1

What is risk management?

The biggest fraud of all time

A number of banks have succeeded in losing huge sums of money in their trading operations, but Société Générale ('SocGen') has the distinction of losing the largest amount of money as the result of a fraud. This took place in 2007, but was uncovered in January 2008. SocGen is one of the largest banks in Europe and the size of the fraud itself is staggering; SocGen estimated that it lost 4.9 billion Euros as a result of unwinding the positions that had been entered into. With a smaller firm this could well have caused the bank's collapse, as happened to Barings in 1995, but SocGen is large enough to weather the storm. The employee responsible was Jérôme Kerviel, who did not profit personally (or at least only through his bonus payments being increased). In effect, he was taking enormous unauthorized gambles with his employer's money. For a while these gambles came off, but in the end they went very badly wrong.

In America the news broke on January 24, 2008, when the *New York Times* reported as follows:

> 'Société Générale, one of the largest banks in Europe, was thrown into turmoil Thursday after it revealed that a rogue employee had executed a series of "elaborate, fictitious transactions" that cost the company more than $7 billion US, the biggest loss ever recorded in the financial industry by a single trader.
>
> Before the discovery of the fraud, Société Générale had been preparing to announce pretax profit for 2007 of €5.5 billion, a figure that Bouton (the Société Générale chairman) said would have shown the company's "capacity to absorb a very grave crisis." Instead, Bouton – who is forgoing his salary through June as a sign of taking responsibility – said the "unprecedented" magnitude of the loss had prompted it to seek

Business Risk Management: Models and Analysis, First Edition. Edward J. Anderson.

Companion website: www.wiley.com/go/business_risk_management

> about €5.5 billion in new capital to shore up its finances, a move that secures the bank against collapse.
>
> Société Générale said it had no indication whatsoever that the trader – who joined the company in 2000 and worked for several years in the bank's French risk-management office before being moved to its Delta One trading desk in Paris – "had taken massive fraudulent directional positions in 2007 and 2008 far beyond his limited authority." The bank added: "Aided by his in-depth knowledge of the control procedures resulting from his former employment in the middle-office, he managed to conceal these positions through a scheme of elaborate fictitious transactions."
>
> When the fraud was unveiled, Bouton said, it was "imperative that the enormous position that he had built, and hidden, be closed out as rapidly as possible." The timing could hardly have been worse. Société Générale was forced to begin unwinding the trades on Monday "under conditions of extreme market volatility," Bouton said, as global stock markets plunged amid mounting fears of an economic recession in the United States.'

A story like this inevitably prompts the question: How could this have happened? Later in this chapter we will give more details about what went wrong. SocGen was a victim of an enormous fraud but the defense lawyers at Kerviel's trial argued that the company itself was primarily responsible. Whatever degree of blame is assigned to SocGen, it clearly paid a heavy price. It is easy to be wise after the event, but good business risk management calls on us to be wise beforehand. Later in this chapter we will discuss the things that can be learnt from this episode (and that need to be applied in a much wider sphere than just the world of banks and traders.)

1.1 Introduction

In essence, *risk management* is about managing effectively in a risky and uncertain world. Banks and financial services companies have developed some of the key ideas in the area of risk management, but it is clearly vital for any manager. All of us, every day, operate in a world where the future is uncertain.

When we look out into the future there is a myriad of possibilities: there can be no comprehension of this in its totality. So our first step is to simplify in a way that enables us to make choices amidst all the uncertainty. The task of finding a way to simplify and comprehend what the future might hold is conceptually challenging and different individuals will do this in different ways. One approach is to set out to build, or imagine, a set of different possible futures, each of which is a description of what might happen. In this way we will end up with a range of possible future scenarios that are all believable, but have different likelihoods.

Though it is obviously impossible to describe every possibility in the future, at least having a set of possibilities will help us in planning.

One way to construct a scenario is to think of chains of linked events: if one thing happens then another may follow. For example, if there is a typhoon in Hong Kong, then the shipment of raw materials is likely to be late, and if this happens then we will need to buy enough to deal with our immediate needs from a local supplier, and so on. This creates a *causal chain*.

A causal chain may, in reality, be a more complicated network of linked events. But in any case it is often helpful to identify a particular *risk event* within the chain that may or may not occur. Then we can consider both the probability of the risk event occurring and also the consequences and costs if it does. In the example of the typhoon in Hong Kong, we need to bear in mind both the probability of the typhoon and the costs involved in finding an alternative temporary source.

Risk management is about seeking better outcomes, and so it is critical to identify different risk events and to understand both their causes and consequences. Usually risk in this context refers to something that has a negative effect, so that our interest in the causes of negative risk events is to reduce their probability or, better still, eliminate them altogether. We are concerned about the consequences of risk events so that we can act beforehand in a way that reduces the costs if a negative risk event does occur. The open-ended nature of this exercise makes it important to concentrate on the most important causal pathways – we can think of this as identifying *risk drivers*.

At the same time as looking at actions specifically designed to reduce risk, we may need to think about the risk consequences of management decisions that we make. For example, we may be considering moving to an overseas supplier who is able to deliver goods at a lower price but with a longer lead time, so that orders will need to be placed earlier: then we need to ask what extra risks are involved in making this change. In later chapters we will give much more attention to the problems of making good decisions in a risky environment.

Risk management involves planning and acting before the risk event. This is proactive rather than reactive management. We don't just wait and see what happens, with the hope that we can manage our way through the consequences; instead we work out in advance what might happen and what the consequences are likely to be. Then we plan what we should do to reduce the probability of the risk event and to deal with the consequences if it occurs.

Sometimes the risk event is not in our control; for example, we might be dealing with changes in exchange rates or government regulation – usually this is called an *external risk*. On other occasions we can exercise some control over the risk events, such as employee availability, supply and operations issues. These are called *internal risks*. The same distinction between what we can and cannot control occurs with consequences too. Sometimes we can take actions to limit negative consequences (like installing sprinklers for a fire), but at other times there are limits to what we can do and we might choose to insure against the event directly (e.g. purchasing fire insurance).

We will use the term risk management to refer to the entire process:

- *Understanding risk:* both its drivers and its consequences.
- *Risk mitigation:* reducing or eliminating the probability of risk events as well as reducing the severity of their impact.
- *Risk sharing:* the use of insurance or similar arrangement so that some of the risk is transferred to another party, or shared between two parties in some contractual arrangement.

The risk framework we are discussing makes it sound as though all risk is bad, but this is misleading in two ways. First we can use the same approach to consider good outcomes as well as bad ones. This would lead us to try to understand the most important causal chains, with the aim of maximizing the probability of a positive chance event, and of optimizing the benefits if this event does occur. Second we need to recognize that sometimes the more risky course of action is ultimately the wiser one. Managers are schizophrenic about risk. Most see risk taking as part of a manager's role, but there is a tendency to judge whether a decision about risk was good or bad simply by looking at the results. Though it is rarely put in these terms, the idea seems to be that it is fine to take risks provided that nothing actually goes badly wrong! Occasionally managers might talk of 'controlled risk' by which they mean a course of action in which there may be negative consequences but these are of small probability and the size of the cost is tolerable.

In their discussion of the agile enterprise, Rice and Franks (2010) say, 'While uncertainty impacts risk, it does not necessarily make business perilous. In fact, risk is critical to any business – for nothing can improve without change – and change requires risk.' Much the same point was made by Prussian Marshall Helmuth von Moltke in the mid-1800s: 'First weigh the considerations, then take the risks.'

Our discussion so far may have implied an ability to list all the risks and discuss the probability that an individual risk event occurs. But often there is no way to identify all the possible outcomes, let alone enter into a calculation of the probability of their occurrence. Some people use the term *uncertainty* (rather than risk) to refer to this idea. Frank Knight was an economist who was amongst the first to distinguish clearly between these two concepts and he used 'risk' to refer to situations where the probabilities involved are computable. In many real environments there may be a total absence of information about, or awareness of, some potentially significant event. In a much-parodied speech made at a press briefing on February 12, 2002, former US Defense Secretary Donald Rumsfeld said:

> 'There are known knowns. These are things we know that we know. There are known unknowns. That is to say, there are things that we now know we don't know. But there are also unknown unknowns. These are things we do not know we don't know.'

In Chapter 8 we will return to the question of how we should behave in situations with uncertainty, when we need to make decisions without being able to assign probabilities to different events.

1.2 Identifying and documenting risk

Many companies set up a formal *risk register* to document risks. This enables them to have a single point at which information is gathered together and it encourages a careful assessment of risk probabilities and likely responses to risk events.

A carefully documented risk management plan has a number of advantages. There is first of all a benefit in making it more likely that risk will be managed appropriately, with major risks identified and appropriate measures taken. Secondly there is an advantage in defining the responsibility for managing and responding to particular categories of risk. It is all too easy to find yourself in a company in which something goes wrong and no person or department admits to being the responsible party.

Moreover, a risk management plan allows stakeholders to approve the risk management approach and helps to demonstrate that the company has exercised an appropriate level of diligence in the event that things do go wrong.

There are really three steps in setting up a risk register:

1. *Identify the important risk events.* The first step is to make some kind of list of different risks that may occur, and in doing this a systematic process for identifying risk can be helpful. A good starting point is to think about the context for the activity: the objectives; the external influences; the stages that are gone through. The next step is to go through each element of the activity asking what might happen that could cause external factors to change, or that could affect the achievement of any objective.
2. *Understand the causes of the risk events.* Risk does not occur in a vacuum. Having identified a set of risk events, the next step is to come to grips with the factors that are involved in causing the risk events. In order to understand what can be done to avoid these risks, we should ask the following questions, for each risk:
   - How are these events likely to occur?
   - How probable are these events?
   - What controls currently exist to make this risk less likely?
   - What might stop the controls from working?
3. *Assess the consequences of the risk events.* The final step is to understand what may happen as a result of these risk events. The aim is to find ways to reduce the bad effects. For each risk we will want to know:
   - Which stakeholders might be involved or affected? For example, does it affect the return on share capital for shareholders? Does it affect the

assurance of payment for suppliers? Does it affect the security that is offered to our creditors? Does it affect the assurance of future employment for our employees?

- How damaging is this risk?
- What controls currently exist to make this risk less damaging?
- What might stop the controls from working?

At the end of this process we will be in a better position to build the risk register. This will indicate, for each risk identified:

- its causes and impacts;
- the likelihood of this risk event;
- the controls that exist to deal with this risk;
- an assessment of the consequences.

Because the risk register will contain a great many different risks, it is important to focus on the most important ones. We want to construct some sort of priority rating – giving the overall level of risk. This then provides a tool so that management can focus on the most important risk events and then determine a risk treatment plan to reduce the level of risk. The most important risks are those with serious consequences that are relatively likely to occur. We need to combine the likelihood and the impact and Figure 1.1 shows the type of diagram that is often used to do this, with risk levels labeled L = Low; M = Medium; H = High; and E = Extreme.

This type of diagram of risk levels is sometimes called a *heat map*, and often red is used for the extreme risk boxes; orange for the high risks; and yellow for the medium risks. It is a common tool and is recommended in most risk management standards. It should be seen as an important first step in drawing

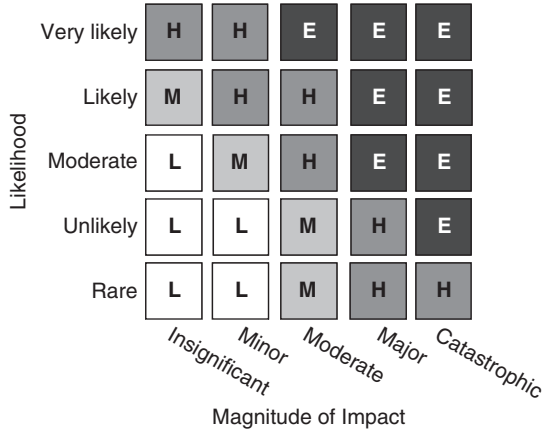

Figure 1.1 Calculating risk level from likelihood and impact.

up a risk management plan, prior to making a much fuller investigation of some specific risks, but nevertheless there are some significant challenges associated with the use of this approach.

One problem is related to the use of a scale based on words like 'likely' or 'rare': these terms will mean very different things to different people. Some people will use a term like 'likely' to mean a more than two thirds chance of occurring (this is the specific meaning that is ascribed in the IPCC climate change report). But in a risk management context, quite small probabilities over the course of a year may seem to merit the phrase 'likely'.

The use of vague terms in a scale of this sort will make misunderstandings far more likely. Douglas Hubbard describes an occasion when he asked a manager 'What does this mean when you say this risk is "very likely"?' and was told that it meant there was about a 20% chance of it happening. Someone else in the room was surprised by the small probability, but the first manager responded, 'Well this is a very high impact event and 20% is too likely for that kind of impact.' Hubbard describes the situation as 'a roomful of people who looked at each other as if they were just realizing that, after several tedious workshops of evaluating risks, they had been speaking different languages all along.' This story illustrates how important it is to be absolutely clear about what is meant when discussing probabilities or likelihoods in risk management.

The heat map method is clearly a rough and ready tool for the identification of the most important risks. But its greatest value is in providing a common framework in which a group of people can pool their knowledge. Far too often the methodology fails to work as well as it might, simply because there has not been any prior agreement as to what the terms mean. A critical point is to have a common view of the time frame or horizon over which risks are assessed. Suppose that there is a 20% probability of a particular risk event occurring in the next year, but the group charged with risk management is using an implicit 10-year time horizon. This would certainly allow them to assess the risk as very likely, since, if each year is independent of the last and the probability does not vary, then the probability that the event does not occur over 10 years is $0.8^{10} = 0.107$. So there is a roughly 90% chance that the event *does* occur at some point over a 10-year period.

More or less the same argument applies to the terms used to identify the magnitude of the impact. It will not be practicable to give an exact dollar figure associated with losses, just as there is little point in trying to ascribe exact probabilities to risk events. But it is worthwhile having a discussion on what a 'minor' or a 'moderate' impact really means. For example, we might initiate a conversation about the evaluation we would give for the impact of an event that led to an immediate 5% drop in the company share price.

1.3 Fallacies and traps in risk management

In this introductory chapter it is appropriate to give some 'health warnings' about the practice of risk management. These are ideas about risk management that can be misleading or dangerous.

It is worth beginning with the observation that society at large is increasingly intolerant of risk which has no obvious owner – no one who is responsible and who can be sued in the event of a bad outcome. Increasingly it is no longer acceptable to say 'bad things happen' and we are inclined to view any bad event as someone's fault. This is associated with much management activity that could be characterized as 'covering one's back'. The important thing is no longer the risk itself but the demonstration that appropriate action has been taken so that the risk of legal liability is removed. The discussion of risk registers in the previous section demonstrates exactly this divergence between what is done because it brings real advantage, and what is done simply for legal reasons. Michael Power makes the case that greater and greater attention is placed on what might be called *secondary risk management*, with the sole aim of deflecting risk away from the organization or the individuals within it. It is fundamentally wrong to spend more time ensuring that we cannot be sued than we do in trying to reduce the dangers involved in our business. But in addition to questions of morality, a focus on secondary risk management means we never face up to the question of what is an appropriate level of risk, and we may end up losing the ability to make sound judgments on appropriate risks: the most fundamental requirement for risk management professionals.

Another trap we may fall into is the feeling that good risk management requires a scenario-based understanding of all the risks that may arise. Often this is impossible, and trying to do so will distract attention from effective management of important risks. As Stulz (2009) argues, there are two ways to avoid this trap. First there is the use of statistical tools (which we will deal with in much more detail in later chapters).

> 'Contrary to what many people may believe, you can manage risks without knowing exactly what they are – meaning that most of what you'd call unknown risks can in fact be captured in statistical risk management models. Think about how you measure stock price risk. ... As long as the historical volatility and mean are a good proxy for the future behavior of stock returns, you will capture the relevant risk characteristics of the stock through your estimation of the statistical distribution of its returns. You do not need to know why the stock return is $+10\%$ in one period and -15% in another.'

The second way to avoid getting bogged down in an unending set of almost unknowable risks is to recognize that important risks are those that make a difference to management decisions. Some risks are simply so low in probability that a manager would not change her behavior even if this risk was brought to her attention. This is like the risk of being hit by an asteroid – it must have some small probability of occurring but it does not change our decisions.

A final word of caution relates to the use of historical statistical information to project forward. We may find a long period in which something appears to be varying according to a specific probability distribution, only to have this change quite suddenly. An example with a particular relevance for the author is in the

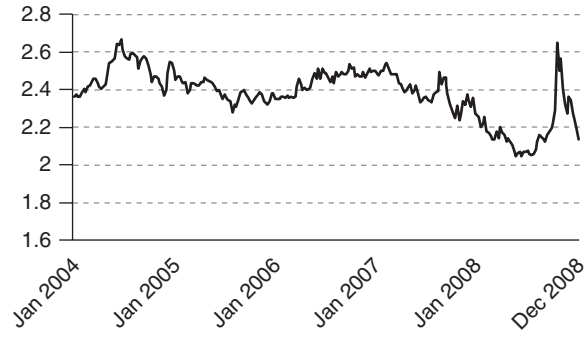

Figure 1.2 Australian dollars to one British pound 2004–2008.

exchange rate between the Australian dollar and the British pound. The graph in Figure 1.2 shows what happened to this exchange rate over a five-year period from 2004 to 2008.

The weekly data here have a mean of 2.38 Australian dollars per pound and the standard deviation is 0.133. Fifteen months later, in March 2010, the rate had fallen to 1.65 (and continued to fall after that date). Now, if weekly exchange rate data followed a normal distribution then the chance of observing a value as low as 1.65 (more than five standard deviations below the mean) would be completely negligible. Obviously the foreign exchange markets do not behave in quite the way that this superficial historical analysis suggests. Looking over a longer period and considering also other foreign exchange rates would suggest that the relatively low variance over the five-year period taken as a base was unusual. In this case the fallout from the global financial crisis quickly led to exchange rate values that reflect historically very high levels for the Australian dollar and a low level for the British pound.

We may be faced with the task of estimating the risk of certain events on the basis of statistical data but without the benefit of a very long view and with no opportunity to compare any related data. In this situation all that we might have to guide us is a set of data like Figure 1.2. Understanding how hard it is in a foreign exchange context to say what the probabilities are of certain outcomes should help us to be cautious when faced with the same kind of task in a different context.

1.4 Why safety is different

This book is about business risk management and is aimed at those who will have management responsibility. There are significant differences between how we may behave as managers and how we behave in matters of our personal safety. Every day as we grow up, and throughout our adult lives, we make decisions which involve personal risk. The child who decides to try jumping off the playground swing is weighing up the risk of getting hurt against the excitement involved. And the driver who overtakes a slower vehicle on the road

is weighing up the risks of that particular road environment against the time or frustration saved. In that sense we are all risk experts; it's what we do every day.

It is tempting to think about safety within the framework we have laid out of different risk events, each with a likelihood and a magnitude of impact. With this approach we could say that a car trip to the shops involves such a tiny likelihood of being involved in a collision with a drunk driver that the overall level of risk is easily outweighed by the benefits. But there are two important reasons why thinking in this way can be misleading.

First we need to consider not only the likelihood of a bad event, but also its consequences. And if I am worried about someone else driving into me, then the consequence might be the loss of my life. Just how does that get weighed up against the inconvenience of not using a car? Most of us would simply be unable to put a monetary value on our own lives, and no matter how small the chance of our being killed in a car crash, the balance will tilt against driving the car if we make the value of our life high enough. But yet we still drive our cars and do all sorts of other things that carry an element of personal risk.

A second problem with treating safety issues in the same way as other risks is that the chance of an accident is critically determined by the degree of care taken by the individual concerned. The probability of dying in a car crash on the way to the shops is mostly determined by how carefully I drive. This makes my decision on driving a car different to a decision on traveling by air, where once on board I have no control over the level of risk. However, there are many situations where being careful will dramatically reduce the risk to our personal safety. Paradoxically, the more dangerous we perceive the activity to be then the more careful we are. The risks from climbing a ladder may end up being greater than from using a chain saw if we believe that the ladder is basically safe, but that the chain saw is extremely dangerous.

A better way to consider personal safety is to think of each of us as having an in-built 'risk thermostat' that measures our own comfort level with different levels of risk. As we go about our lives there comes a time with certain activities when we start to feel uncomfortable with the risk we are taking; this happens when the amount of risk starts to exceed our own risk thermostat setting. The risk we will tolerate varies according to our own personalities, our age, our experience of life, etc. But if the level of risk is below this personal thermostat setting then there is very little that holds us back from increasing the risk. So, if driving seems relatively safe then we will not limit our driving to occasions when the benefits are sufficiently large. John Adams points out that some people will actively seek risk so that they return to the risk thermostat setting which they prefer. So, in discussing the lives that might be saved if motorcycling was banned, he points out that, 'If it could be assumed that all the banned motorcyclists would sit at home drinking tea, one could simply subtract motorcycle accident fatalities from the total annual road accident death toll. But at least some frustrated motorcyclists would buy old bangers and try to drive them in a way that pumped as much adrenaline as their motorcycling'.

These are important issues and need to faced by businesses in which health and safety are big concerns, such as mining. If the aim is to get as close as possible to eliminating accidents in the workplace, then it is vital to pay attention to the workplace culture, which can have a role in resetting the risk thermostat of our employees to a lower level.

1.5 The Basel framework

The Basel Accords refer to recommendations made by the Basel Committee on Banking Supervision about banking regulations. The second of these accords (Basel II) was first published in 2004 and defines three different types of risk for banks – but the framework is quite general and can apply to any business.

Market risk. Market risk focuses on the uncertainties that are inherent in market prices which can go up or down. Market risk applies to any uncertainty where the value is dependent on prices that cannot be predicted fully in advance. For example, we might build a plant to extract gold from a low-yield resource, but there is a risk that the gold price will drop and our plant will no longer be profitable. This is an example of a *commodity risk*. Other types of market risk are *equity risk* (related to stock prices and their volatility); *interest rate risk*; and *currency risk* (related to foreign exchange rates and their volatility).

Credit risk. Any business will be involved in many different contractual arrangements. If the counterparty to the contract does not deliver what is promised then legal means can be used to extract what is owed. But this assumes that the counterparty still has funds available. Credit risk is the risk of a counterparty to a contract going out of business. For example, a business might deliver products to its customers and have 30-day payment terms. If the customer goes out of business there may be no way of getting back more than a small percentage of what is owed. In its most direct form, the contract is a loan made to another party and credit risk is about not being repaid due to bankruptcy.

Operational risk. Operational risk is about something going badly wrong. This category of risk includes many of the examples we have discussed so far that are associated with negative risk events. Operational risk is defined as arising from failures in internal processes, people or systems, or due to external events.

Since we are interested in more general risk management concerns, not just risk for banks, it is helpful to add a fourth category to the three discussed by Basel II.

Business risk. Business risk relates to those parts of our business value proposition where there is considerable uncertainty. For example, there may be

a risk associated with changes in costs, or changes in customer demand, or changes in the security of supply of raw materials. Business risk is like market risk but does not relate directly to prices.

Both market risk and credit risk are, to some extent, entered into deliberately as a result of calculation. Market risk is expected, and we can make calculations on the basis of the likelihood of different market outcomes. Business risk also often has this characteristic: for example, most businesses will have a clear idea of what will happen under different scenarios for customer demand. Credit risk is always present, and in many cases we assess credit risk explicitly through credit ratings. But operational risk is different: it is not entered into in the expectation of reward. It is inherent and is, in a sense, the unexpected risk in our business. It may well fit into the 'unknown unknown' description in the quotation from Rumsfeld that we gave earlier. Usually operational risk involves low-probability and high-severity events and this makes it particularly challenging to deal with.

1.6 Hold or hedge?

When dealing with market or business risk a manager is often faced with an ongoing risk, so that it recurs from day to day or month to month. In this case there is the need to take strategic decisions related to these risks.

An example of a recurring risk occurs with airlines who face ongoing uncertainty related to the price of fuel (which can only be partially offset by adding fuel surcharges). The question that managers face is: when to hold on to that risk, when to insure or hedge it, and when to attack the risk so that it is reduced?

A financial hedge is possible when we can buy some financial instrument to lessen the risk of market movements. For example, a power utility company might trade in futures for gas prices. If the utility is buying gas and selling electricity then it is exposed to a market risk if the price of gas rises and it is not able to raise the price of electricity to the same extent. By holding a futures contract on the gas price, the company can obtain a benefit when the price of gas increases: if the utility knows how much gas it will purchase then the net effect will be to fix the gas price for the period of the contract and eliminate this form of market risk. Even if the utility cannot exactly predict the amount of gas it will burn, there will still be the opportunity to hedge the majority of its potential losses from gas price rises.

Sometimes we have an operational hedge which achieves the same thing as a financial hedge through the way that our operations are organized. For example, we may be concerned about currency risk if our costs are primarily in US dollars but our sales are in the Euro zone. Thus, if the Euro's value falls sharply relative to the US dollar, then we may find our income insufficient to meet our manufacturing expenses even though our sales have remained strong. An option is to buy a futures contract which has the effect of locking in an exchange rate. However, another 'operational hedge' could be achieved by moving some

of our manufacturing activity into a country in the Euro zone, so that more of our costs occur in the same currency as the majority of our sales.

In holding on to a risk the company deliberately decides to accept the variation in profit which results. This may be the best option when a company has sufficient financial resources, and when it has aspects of its operations that will limit the consequences of the risk. For example, a vertically integrated power utility company that sets the price of electricity for its customers may decide not to fully hedge the risks associated with rises in the cost of gas if there are opportunities to quickly change the price of the electricity that it sells in order to cover increased costs of generation.

1.7 Learning from a disaster

We began this chapter with the remarkable story of Jérôme Kerviel's massive fraud at Société Générale, which fits into the category of operational risk. Now we return to this example with the aim of seeing what can be learnt. To understand what happened we will start by giving some background information on the world of bank trading. A bank, or any company involved in trading in a financial marketplace, will usually divide its activities into three areas. First the traders themselves: these are the people who decide what trades to make and when to make them (the 'front office'). Second, a risk management area responsible for monitoring the traders' activity measuring and modeling risk levels etc. (the 'middle office'). And finally an area responsible for carrying out the trades, making the required payments and dealing with the paperwork (the 'back office').

The trading activities are organized into *desks*: groups of traders working with a particular type of asset. The Kerviel story takes place in SocGen's Delta One desk in Paris. Delta One trading refers to buying and selling straightforward derivatives that do not involve any options. Options are derivatives which give 'the right but not the opportunity' to make a purchase or sale. The trading of options gives a return that depends non-linearly on whatever is the underlying security (we explain more about this in Chapter 9), but trading activities for a Delta One desk are simpler than this – the returns just depend directly on what happens to the underlying security. In fact, the delta in the terminology refers to the first derivative of the return as a function of the underlying security, and 'Delta One' is shorthand for 'delta equals one,' implying this direct relationship.

For example, a trade might involve buying a future on the DAX, which is the main index for the German stock market and comprises the 30 largest and most actively traded German companies. Futures can be purchased in relation to different dates (the end of each quarter) and are essentially a prediction of what the index will be at that date. One can also buy futures in the individual stocks that make up the index and by creating a portfolio of these futures in the proportions given by the weights in the DAX index, one would mimic the behavior of the future for the index as a whole. However, over time the weights in the DAX index are updated (in an automatic way based on market capitalization),

so holding the portfolio of futures on individual stocks would lead to a small divergence from the DAX future over a period of time.

The original purpose of a Delta One trading desk is to carry out trades for the bank's clients, but around that purpose has grown up a large amount of proprietary trading where the bank intends to make money on its own account. One approach is for a trader to make a bet on the convergence of two prices that (in the trader's view) should be closer than they are. If the current price of portfolio A is greater than that of portfolio B and the two prices will come back together before long, then there will be an opportunity to profit by buying B and selling A, and then reversing this transaction when the prices move closer together. Since both portfolios are made up of derivatives, the 'buying' and 'selling' here need not involve ownership of the underlying securities, just financial contracts based on their prices. This type of trading, which intends to take advantage of a mis-pricing in the market, is called an arbitrage trade, and since trades of one sort are offset by trades in the opposite direction, the risk involved should, in theory, be very low.

Many of these trading activities take advantage of quite small opportunities for profit (in percentage terms) and therefore, in order to make it worthwhile, they require large sums of money to be involved. Kerviel was supposed to act as an arbitrageur, looking for small differences in price between different stock index futures. In theory this means that trades in one direction are offset by balancing trades in the other direction. But Kerviel was making fictitious trades: reporting trades that did not occur. This enabled him to hold one half of the combined position but not the other. The result of the fictitious trade is to change an arbitrage opportunity with large nominal value but relatively small risk into a simple (very large) bet on the movement of the futures price.

When Kerviel started on this process in 2006 things went reasonably well–his bets came off and the bank profited. Traders are allowed to make some speculative trades of this sort, but there is a strict limit on the amount of risk they take on: Kerviel breached those limits repeatedly (and spectacularly). Over time the amounts involved in these speculations became greater and greater, and things still went well. During 2007 there were some ups and downs in the way that these bets turned out, but by the end of the year Kerviel was well ahead. He has claimed that his speculation made 1.5 billion Euros in profits for SocGen during 2007. None of this money made its way to him personally; he would only have profited through earning a large bonus that year.

In January 2008, however, his good fortune was reversed when some large bets went very wrong. The senior managers at the bank finally discovered what was happening on January 18th 2008. There were enormous open positions and SocGen decided that it had no option but to close off those positions and take the losses, whatever these turned out to be. The timing was bad and the market was in any case tumbling; the net result was that SocGen lost 4.9 billion Euros. The news appeared on January 24th. The sums of money involved are enormous and a smaller bank would certainly have been bankrupted by these losses, but SocGen is very large and some other parts of its operation had been going well.

Nevertheless, the bank was forced into seeking an additional 5.5 billion Euros in new capital as a result of the losses.

Banks such as SocGen have elaborate mechanisms to ensure that they do not fall into this kind of situation. Outstanding positions are checked on a daily basis, but each evening Kerviel, working late into the night, would book offsetting fictitious transactions, without any counterparties, and in this way ensure that his open positions looked as if they were appropriately hedged. Investigations after the event revealed more than a thousand fake trades; there is no doubt that these should have been picked up.

Kerviel, who was 31 when the scandal broke, was tried in June 2010. He acknowledged what he had done in booking fake trades, but he argued that his superiors had been aware of what he was doing and had deliberately turned a blind eye. He said 'It wasn't me who invented these techniques – others did it, too.' Finally, in October 2010, Kerviel was found guilty of breach of trust, forging documents and entering false data into computers; he was sentenced to three years in prison and ordered to repay SocGen's entire trading loss of 4.9 billion Euros. The judge held Kerviel solely responsible for the loss and said that his crimes had threatened the bank's existence. The case came to appeal in October 2012 and the original verdict was upheld. There is, of course, no possibility of Kerviel ever repaying this vast sum, but SocGen's lawyers have said that they will pursue him for any earnings he makes by selling his story.

1.7.1 What went wrong?

There is no doubt that what happened at SocGen came about because of a combination of factors. First there was Kerviel himself, who had some knowledge of the risk management practices of the middle office through previously having worked in this area. It seems that he kept some access appropriate to this, even when he became a trader. This is exactly what happened with Nick Leeson at Barings – another famous example of a trader causing enormous losses at a bank. Kerviel was someone whose whole world was the trading room and, over the course of a year or so, he was drawn into making larger and larger bets with the company's money. There remains a mystery about what might have been his motivation. In his appeal he offered no real explanation, simply describing his actions as 'moronic', but maintaining that he was someone trying 'to do his job as well as possible, to make money for the bank'.

A second factor was the immediate supervision at the Delta One desk. Whether or not one accepts Kerviel's claims that his bosses knew what was going on, they certainly should have known and done something about it. It is hard at this point to determine what is negligence and what is tacit endorsement. Eric Cordelle, who was Kerviel's direct superior, was only appointed head of the Delta One desk in April 2007, and did not have any trading experience. He was sacked for incompetence immediately after the fraud was discovered. He claims that during this period his team was seriously understaffed and he had insufficient time to look closely at the activities of individual traders.

A third important factor is the general approach to risk management being taken at SocGen in this period. It is easy to take a relaxed attitude to the risk of losses when everything seems to be going well. During 2007 there was an enormous increase in trading activity on the Delta One desk and large profits were being made. The internal reports produced by SocGen following the scandal were clear that there had been major deficiencies in the monitoring of risks by the desk. The report by PriceWaterhouseCoopers on the fraud stated that: 'The surge in Delta One trading volumes and profits was accompanied by the emergence of unauthorized practices, with limits regularly exceeded and results smoothed or transferred between traders.' Moreover, 'there was a lack of an appropriate awareness of the risk of fraud.'

In fact there were several things which should have alerted the company to a problem:

- there was a huge jump in earnings for Kerviel's desk in 2007;
- there were questions which were asked about Kerviel's trades from Eurex, the German derivatives exchange, who were concerned about the huge positions that Kerviel had built up;
- there was an unusually high level of cash flow associated with Kerviel's trading;
- Kerviel did not take a vacation for more than a few days at a time – despite a policy enforcing annual leave;
- there was a breach of Kerviel's market risk limit on one position.

We can draw some important general lessons from this case. I list five of these below.

1. *Company culture is more important than the procedures.* The organizational culture in SocGen gave precedence to the money-making side of the business (trading) over the risk management side (middle office), and this is very common. Whether or not procedures are followed carefully will always depend on cultural factors, and the wrong sort of risk culture is one of the biggest factors leading to firms making really disastrous decisions.
2. *Good times breed risky behavior.* In the SocGen case the fact that Kerviel's part of the operation was doing well made it easy to be lax in the care with which procedures were carried out. It may be true that the reverse of this statement is also true: in bad times taking risks may seem the only way through, but whether wise or not these are at least a conscious choice. Risks that managers enter into unconsciously seem to generate the largest disasters.
3. *Companies often fail to learn from experience.* One example occurs when managers ignore risks in similar companies, such as we see in the uncanny

resemblance between SocGen and Barings. But it can also be surprisingly hard to learn from our own mistakes in a corporate setting. Often a scapegoat is found and moved on, without a close look at what happened and why. Dwelling on mistakes is a difficult thing to do and will inevitably be perceived as threatening, and perhaps that is why a careful analysis of bad outcomes is often ducked.

4. *Controls need to be acted upon.* On many occasions risks have been considered and controls put in place to avoid them. The problem occurs when the controls that are in place are ignored in practice. SocGen had a clear policy on taking leave (as is standard in the industry) but failed to act upon it.

5. *There must be adequate management oversight.* Inadequate supervision is a key ingredient in poor operational risk management. In the SocGen case, Kerviel's supervisor had inadequate experience and failed to do his job. More generally, risks will escalate when a single person or a small group can make decisions that end with large losses, either through fraud or simple error. Companies need to have systems that avoid this through having effective oversight of individuals by managers, who need to supervise their employees sufficiently closely to ensure that individuals do what they are supposed to do.

This book is mostly concerned with the quantitative tools that managers can use in order to deal with risk and uncertainty. It is impossible to put into a single book everything that a manager might need to know about risk. In fact, the most important aspects of risk management in practice are things that managers learn through experience better than they learn in an MBA class. But paying attention to the five key observations above will be worthwhile for anyone involved in risk management, and may end up being more important than all the quantitative methods we are going to explore later in this book.

It is hard to overstate the importance of the culture within an organization: this will determine how carefully risks are considered; how reflective managers are about risk issues; and whether or not risk policies are followed in practice. A culture that is not frightened by risk (where employees are prepared to discuss risk openly and consider the appropriate level of risk) is more likely to avoid disasters than a culture that is paranoid about risk (where employees are uncomfortable in admitting that risks have been not been eliminated entirely). It seems that when we are frightened of risk we are more likely to ignore it, or hide it, than to take steps to reduce it.

Notes

This chapter is rather different than the rest of the book: besides setting the scene for what follows, it also avoids doing much in the way of quantification. I have tried to distill some important lessons rather than give a set of models to be

used. I have found the book by Douglas Hubbard one of the best resources for understanding the basics of risk management applied in a broad business context. His book covers not only some of the material in this chapter but also has useful things to say about a number of topics we cover in later chapters (such as the question of how risky decisions are actually made, which we cover in Chapter 6).

A good summary of approaches which can be used to generate scenarios and think about causal chains as well as the business responses can be found in Miller and Waller (2003). The discussion on why we need to think differently about safety issues is taken from the influential book by John Adams, who is a particular expert on road safety.

The material on the Société Générale fraud has been drawn from a number of newspaper articles: Société Générale loses $7 billion in trading fraud, *New York Times*, January 24, 2008; Bank Outlines How Trader Hid His Activities, *New York Times*, January 28, 2008; A Société Générale Trader Remains a Mystery as His Criminal Trial Ends, *New York Times*, June 25, 2010.; Rogue Trader Jerome Kerviel 'I Was Merely a Small Cog in the Machine' *Der Spiegel* Online, November 16, 2010.

We have said rather little about company culture and its bearing on risk, but this is by no means a commentary on the importance of this aspect of risk, which probably deserves a whole book to itself (some references on this are Bozeman and Kingsley, 1998; Flin *et al*., 2000; Jeffcot *et al*., 2006 as well as the papers in the book edited by Hutter and Power, 2005).

References

Adams, J. (1995) *Risk*. UCL Press.

Bozeman, B. and Kingsley, G. (1998) Risk Culture in Public and Private Organizations. *Public Administration Review*, **58**, 109–118.

Flin, R., Mearns, K., O'Connor, P. and Bryden, R. (2000) Measuring safety climate: identifying the common features. *Safety Science*, **34**, 177–192.

Hubbard, D. (2009) *The Failure of Risk Management*. John Wiley & Sons.

Hutter, B. and Power, M. (2005) *Organizational Encounters with Risk*. Cambridge University Press.

Jeffcott, S., Pidgeon, N., Weyman, A. and Walls, J. (2006) Risk, trust, and safety culture in UK train operating companies. *Risk Analysis*, **26**, 1105–1121.

Miller, K. and Waller, G. (2003) Scenarios, real options and integrated risk management. *Long Range Planning*, **36**, 93–107.

Power, M. (2004) The risk management of everything. *Journal of Risk Finance*, **5**, 58–65

Rice, J. and Franks, S. (2010) Risk Management: The agile enterprise. *Analytics Magazine*, INFORMS.

Ritchie, B. and Brindley, C. (2007) Supply chain risk management and performance. *International Journal of Operations & Production Management*, **27**, 303–322.

Stulz, R. (2009) Six ways companies mismanage risk. *Harvard Business Review*, **87** (3), 86–94

Exercises

1.1 Supply risk for valves

DynoRam makes hydraulic rams for the mining industry in Australia. It obtains a valve component from a supplier called Sytoc in Singapore. The valves cost 250 Singapore dollars each and the company uses between 450 and 500 of these each year. There are minor differences between valves, with a total of 25 different types being used by DynoRam. Sytoc delivers the valves by air freight, typically about 48 hours after the order is placed. Deliveries take place up to 10 times a month depending on the production schedule at DynoRam. Because of the size of the order, Sytoc has agreed a low price on condition that a minimum of 30 valves are ordered each month. On the 10th of each month (or the next working day) DynoRam pays in advance for the minimum of 30 valves to be used during that month and also pays for any additional valves (above 30) used during the previous month.

(a) Give one example of market risk, credit risk, operational risk and business risk that could apply for DynoRam in relation to the Sytoc arrangement.

(b) For each of the risks identified in part (a) suggest a management action which would have the effect either of reducing the probability of the risk event or minimizing the adverse consequences.

1.2 Connaught

The following is an excerpt from a newspaper report of July 21, 2010 appearing in the UK *Daily Telegraph*.

> 'Troubled housing group Connaught has been driven deeper into crisis after it discovered a senior executive sold hundreds of thousands of pounds worth of shares ahead of last month's shock profit warning.
>
> The company which lost more than 60% of its value in just three trading days in June, and saw its chief executive and finance director resign, has launched an internal investigation into the breach of city rules... Selling shares with insider information when a company is about to disclose a price-sensitive statement is a clear breach of FSA rules.
>
> Connaught, which specializes in repairing and maintaining low cost (government owned) housing, has fallen a total of 68% since it gave a warning that a number of public sector clients had postponed capital expenditure, which would result in an 80 million pound fall in expected revenue this year.

> The group said that it had been hit by deferred local authority contracts which would knock 13m pounds off this year's profits and 16m pounds from next year's. It also scaled back the size of its order book from the 2.9 billion pounds it said it was worth in April to 2.5 billion.
>
> The profit warning also sparked renewed concerns about how Connaught accounts for its long-term repair and maintenance contracts. Concerns first surfaced late last year with city analysts questioning whether the company was being prudent when recognizing the revenue from, and costs of, its long term contracts.
>
> The company vehemently defended its accounting practices at the time and continues to do so. Chairman Sir Roy Gardner has tried to steady the company since his arrival earlier this year.'

(a) How would you describe the 'profits warning' risk event: is it brought about by market risk, credit risk, operational risk or business risk?

(b) From the newspaper report can you make any deductions about risk management strategies the management of the company could have taken in advance of this in order to reduce the loss to shareholders?

1.3 Bad news stories

Go through the business section of a newspaper and find a 'bad news' story, where a company has lost money.

(a) Can you identify the type of risk event involved: market risk, credit risk, operational risk or business risk?

(b) Look at the report with the aim of understanding the risk management issues in relation to what happened. Was there a failure to anticipate the risk event? Or a failure in the responses to the event?

1.4 Product form for heat map

Suppose that the risk level is calculated as the expected loss and that the likelihoods are converted into probabilities over a 20-year period as follows: 'very likely' $= 0.9$; 'likely' $= 0.7$; 'moderate' $= 0.4$; 'unlikely' $= 0.2$; and 'rare' $= 0.1$. Find a set of dollar losses associated with the five different magnitudes of impact such that the expected losses are ordered in the right way for Figure 1.1: in other words, so that the expected losses for a risk level of low are always lower than the expected losses for a risk level of medium, and these are lower than the expected losses for a risk level of high, which in turn are lower than the expected losses for a risk level of extreme. You should set the lowest level of loss ('insignificant') as \$10 000.

1.5 Publication of NHS reform risk register

Risk registers may take various forms, but the information they contain can sometimes be extremely sensitive. In 2012 the UK government discussed whether or not to release the full risk register that had been created for the highly controversial reform of the National Health Service. The health secretary, Andrew Lansley, told parliament in May 2012 that only an edited version of this document would be made available, on the principle that civil servants should be able to use 'direct language and frank assessments' when giving advice to ministers. Lansley argued that if this advice were to be released routinely then 'future risk registers [would] become anodyne documents of little use.' The net result would be that 'Potential risks would be more likely to develop without adequate mitigation. That would be detrimental to good government and very much against the public interest.' Using this example as an illustration, discuss whether there is a tension between realistic assessment of risk and the openness that may be implicit in a risk register.