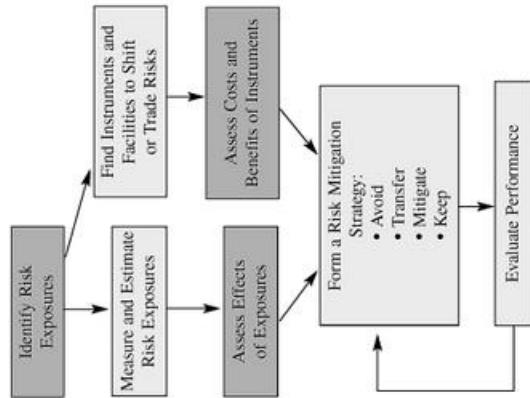


CHAPTER 1

Risk Management—A Helicopter View¹

FIGURE 1-1
The Risk Management Process



reward lies at the heart of the management process of all enduringly successful corporations.

Yet the rise of financial risk management as a formal discipline has been a bumpy affair, especially over the last 10 years. On the one hand, we've seen an extraordinary growth in new types of institutions that earn their keep by taking and managing risk (e.g., hedge funds), as well as some extraordinary successes in risk management mechanisms: the lack of financial institution bankruptcies during the violent downturn in credit quality in 2001–2002 is often claimed to be the result of better credit-risk management processes at banks.

Risk management is also now widely acknowledged as the most creative force in the world's financial markets. A striking recent example is

The future cannot be predicted. It is uncertain, and no one has ever been successful in forecasting the stock market, interest rates, or exchange rates consistently—or credit, operational, and systemic events with major financial implications. Yet, the financial risk that arises from uncertainty can be managed. Indeed, much of what distinguishes modern economies from those of the past is the new ability to identify risk, to measure it, to appreciate its consequences, and then to take action accordingly, such as transferring or mitigating the risk.

This simple sequence of activities, shown in more detail in Figure 1-1, is often used to define risk management as a formal discipline. But it's a sequence that rarely runs smoothly in practice: sometimes simply identifying a risk is the critical problem, while at other times arranging an efficient economic transfer of the risk is the skill that makes one risk manager stand out from another. (In Chapter 2 we discuss the risk management process from the perspective of a corporation.)

To the contrary, Figure 1-1 might suggest that risk management is a continual process of corporate risk reduction. But we mustn't think of the modern attempt to master risk in defensive terms alone. Risk management is really about how firms actively select the type and level of risk that it is appropriate for them to assume. Most business decisions are about sacrificing current resources for future uncertain returns.

In this sense, *risk management* and *risk taking* aren't opposites, but two sides of the same coin. Together they drive all our modern economies: the capacity to make forward-looking choices about risk in relation to

Part of the reason for risk management's mixed record here lies with the double-edged nature of risk management technologies. Every financial instrument that allows a company to transfer risk also allows other corporations to assume that risk as a counterparty in the same market—wisely or not. Most importantly, every risk management mechanism that allows us to change the shape of cash flows, such as deferring a negative outcome into the future, may work to the short-term benefit of one group of stakeholders in a firm (e.g., managers) at the same time that it is destroying long-term value for another group (e.g., shareholders or pensioners). In a world that is increasingly driven by risk management concepts and technologies, we need to look more carefully at the increasingly fluid and complex nature of risk itself, and at how to determine whether any change in a corporation's risk profile serves the interests of stakeholders. We need to make sure we are at least as literate in the language of risk as we are in the language of reward.

The nature of risk forms the topic of our next section, and it will lead us to the reason we've tried to make this book accessible to everyone, from shareholders, board members, and top executives to line managers, legal and back-office staff, and administrative assistants. We've removed from this book many of the complexities of mathematics that act as a barrier to understanding the essential principles of risk management, in the belief that, just as war is too important to be left to the generals, risk management has become too important to be left to the "rocket scientists" of the world of financial derivatives.

WHAT IS RISK?

We're all faced with risk in our everyday lives. And although risk is an abstract term, our natural human understanding of the trade-offs between risk and reward is pretty sophisticated. For example, in our personal lives, we intuitively understand the difference between a cost that's already been budgeted for (in risk parlance, a predictable or expected loss) and an unbudgeted cost (at its worst, a catastrophic loss of a magnitude well beyond losses seen in the course of normal daily life).

In particular, we understand that risk is not synonymous with the size of a cost or of a loss. After all, some of the costs we expect in daily life are very large indeed if we think in terms of our annual budgets: food, fixed mortgage payments, college fees, and so on. These costs are big, but they are not a threat to our ambitions because they are reasonably predictable and are already allowed for in our plans.

the development of a huge market for credit derivatives, which allows institutions to obtain insurance to protect themselves against credit default (or, alternatively, to get paid for assuming credit risk as an investment). Credit derivatives can be used to redistribute part or all of an institution's credit-risk exposures to banks, hedge funds, or other institutional investors, and they are a specific example of a broader, beneficial trend in financial markets summed up by Alan Greenspan, chairman of the U.S. Federal Reserve Board:

The development of our paradigms for containing risk has emphasized dispersion of risk to those willing, and presumably able, to bear it. If risk is properly dispersed, shocks to the overall economic system will be better absorbed and less likely to create cascading failures that could threaten financial stability.²

On the other hand, the last 10 years have seen some extraordinary and embarrassing failures of risk management in its broadest definition. These range from the near failure of the giant hedge fund Long-Term Capital Management (LTCM) in 1998 to the string of financial scandals associated with the millennial boom in the equity and technology markets (from Enron, WorldCom, Global Crossing, and Qwest in the United States to Parmalat in Europe).

Unfortunately, risk management has not consistently been able to prevent market disruptions or to prevent business accounting scandals resulting from breakdowns in corporate governance. In the case of the former problem, there are serious concerns that derivative markets make it easier to take on large amounts of risk, and that the "herd behavior" of risk managers after a crisis gets underway (e.g., selling risky asset classes when risk measures reach a certain level) actually increases market volatility.

Sophisticated financial engineering, supplied by the banking, securities, and insurance industries, also played a role in covering up the true economic condition of poorly run companies during the equity markets' millennial boom and bust. Alongside rather simpler accounting mistakes and ruses, this type of financial engineering was one reason that some of these companies violently imploded after years of false success (rather than simply fading away or being taken over at an earlier point).

2. Remarks by Chairman Alan Greenspan before the Council on Foreign Relations, Washington, D.C., Nov. 19, 2002.

BOX 1-1**RISK FACTORS AND THE MODELING OF RISK**

In order to measure risk, the risk analyst first seeks to identify the key factors that seem likely to cause volatility in the returns from the position or portfolio under consideration. For example, in the case of an equity investment, the risk factor will be the volatility of the stock price (categorized in the appendix to this chapter as a market risk), which can be estimated in various ways.

In this case, we identified a single risk factor. But the number of risk factors that are considered in a risk analysis—and included in any risk modeling—varies considerably depending on both the problem and the sophistication of the approach. For example, in the recent past, bank risk analysts might have analyzed the risk of an interest-rate position in terms of the effect of a single risk factor—e.g., the yield to maturity of government bonds, assuming that the yields for all maturities are perfectly correlated. But this one-factor model approach ignored the risk that the dynamic of the term structure of interest rates is driven by more factors, e.g., the forward rates. Nowadays, leading banks analyze their interest-rate exposures using at least two or three factors, as we describe in Chapter 6.

Further, the risk manager must also measure the influence of the risk factors on each other, the statistical measure of which is the “covariance.” Disentangling the effects of multiple risk factors and quantifying the influence of each is a fairly complicated undertaking, especially when covariance alters over time (i.e., is *stochastic*; in the modeler’s terminology). There is often a distinct difference in the behavior and relationship of risk factors during normal business conditions and during stressful conditions such as financial crises.

Under ordinary market conditions, the behavior of risk factors is relatively less difficult to predict because it does not change significantly in the short and medium term: future behavior can be extrapolated, to some extent, from past performance. However, during stressful conditions, the behavior of risk factors becomes far more unpredictable, and past behavior may offer little help in predicting future behavior. It’s at this point that statistically measurable risk threatens to turn into the kind of unmeasurable uncertainty that we discuss in Box 1-2.

The real *risk* is that these costs will suddenly rise in an entirely unexpected way, or that some other cost will appear from nowhere and steal the money we’ve set aside for our expected outlays. The risk lies in how *variable* our costs and revenues really are. In particular, we care about how likely it is that we’ll encounter a loss big enough to upset our plans (one that we have not defused through some piece of personal risk management such as taking out a fixed-rate mortgage, setting aside savings for a rainy day, and so on).

This day-to-day analogy makes it easier to understand the difference between the risk management concepts of *expected loss* (or expected costs) and *unexpected loss* (or unexpected cost). Understanding this difference is a task that has managed to confuse even risk-literate banking regulators over the last few years, but it’s the key to understanding modern risk management concepts such as economic capital attribution and risk-adjusted pricing. (However, this is not the only way to define risk, as we’ll see in Chapter 5, which discusses various academic theories that shed more light on the definition and measurement of risk.)

The main difference between our intuitive conception of risk and a more formal treatment of it is the use of statistics to define the extent and potential cost of any exposure. To develop a number for unexpected loss, a bank risk manager first identifies the risk factors that seem to drive volatility in any outcome (Box 1-1) and then uses statistical analysis to calculate the probabilities of various outcomes for the position or portfolio under consideration. This probability distribution can be used in various ways; for example, the risk manager might pinpoint the area of the distribution (i.e., the extent of loss) that the institution would find worrying, given the probability of this loss occurring (e.g., is it a 1 in 10 or a 1 in 10,000 chance?).

The distribution can also be related to the institution’s stated “risk appetite” for its various activities. For example, as we discuss in Chapter 4, the senior risk committee at the bank might have set boundaries on the institution’s future risk that it is willing to take by specifying the maximum loss it is willing to tolerate at a given level of confidence, such as, “We are willing to countenance a 1 percent chance of a \$50 million loss from our trading desks on any given day.”

The formality of this language and the use of statistical concepts can make risk management sound pretty technical. But the risk manager is simply doing what we all do when we ask ourselves in our personal lives, “How bad, within reason, might this problem get?”

What does our distinction between expected loss and unexpected loss