific management alternatives available to a company with financial or operating risks.

1.2.2 Risk management techniques

A company with any degree of risk exposure is wise to develop a philosophy that explicitly indicates its approach to risk and the resources it is willing to allocate (and potentially lose) in its endeavors. Best practice governance calls for a firm's board of directors to clearly express risk tolerance (or appetite) by relating exposures to overall corporate goals, stakeholder expectations, and financial/technical resources. Firms that are in business primarily to take risks, and have the financial resources to support potentially large losses, might choose to take a large amount of financial and operating risk. For instance, a bank might assume a considerable amount of credit and market risk as the core of its operation; given sufficient financial resources

and proper controls, it should be able to actively retain and manage such exposures. Those that are in business primarily to produce goods or services that are not based on active risk-taking, or those that lack sufficient financial resources to absorb large losses, are unlikely to favor significant risk exposure. For instance, a company that produces automobiles might be exposed to a series of input risks, such as steel and rubber; these form part of the core business and the board might wish to manage them by retaining them or hedging a portion of them. However, in order not to be distracted from its primary operations, it may not want to assume any risks related to non-core business activities, such as foreign exchange risk from sourcing raw materials or selling completed automobiles in other countries; these might not only be a distraction, but they might fall outside the firm's technical expertise. Assuming that the costs of doing so are consistent with its risk/return goals, the company may eliminate non-core risks.

It is common to consider three broad approaches to the management of risks, including loss control, loss financing, and risk reduction.

- **Loss control** Under this process (sometimes also referred to as loss prevention) a firm takes necessary *precautions* in order to reduce the threat of a particular risk. For instance, to diminish the likelihood of financial damage arising from a fire within a factory, a company might install a sprinkler system. Alternatively, a company dealing with hazardous material might reduce the chance of worker injury by introducing a comprehensive safety program. Loss control techniques vary by form of risk and potential threat, but typically involve an upfront investment and/or ongoing cost (e.g., paying for the sprinkler system, training personnel in safety procedures). As we shall see, the costs and benefits must be weighed in order to arrive at an appropriate decision.
- **Loss financing** This broad category of risk techniques, which involves the *transfer*, *retention*, or *hedging* of exposures, is primarily concerned with ensuring the availability of funds in the event of a loss. For instance, rather than installing a sprinkler system, a firm may choose to protect against potential fire damage by transferring risk through the purchase of an insurance policy that provides compensation if a fire occurs. Alternatively, the company exposed to \$/¥ foreign exchange risk might purchase a currency option as a hedge. Or, if a company feels that its risk exposures are particularly 'well-behaved' – reasonable in size and predictable with some degree of certainty – it may retain a portion. There are special instances where a company might choose to bundle together various techniques to produce a hybrid, or customized, solution. For instance, it might want to retain a portion of its \$/¥ risk and transfer the balance through a hedge, or it might wish to combine disparate risks – such as its property exposure from fire risk and its \$/¥ risk – into a single transfer mechanism. In fact, the hybrid management of risk is a cornerstone of the ART market, as well shall discover in later chapters. Regardless of the specific technique used, the relative costs of retention, transfer, hedging, or some hybrid must be weighed against possible benefits before a decision can be made.
- **Risk reduction** In some instances the risks may be too idiosyncratic or misaligned for a company to consider loss control or loss-financing methods. Accordingly, it might employ *risk reduction* techniques that involve partial or complete *withdrawal* from a business with particular characteristics or the *diversification* of exposures through a pooling or portfolio concept. Either can lead to a reduction in risk levels. Again, the risk reduction process has an associated cost and must therefore be considered in the cost/benefit framework before a decision is taken.

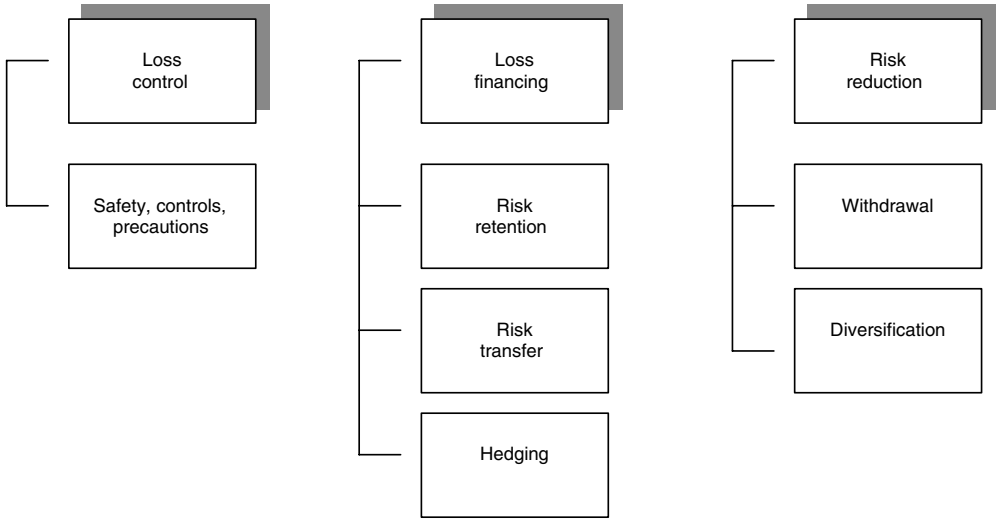

Figure 1.2 Risk management techniques

Risk exposures that are not eliminated must be managed through retention, transfer or financing (while loss control measures may be beneficial, they are generally applied to risks that are retained, e.g., loss control measures are more likely to be dependent on retention levels rather than vice versa). In fact, the general category of loss financing is a major focus of active risk management. Loss-financing techniques – including use of retained earnings, self-insurance, captives, contingent capital, and so on – can be managed from an internal or external perspective and may be funded or unfunded prior to a loss. We shall discuss a number of these techniques in subsequent chapters, as they form an essential part of the ART market. Figure 1.2 summarizes some common **risk management techniques**.

In practice, financial and non-financial corporations can turn to a range of instruments to execute active risk management strategies. Firms often use a combination of tools and may even bundle them together in order to produce a more efficient and cost-effective solution. For instance, an insurance company, which is in the business of underwriting risks, must manage its own risk profile actively and continuously, and may do so by:

- retaining some amount of risk, after having assessed the likelihood of loss and charged an appropriate premium (that covers expected losses and provides a fair return);
- identifying risks where it feels it must raise premiums in order to compensate for increased risks;
- ceasing to underwrite risks where it does not feel it is earning a proper return;
- creating additional reserves to cover unexpected losses;
- diversifying its portfolio further by expanding its underwriting efforts into new, uncorrelated, and profitable markets;
- purchasing reinsurance cover for portions of its portfolio from a reinsurer;
- issuing an insurance-linked security or structuring a contingent capital facility to provide additional funded or unfunded cover.

There are obviously many possibilities to consider that are applicable to both industrial and financial corporations, and most sectors enjoy access to multiple risk management solutions.

Each scheme has specific costs and benefits, but many can be applied in the structuring of an appropriate risk management program. In many cases it takes time to reshape the risk characteristics of a portfolio of businesses, and although some solutions can be enacted quickly, processes such as increasing premium rates, diversifying a portfolio or issuing an insurance-linked security might take several months (or longer). Therefore, companies must always be aware of the time dimension of the risk management process.

A convenient "rule of thumb" related to risk management techniques suggests that core risks – those that are central to a firm's daily business – should be retained, while non-core risks – those that are a byproduct of daily business – should be transferred or hedged. The premise is that a company has information and expertise regarding its core risks and, therefore, greater ability to manage its exposures intelligently (e.g., safely, efficiently, and cost-effectively). Exposure to risks where it lacks knowledge or competitive advantage can be more dangerous and costly. The generalization is interesting, but is complex and often nebulous. For instance, should an aircraft manufacturer view the price of steel, one of its key inputs, as a core or non-core risk? If it is a core risk should it actively retain and manage the exposure by dedicating resources and time to the effort? Should it transfer, hedge or eliminate a core risk if there is a remote possibility of an excessively large loss? If it is a non-core risk should the firm ignore the price of steel by simply locking in a price for future steel delivery, or should it be more dynamic about its hedging? Many other issues can obviously influence the decision, so the rule of thumb may be seen as somewhat simplistic.

In fact, while the core/non-core distinction may be applicable in some instances, it may not necessarily result in the best decision for every company in every scenario. The risk management decision process is complicated and must generally be considered through a rigorous analytical framework, such as a cost/benefit analysis. This can help a company to determine how it should manage its individual and aggregate risk exposures in order to maximize value (which, as we shall note in Chapter 2, is a general corporate goal). The cost/benefit tradeoff, characteristic of every risk-related decision a firm must make, is straightforward:

- Pay a cost and gain a benefit by eliminating or reducing NCF uncertainty.
- Pay nothing but accept the NCF uncertainty and remain exposed to potential cash flow volatility.

Since every risk has a theoretical price, it is possible to create a risk-free company by paying all the costs associated with eliminating every aspect of risk (e.g., through premiums, safety measures, diversification, withdrawal from businesses, and so on); the uncertainty associated with expected NCFs will then be eliminated. This, as we shall note later, is likely to be prohibitively expensive and impractical, and will almost certainly not result in a maximization of enterprise value. Accordingly, risk management solutions, consistent with the firm's appetite and philosophy, must focus on the tradeoffs between costs and benefits; only when this is thoroughly understood can a solution that leads toward enterprise value maximization be developed.

1.2.3 General risk management considerations

Risk management is concerned with the best and most efficient way of coping with financial and operating uncertainties. When crafting a risk strategy, firms often consider the process

in two different stages: pre-loss management and post-loss management. **Pre-loss management** prepares a firm for possible losses in a way that maximizes corporate value and covers legal and contractual obligations. **Post-loss management** ensures that a firm operates as a "going concern" with stable earnings and a minimal possibility of financial distress. The techniques we discuss in this book can be categorized as pre-loss or post-loss risk management tools.

The corporate governance process demands that a company, in fulfilling its responsibilities to shareholders, consider and define its tolerance for operating and financial risks. Directors must ensure that executives and independent control functions monitor, manage, and control exposures on an ongoing basis. In addition, shareholders must be made aware of the risks the company is retaining, eliminating, or transferring.

As noted earlier, a key element of the process is the firm's definition of a **risk philosophy**, a statement that reflects the firm's objectives related to the management of risk. Ideally, this should correlate with the specific type and amount of exposure the firm intends to take, retain, transfer, or reduce. For instance, in a pre-loss state a company might want to implement a risk management strategy that allows it to reduce the possibility of financial catastrophe, meet regulatory requirements, and operate more efficiently. In a post-loss state a company may want to ensure that its strategy allows it to operate as a going concern, and to continue expanding revenues and stabilize earnings.

We shall see throughout the course of the text that there are many ways of considering and managing risks, and the construction of a standard 'template' to fit every situation is simply not feasible. All companies are different. They engage in a wide range of businesses, have unique financial profiles and mandates, and are subject to unique internal and external pressures – meaning that there is no universal paradigm when it comes to creating a risk management program. We can therefore propose only general approaches to the management of risks for a generic 'risk averse' company. In particular, we note that it is often advantageous from a cost/benefit perspective for a firm exposed to low-frequency risks (i.e., those that are highly improbable) and low-severity risks (i.e., those with a small financial impact) to simply retain the exposures and fund them as losses occur (or fund them in advance through a self-insurance fund). Low-frequency but high-severity risks (i.e., infrequent, but with a large financial impact) are often good candidates for some type of loss financing (e.g., insurance, hedging). High-frequency but low-severity risks (i.e., those that are highly probable but not especially damaging) can often be accommodated via loss prevention and/or loss retention programs. High frequency and high severity risks (i.e., highly likely and highly damaging risks that can lead a company into financial distress) must typically be avoided. Although these are simple generalizations, they will help to focus our discussion at various points in the text. Table 1.1 and Figure 1.3 summarize these guidelines.

Table 1.1 Generalized risk management guidelines

Frequency	Severity	Guideline
Low	Low	Retention
Low	High	Loss financing (insure, hedge)
High	Low	Prevention, retention
High	High	Avoidance

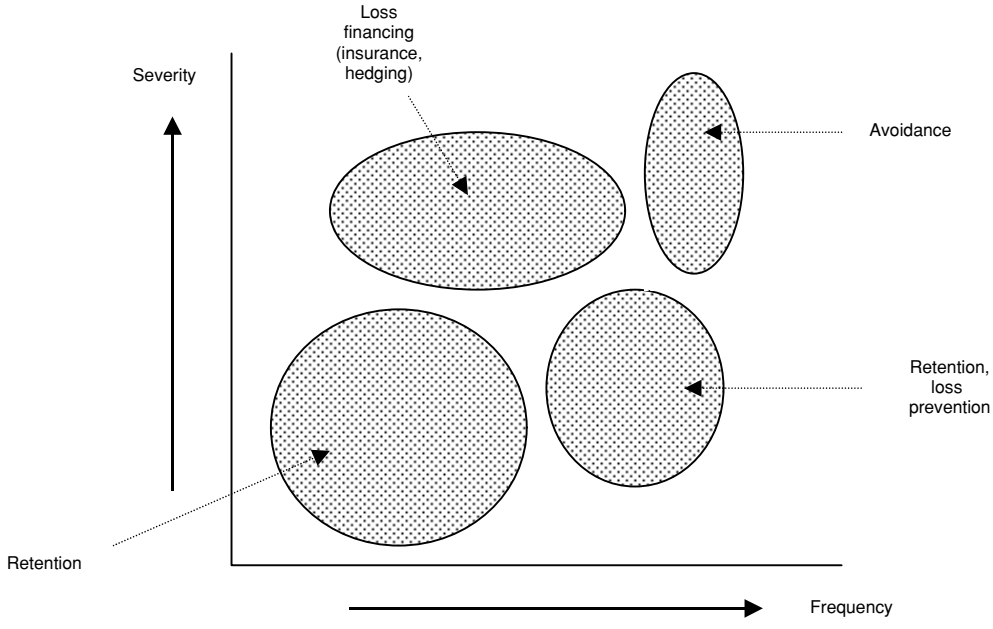

Figure 1.3 Generalized risk management guidelines

1.3 RISK CONCEPTS

To further frame aspects of our discussion we introduce several fundamental risk concepts in this section. This will assist in the chapters that follow when we consider the specific benefits provided by a variety of risk management instruments and techniques. Risk concepts can quickly become highly technical (with a great deal of intricate mathematics and statistics) but we have chosen to focus our discussion on basic ideas; readers interested in a detailed, technical treatment of these topics may wish to consult the references listed at the end of the book.

1.3.1 Expected value and variance

We begin with the concept of a **random variable**, which is simply a variable with an uncertain outcome. The variable can be discrete (appearing at specified time intervals) or continuous (appearing at any time), and may carry a defined value or any value at all; the result of a fair coin toss is thus a random variable with one of two possible values. By drawing many samples of random variables we can create a distribution that identifies all possible outcomes and their probability of occurrence. Distributions can take different shapes, but we shall concentrate primarily on the normal distribution, with its traditional bell shape. All the information regarding a random variable is summarized in the statistical distribution, which then becomes a useful tool when trying to estimate, *ex-ante*, the possibility of some event occurring. For instance, we can use statistical properties to obtain information about the likelihood that a particular event (e.g., a loss) will occur and the magnitude of the event that occurs.

Expected value (EV), the value that is obtained given a certain probability of occurrence, is a central element of statistics and of considerable use in risk evaluation. EV is determined

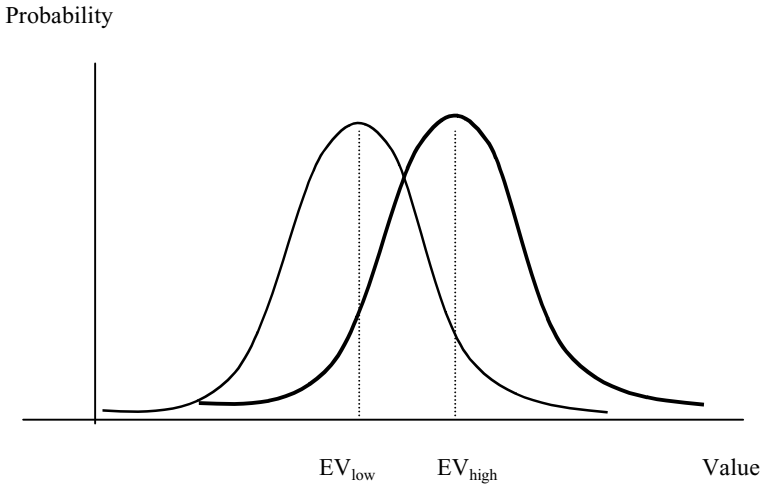

Figure 1.4 Distributions and expected values

by multiplying the probability of occurrence by the outcome of an event; in risk management terms this is often summarized as frequency (probability) times severity (outcome). More formally, we can say

$$\text{EV} = (\text{Probability} \times \text{Outcome}) + ((1 - \text{Probability}) \times \text{Outcome})$$

or

$$\text{EV} = \sum_{i=1}^{N} x_i \, p_i$$

where x_i is the outcome and p_i is the probability.

Thus, a payoff of \$80 occurring with 20% probability, and a payoff of \$100 occurring with 80% probability generates an EV of \$96. The EV of a probability distribution provides information of average outcomes. A distribution with a higher EV will have a higher outcome, on average, than one with a lower EV; this relationship, for a normal distribution, is depicted in Figure 1.4.

From a pure risk perspective we can create a probability distribution that focuses strictly on losses; the EV of the loss distribution is equivalent to the **expected loss**. Creating a loss distribution can be done through historical loss experience (this is possible for insurance companies and other financial intermediaries with a long history of risk management data) but still demands considerable geographic depth and breadth. Alternatively, certain simulation techniques or non-statistical estimates (e.g., those that might be found via technical or economic studies) can be used.

Next we introduce the **variance** (or standard deviation, which is equal to the square root of the variance); this is a measure of the magnitude by which an outcome differs from the EV and is given as:

$$\text{Var} = \text{Probability} \times (\text{Outcome} - \text{EV})^2$$

or

$$\text{Var} = \sum_{i=1}^{N} p_i (x_i - \mu)^2$$

where μ is the expected value and all other terms are as defined above.

Standard deviation is simply:

$$\text{SD} = \sqrt{\text{Var}}$$

When variance is low, the actual outcome is likely to be close to the EV, and when it is high it may be quite far away and difficult to predict. Not surprisingly, since variance is a measure of the difference between actual and expected outcomes, it serves as an important measure of risk – indeed, it reflects variability against expectations, which is the essence of risk. Standard deviation is useful when we are trying to consider the likelihood that an observation will lie within a particular range of values. Using the normal distribution, an observation falling within ± 1 standard deviation is expected to occur 68% of the time; ± 1.96 standard deviations includes 95% of observations, and so on. With this information we can construct a loss distribution to determine possible losses arising from risky activities, adjusted to a specified statistical confidence level (e.g., 90%, 95%, 99%). It is also possible to compute the **probability of ruin**, or the chance that the distribution of average losses will exceed a solvency benchmark value (e.g., some minimum surplus or tangible net worth amount); this, again, is an important measure in risk management.

Since representation of an entire population of observations is not realistic, we need to rely on smaller samples; accordingly, we use the sample mean (μ) and sample standard deviation (σ) as appropriate representations. Assuming that the correct sampling techniques are used, then the greater the sample size the narrower the range of error at particular statistical confidence intervals. Figure 1.5 summarizes the normal distribution, expected value, and standard deviation parameters.

Statistical loss forecasting, which is an important dimension of risk management, can be accomplished through probability analysis, regression analysis, loss distribution analysis, and other techniques. For instance, probability analysis focuses on the number of events that could give rise to risk exposure and considers the dependence/independence characteristics associated with each. We shall consider this at greater length in the risk-pooling example below. Regression analysis makes use of historical data to determine how a dependent variable is impacted by a series of independent variables (e.g., damage to a fleet of automobiles (dependent variable) based on the number of inches of snow or rain (independent variable)).

1.3.2 Risk aversion

Risk aversion is characteristic of a company that prefers less, rather than more, risk, and is willing to pay a price for protection (via reduction, transfer, hedging). The existence of risk aversion can be demonstrated by the demand for insurance and other risk mitigants: individuals and institutions are willing to pay for risk management because they are averse, to varying degrees, to the risk of loss. If risk aversion did not exist, there would be no willingness to pay for mitigation; individuals and firms would simply bear the risk of loss to which they were exposed.

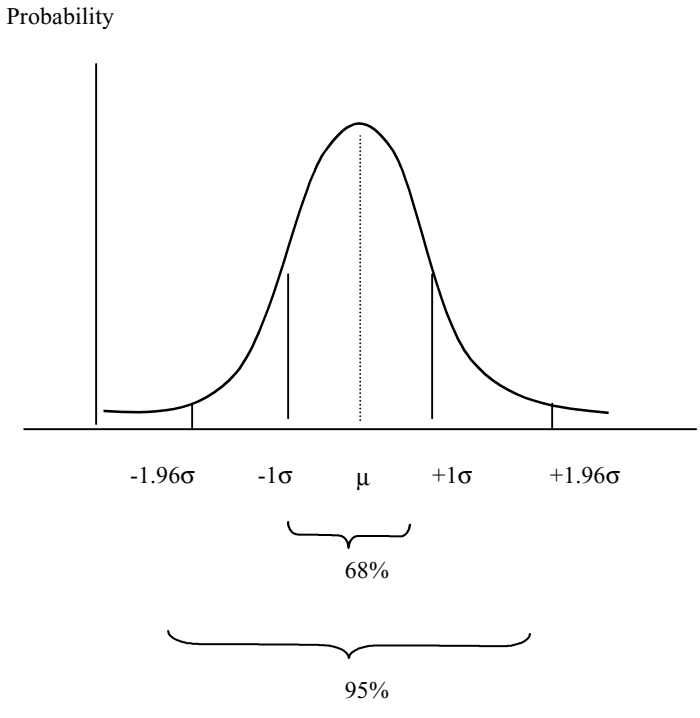

Figure 1.5 Normal distribution

We know that the greater the variability in potential outcomes, the greater the risk; this stems primarily from lack of *ex-ante* knowledge about the outcome that will occur. In the absence of risk, decision-making is simple: outcomes that generate the highest value are preferred and the rational firm will select the outcome that yields the greatest EV. Relating EV to the economic concept of **expected utility**, or the weighted average utility value (e.g., satisfaction from income or wealth) derived from some activity, we can consider the law of diminishing marginal utility, which indicates that the utility derived from an incremental (or marginal) unit of wealth begins to diminish at some point. The risk averse firm faces a concave utility function, such as depicted in Figure 1.6, and will attempt to protect against risk of loss if the **risk premium**, or protection payment, it must pay is less than or equal to the EV of the loss. The expected utility of not protecting appears as a point below the utility function; if the risk premium is no greater than the EV of the loss, then acquiring protection will move the expected utility point of the risk averse firm up to the frontier of the curve. As a result of the concave utility function, parties that are risk averse demonstrate a willingness to pay to avoid risk that would jeopardize wealth. They may choose to do so through any of the risk management techniques summarized above, including loss control, loss financing, and risk reduction. Knowing this, it should be clear that the risk-seeking firm faces a convex utility function, as the marginal utility of wealth increases as wealth increases.

While utility functions can be interesting to consider in a theoretical sense, they are seldom used in practical corporate risk management applications as constructing a meaningful utility function is challenging, if not impossible. However, the notion of the risk averse firm is fundamental to the working of the risk management markets.

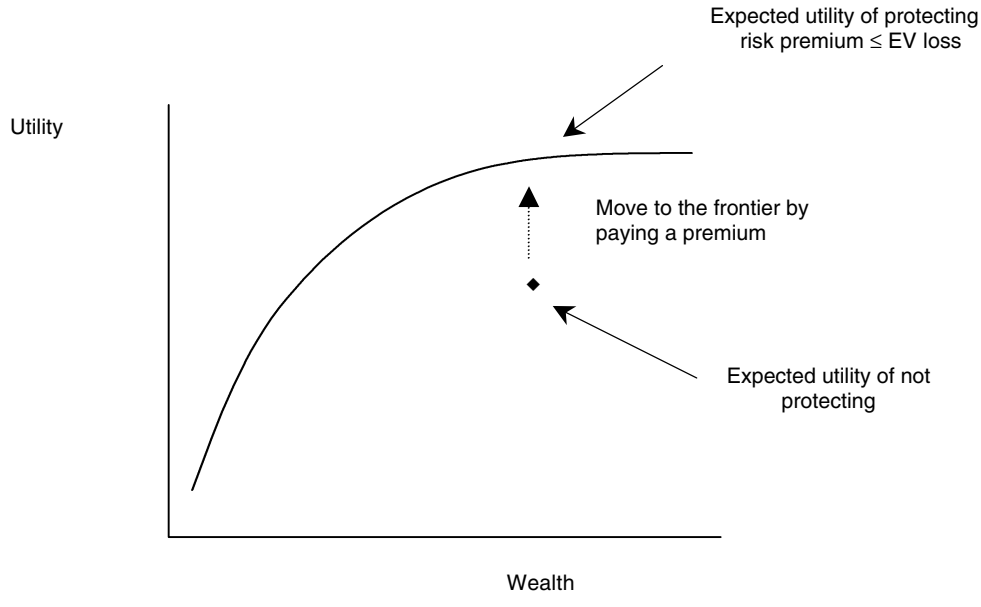

Figure 1.6 Risk premium and utility of a risk averse firm

1.3.3 Risk transfer and the insurance mechanism

The insurance market is premised on two fundamental characteristics: the transfer of exposure from a single party to a broad group, and the sharing of losses by all those in the group. **Risk transfer**, as the name suggests, occurs when one party pays a second party a small, certain cost (e.g., a risk premium) in exchange for coverage of uncertain losses; this is equal to a shifting of exposures. The risk averse firm, in creating its risk philosophy, may decide to shed an exposure by transferring it through one of several different mechanisms, including insurance/reinsurance, derivatives or hybrid structures. The amount of risk that a firm transfers is a function of overall tolerance (i.e., its level of risk aversion), the specific benefits it hopes to derive, and the total cost; this is often determined in a cost/benefit analysis framework.

An insurer can generally predict, within fairly tight ranges, the amount of losses that will occur for a given type of risk exposure. A large sample improves the estimate of the underlying probability of occurrence. Thus, when an insurer has a very large portfolio of relatively homogeneous policies, its ability to estimate losses improves. The process works on the basis of two statistical principles: the **Law of Large Numbers**, which indicates that as the number of participants (N) gets very large, the average outcome approaches the EV; and the **Central Limit Theorem**, which indicates that the distribution of the average outcome approaches the normal distribution as N gets very large.[3]

An **insurance contract** is an agreement between two parties (the insurer, as protection provider, and the **cedant** (also known as insured, beneficiary), as protection purchaser) that exchanges an *ex-ante* premium for an *ex-post* claim, with no ability to readjust the claim amount once it has been agreed. Insurance contracts are governed by the principle of **indemnity**,

[3] With mean μ and standard deviation $\sigma/\sqrt{N}$.

which indicates that the cedant cannot profit from its insurance activities; that is, insurance exists to cover a loss, not to generate a speculative profit. Coverage can be created through an **indemnity contract** (covering actual losses sustained) or a **valued contract** (covering a specific amount agreed upfront). A contract covering actual fire damage is an example of an indemnity contract, while a life insurance policy paying out a stated amount on death of the cedant is a valued contract. In order for a contract to qualify as insurance, the cedant must generally demonstrate an **insurable interest** – that is, it must prove that it has suffered an economic loss once the defined event occurs. Insurable interest exists to reduce or prevent gambling and moral hazard (as discussed below). An insurer, as cedant, may seek protection through a **reinsurance contract**; likewise, a reinsurer can obtain protection from another reinsurer through a **retrocession contract**.

A company may opt for **full insurance** (complete coverage of a risk exposure in exchange for a higher risk premium), or **partial insurance** (fractional coverage of risk for a lower risk premium). A cedant can create partial insurance by including a **deductible** (a 'first loss' amount paid by the cedant before the insurer makes a payment), a **coinsurance** feature (a 'shared loss' component between cedant and insurer), and/or a **policy cap** (a maximum amount payable by the insurer). We shall consider these, and associated technical details, in Chapter 4.

If it is economically sensible for the firm to pay the larger risk premium to secure full insurance (and consistent with its risk philosophy), it will do so. Alternatively, it may select from one of the partial insurance options. When a firm can clearly identify an optimal EV loss scenario that is preferable, the choice of protection becomes relatively straightforward. However, it is possible to create a range of full and partial insurance options with EV loss rankings; in such cases a firm needs to examine its utility function to determine whether one option dominates. In practice, since it is difficult for a company facing a complex set of businesses with varying priorities and goals to know the slope of its utility function, it must turn to alternative techniques (e.g., a cost/benefit review, a mean-variance analysis that takes specific account of variance/standard deviation and does not require *ex-ante* identification of a utility function, and similar "practical" measures).

1.3.4 Diversification and risk pooling

Diversification – a spreading or diffusion of risk exposures – is a common technique of risk management that seeks to lower risk by combining exposures that are not related (correlated) to one another. Much of this work has its foundation in Markowtiz' 1952 work related to capital markets portfolio theory, which demonstrates how diversification permits the risk averse investor to create portfolios that optimize various levels of risk and return. The intent is to create a portfolio on the **efficient frontier**, or the boundary that provides the maximum possible return for a given level of risk. Any portfolio that is below the efficient frontier fails to maximize value for a given level of risk, and can be enhanced through diversification (note also that superior portfolios of risk/return, along the "capital markets line" can be obtained by borrowing and lending at the risk-free rate).[4]

[4] From a pure investment perspective, an investor must look at the individual μ and σ of each security, as well as the correlation between the two, to determine how to construct an optional portfolio that minimizes risk and maximizes return. When all possible combinations of weighted portfolios have been found, the efficient frontier can be created. Portfolios on the efficient frontier represent the best possible tradeoffs between risk and return; those below the frontier are suboptimal (e.g., too much risk for a given return). Borrowing or lending an unlimited amount at some risk-free rate $r(\mathrm{f})$ means the investor can invest in both risky and non-risky securities and create leveraged portfolios that are better than those on the efficient frontier, regardless of the level of risk aversion; the "capital markets line" depicts these. Thus, investors that are risk averse can do better holding part of their capital in X and lending at $r(\mathrm{f})$.

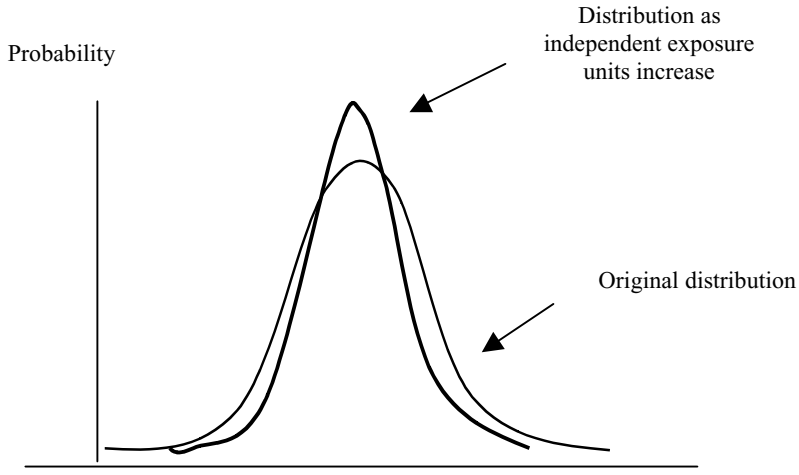

Figure 1.7 Distribution changes with independent exposure units

Risk pooling – a practical implementation of diversification and a fundamental mechanism of the risk management markets – is based on the idea that independent risks can be combined to reduce the overall level of risk. In addition to the Law of Large Numbers and CLT cited above, pooling relies on correlation to measure how random variables – i.e., individual risk exposure units, such as insurance policies – relate to one another. Correlation between two random variables, formally defined as the covariance of the two variables divided by the standard deviation of each one,[5] is measured on a scale of $+1$ to -1, where $+1$ implies perfect positive correlation and -1 perfect negative correlation; a correlation of 0 implies no relationship, indicating that the variables are independent. Thus, if two random variables have a correlation of $+0.7$, a movement of $+1$ in one leads to a movement of $+0.7$ in the other. Risk pooling reduces risks if expected losses are uncorrelated; when this occurs there is no change in the expected loss (or cost) but there is a reduction in the standard deviation. Consider the following simple example:

An automobile driver (A) has a 20% probability of being in an accident that will cost \$2500. By the equations introduced earlier, the EV is \$500 (e.g., (80% $\times$ \$0) + (20% $\times$ \$2500)) and the standard deviation is \$1000 (e.g., 80% $\times$ $(0 - 500)^2 +$ 20% $\times$ $(2500 - 500)^2$). Assume

Those who are more aggressive can borrow at $r(\mathrm{f})$ and invest in X. Accordingly, all investors choose the same portfolio, but vary the level of financing.

Let us assume that an insurance company has individual units of risk exposure (e.g., individual policies) that are independently exposed to the risk of loss; thus, if a loss occurs on one policy, it need not necessarily occur on others. Each unit of risk exposure has some probability of loss, and the sum of all units represents the insurance company's total liability. The statistical distribution of the entire group of independent risk units depends on the distribution of each individual unit (which might take any specific form); however, if they are truly independent, then the distribution of the average loss (e.g., all units of exposure) approaches the normal distribution. This means that we can draw some conclusions about the expected loss and variance of loss. In particular, as the number of units increases to some large number N, the actual loss experience approaches the expected loss experience, and the variance around the expected loss declines, as illustrated in Figure 1.7. This means that if an insurer can diversify its risks sufficiently (i.e., if it can create enough independent risk units), it can reduce the riskiness of its operations. In practice, the degree of independence is measured through correlation and is implemented through pooling techniques. Pooling is applicable to a broad range of risk classes; while it is commonly associated with risks arising from automobile accidents, worker safety, or health claims, it is equally applicable to financial risks, such as credit risks generated by corporate loans (indeed, insurers have become key players in the credit risk transfer market through their application of these techniques). The properties of a portfolio of risk exposure units are different from the sum of the individual units, so a focus on portfolio characteristics is important. If a firm has only a small number of units the portfolio risk profile will not change markedly, and the number of risk units is therefore a key driver in diversification. However, some benefit will still accrue if N is not particularly large, as long as the units are not perfectly correlated.

[5] More formally, $\rho(x, y) = \mathrm{Cov}(x, y)/(\sigma(x)\sigma(y))$, where $\mathrm{Cov}(x, y)$ is $\Sigma p((x - \mu(x)) \times (y - \mu(y)))$.

Table 1.2 Accident scenarios

Accident claim	Cost ($)	Cost per driver ($)	Probability (%)
0	0	0	$80 \times 80 = 64$
1 (A)	2500	1250	$20 \times 80 = 16$
1 (B)	2500	1250	$20 \times 80 = 16$
2 (A and B)	5000	2500	$20 \times 20 = 4$

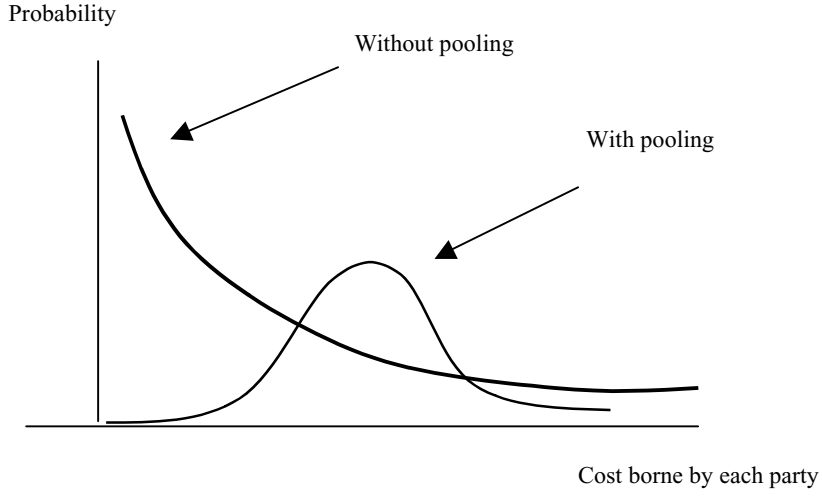

Figure 1.8 Pooling and costs

that another driver (B) faces the same accident parameters, and that driving events/behavior are uncorrelated (i.e., an accident by A will not lead to an accident by B, and vice versa). Under a pooling concept, both drivers agree to share the costs of an accident equally. Thus, if A has an accident he will only pay $1250 (B will pay the balance), and vice versa. We can now summarize various accident scenarios and costs in Table 1.2.

Through pooling, the probability distribution of costs for each participant has changed and the standard deviation, as a proxy of risk, has declined. For instance, a loss of $2500 now occurs 4% instead of 20% of the time, since two accidents, rather than just one, must happen. It is easy to extend the logic and demonstrate that the more participants in the pool, the lower the risk – as long as the exposures of the participants are not correlated. In addition, the probability of extreme outcomes declines. Risk pooling is not a risk transfer mechanism, but a risk reduction method, as long as the events are uncorrelated. If exposures are positively correlated to some degree, risk reduction is still possible, although it will not be as great (i.e., diversification helps but the beneficial effects are limited); when they are strongly positive, little (or no) benefit can be obtained. When exposures are negatively correlated they will not reduce risk to the same degree as independent exposures, but they can be used as 'counter cyclic covers' (and thus have favorable risk reduction characteristics). Summarizing, then, we note that when losses are uncorrelated, the risk in the pool (as measured by standard deviation) approaches zero as the number of pool participants increases; when losses are perfectly correlated risk remains unchanged. Figure 1.8 illustrates the effects of costs borne by each party with and without pooling.

It is worth noting that while risk transfer and pooling are often considered jointly when discussing insurance techniques, they are not synonymous or, indeed, mutually dependent. That is, in transferring risk, pooling can occur, but may not be necessary. For instance, there are times when an insurance company will accept a risk that it does not pool with others. Risk transfer, in contrast, must occur; that is the essence of the insurance mechanism.

1.3.5 Hedging

Insurance is generally associated with the transfer of an insurable risk and can result in a reduction of exposure. **Hedging**, in contrast, is generally associated with risks that are uninsurable through a standard contractual insurance framework, and typically result in transfer rather than reduction. Through hedging a firm transfers named risks to another party (via standard agreements rather than the more complex contracts that characterize insurance dealings). **Derivatives**, or financial transactions that derive their value from a market reference, are commonly used to hedge financial risks. They may be traded on a standardized basis through an exchange (as a listed contract) or in customized form through the over-the-counter (OTC) market. Unlike insurance contracts, derivatives represent an optionable, rather than an insurable, interest, meaning that a party to a contract does not need to be exposed to risk of loss. This suggests that derivatives can generate profits, and can be used to speculate rather than hedge. Derivatives are available in the form of:

- **Futures** Standardized exchange contracts that enable participants to buy or sell an underlying asset at a predetermined forward price.
- **Forwards** Customized off-exchange contracts that permit participants to buy or sell an underlying asset at a predetermined forward price.
- **Swaps** Customized off-exchange contracts that enable participants to exchange periodic flows based on an underlying reference.
- **Options** Standardized exchange or customized off-exchange contracts that grant the buyer the right, but not the obligation, to buy or sell an underlying asset at a predetermined strike price.

Insurance and derivatives have different features that can make one or the other more suitable in a given situation. For instance:

- Derivative contracts are linked to specific market references (or indexes) and are not limited by a cap or subject to the indemnity principle. Since derivatives are generally related to an index rather than a specific loss exposure, they are subject to **basis risk**, or the risk of loss arising from an imperfect match between the loss-making exposure and the compensatory hedge payment (as we discuss at greater length below). Derivatives are typically valued (e.g., marked-to-market) on a periodic basis and can often be traded/transferred between counterparties; in some instances credit exposures arising between two derivative parties are secured by collateral.
- Insurance contracts are based on specific losses or agreed amounts and are generally capped at an upper limit. Cedants must disclose all relevant information in "utmost good faith" through insurance documents and prove an insurable interest in order to make the contract valid and enforceable. Since most insurance contracts are related to specific risks, they

feature no basis risk (but there are some exceptions, as we shall note later). Insurance is not traded or marked-to-market, and credit exposures (i.e., those where the cedant is exposed to the credit of the insurer) are not generally secured.

Given these differences, derivatives are often more suitable when information about risk is well known, or where a company's exposure can be well correlated with a reference index (i.e., basis risk is not a concern). Insurance might be more suitable when the insured has private information about a particular risk exposure and the loss cannot easily be correlated to an external index. Ultimately, however, the relative costs and benefits (e.g., fees, premiums, bid-offer spreads, tax benefits, post-loss financial benefits) are likely to be the most decisive factors.

1.3.6 Moral hazard, adverse selection and basis risk

We now consider several additional concepts that are prevalent in the risk management markets, including moral hazard, adverse selection and basis risk. In its simplest form **moral hazard** can be regarded as a change in behavior arising from the presence of insurance or other forms of risk protection.[6] Theory and practice suggest that the availability of a compensatory payment in the event of loss removes a firm's incentive to behave prudently. For instance, a firm might be exposed to the risk of fire in its operations; if fire strikes its factory and destroys its equipment, it will be unable to produce its goods and thus suffer a loss of sales revenue. Accordingly, it may purchase a policy that covers losses attributable to fire damage. Once in possession of the policy, however, it may behave more carelessly – perhaps leaving flammable material on the factory floor, not upgrading its fire extinguishers and sprinklers when they become outdated, and so on. It will do so because it knows that it is protected: the insurance policy will cover any fire-related losses, so it no longer needs to be too careful. The same behavior can be found in many other types of risk exposure/risk protection schemes and is often a key concern of intermediaries providing alternative forms of risk protection. To combat moral hazard, insurance firms and other financial institutions providing protection modify the terms of their coverage to ensure that the firm bears some of the economic loss. This can occur through use of deductibles (e.g., the cedant bears the first losses, either per event or in aggregate), co-payments/coinsurance (e.g., the cedant and insurer share losses on some pre-arranged basis) or policy caps (e.g., the insurer limits the amount of cover granted to the cedant). Moral hazard may also be explicitly or implicitly priced into the premium, becoming a cost of risk borne by the company and its shareholders.

Adverse selection is defined as the mispricing of risk as a result of information asymmetries, and occurs when a protection provider cannot clearly distinguish between different classes of risk. The end result is that the protection provider supplies too much or too little risk cover at a given price, leading ultimately to an excess of losses or dearth of business. For instance, if an insurer is unable to distinguish between the risk characteristics of two groups of cedants – a high-risk group and a low-risk group – one of two scenarios will emerge: it will price all risk at the low-loss level, meaning that the high-loss group will purchase large quantities of

[6] Moral hazard is generally associated with *ex-ante* behavior, and might be considered *ex-ante* moral hazard (e.g., failing to take actions to prevent losses knowing that insurance coverage exists). There is also a form of *ex-post* moral hazard that can arise from the presence of reinsurance; under this concept, an insurer might relax its loss settlement/claims adjustment procedures in the aftermath of a loss, knowing that it has reinsurance coverage. This can lead to an accumulation of claims/fraud.

cover and generate excessive losses for the insurer; or, it will price all risk at the high-loss level and write no cover for the low-loss group, thus losing business. Risks, in either case, are said to be 'adversely selected', which will have a detrimental effect on the insurer. To protect against adverse selection the insurer must thoroughly understand the nature of its portfolio; this typically means devoting proper resources to identifying, classifying, and tracking the loss experience of each of the parties it is protecting, enabling it to properly stratify and, then price, the protection it is offering.

As indicated above, basis risk is the risk that arises between an exposure and a risk transfer/hedge mechanism that is imperfectly correlated with the exposure. Basis risk arises in derivative and insurance contracts when a company attempts to protect a particular exposure with a proxy that is not precisely matched with the potential loss. An indemnity-based insurance contract, which provides a payment that matches precisely the losses sustained by the insured, features no basis risk. A derivative contract that provides a payment to a hedger based on a proxy has basis risk; the degree of risk depends on the correlation between the exposure and the hedge, and how that correlation performs over time. Of course, not all derivative contracts carry basis risk (e.g., it is possible for a corporate hedger to find a market reference that covers an exposure precisely) and not all insurance contracts are free from basis risk (e.g., a reinsurance contract that provides loss coverage based on an index or parametric trigger, rather than a specific indemnity, has basis risk). All else being equal, a contract that has basis risk is cheaper than one that provides a perfect match; this is logical as the hedger is bearing an incremental amount of risk and the protection provider is not including any premium for moral hazard.

1.3.7 Non-insurance transfers

In addition to some of the risk management mechanism we have summarized above, there are other ways of transferring pure risks, including hold-harmless agreements, indemnity agreements, and leases. In fact, these can allow coverage of risks that might not normally be insurable through standard mechanisms, and they may be a cost-effective way of protecting business. However, coverage can be ambiguous and the level of credit risk the company assumes necessarily rises. While these are certainly valid risk transfer mechanisms, we shall not discuss them in detail in this book.

1.4 OUTLINE OF THE BOOK

With this brief overview of basic risk management issues, we are now prepared to consider how and why the ART market exists and the specific products, vehicles, and solutions that are available to those who actively manage their risks. The balance of the book is structured as follows:

- In Chapters 2 and 3, we discuss important theoretical market drivers that promote growth and innovation in the ART market, the scope, development, and evolution of the marketplace, the nature of convergence, and the role that key participants play in promoting activity.
- In Part II (Chapters 4, 5, and 6) we turn our attention to specific insurance/reinsurance-based ART products and vehicles, including risk transfer and risk financing contracts, captives, and multi-risk (multiple peril and trigger) products.

- In Part III (Chapters 7, 8, and 9) we consider the capital markets dimension of the marketplace, with a particular focus on insurance-based securitization, contingent capital structures and derivatives.
- In Part IV (Chapters 10 and 11) we focus on the future of ART by reviewing the nascent, but increasingly important, field of enterprise risk management (based on integrated risk products and corporate solutions) and conclude with thoughts on future growth prospects.