

Chapter 1

Living with Risk

In This Chapter

- ▶ Exploring the idea of risk
 - ▶ Managing financial risk
 - ▶ Informing people about risk
-

Life is risk, and risk is life. Nobody knows what tomorrow may bring. As the poet Robert Burns famously put it, ‘The best-laid schemes o’ mice an’ men, gang aft agley, an’ lea’e us nought but grief an’ pain, for promis’d joy!’ (Roughly translated, Burns warns that careful plans can come to nothing.)

While most of us instinctively first think about bad risk, good surprises happen as well. ‘Fortune favours the bold,’ we are told, and, ‘Sometimes things just go your way.’ In fact, risk is more than just sometimes good, it is essential. As another saying goes, ‘The only place with people and no risk is a graveyard.’ Religions, philosophies and especially superstitions are deeply rooted in ideas about risk.

My topic is managing risk, not risk itself, which means that I don’t cover all the risks you can’t control – the sun going supernova tomorrow or being diagnosed with a genetic heart condition, for examples. Also, my topic is financial risk, so I don’t talk about risks that aren’t priced in the financial markets. That still leaves me with a large topic, but one I can cover in enough detail to be useful.

Understanding the Scope of Risk

Finance professor Elroy Dimson defined *risk* as meaning that more things can happen than will happen. Although stated in a folksy way, this idea is a deep one that comes from information theory and statistical thermodynamics. The tremendous range of future possibilities creates a kind of force – a tendency to disorder, a decay of information – called *entropy*. Entropy isn’t a physical

force like gravity or magnetism, yet in the long run it determines both the fate of the universe and whether the ‘best-laid schemes o’ mice an’ men’ bring grief and pain or promised joy.

Everything humans try to do can be thought of as attempts to influence what will happen, but even the most precise and complicated plans are vastly simpler than the range of things that might happen. This essential feature of risk is lost when risk is reduced to probability distributions. These distributions require that the range of future outcomes is known exactly. In most cases of practical interest, probabilities can be estimated reliably only for outcomes that have actually happened in the past, and they only have much use if decisions are repeated often enough that each potential outcome actually happens.

This doesn’t mean that conventional statistical analysis is useless – far from it. I’m a big fan of quantitative reasoning. But the risk in risk management is something distinct from the risk that can be modelled with probability distributions.

One popular approach is to model risk as a casino game. This *frequentist* approach can yield insights, but it is very limited. Casino games can be played over and over, and have a known range of outcomes with known probabilities. Real risks only happen once, and you can only guess at the range of outcomes and probabilities. Author Nassim Taleb has dubbed this approach the *Ludic Fallacy*. If all risks were playing roulette or drawing cards, we wouldn’t need risk managers.

Another popular approach, called *Bayesian*, treats all risk like bets on a sporting event. This is more accurate than the frequentist approach because it can handle events that only happen once, with some unknown potential outcomes and only guesses about probability. But it is still a limited model that does not capture all important aspects of risk. Risk managers draw on a broad spectrum of risk models, frequentist and Bayesian, plus models drawn from evolution, statistical thermodynamics, behavioural studies and game theory. And they know that even with all the different analytic approaches, important aspects of risk are missed.

Consider a teacup. You know that teacups can shatter into shards and dust, and also that shards and dust never spontaneously recombine into a teacup. Why? Because of all the possible arrangements of the atoms that make up a teacup, only a negligible fraction actually are a teacup. That’s all you have to know to predict that a teacup is fragile. It can shatter, but it can’t self-construct. Any sufficiently large change in conditions – impact, temperature or others – will destroy it. If I have a china shop, I know that it won’t last forever; I don’t need a bull to destroy it. Risk and time are enough.

Some things in the universe do come into being spontaneously – stars, for example, and people and crystals. In many cases these things gain from disorder and change. They can be destroyed, but they can also recreate without outside help.

The same thing is true of human plans and institutions. Some are fragile. Disorder and change only hurt them. Such plans will fail, however solid they seem. Perversely, people often respond to risk by building in more fragility, making the teacup heavier and stronger but no less exposed to risk and time. Risk managers don't ask how strong your teacup is, they ask how it will respond to the unexpected events that the future will bring. Will it gain or lose? That's what really matters, because although the events are individually unexpected, you can be certain that unexpected events will occur.



Risk management isn't about predicting or preventing disaster. Risk management isn't about estimating probabilities or outcomes. It *is* about constructing plans or institutions that will thrive under disorder. It's not about guessing what will happen – in fact, people who guess are the enemies of risk management. Risk management is preparing for anything that might happen. Preparing not just in the sense of having contingency plans to avoid problems, but also in the sense of being ready to take maximum advantage of opportunities.


Measuring risk

I don't talk much about measuring risk. For the most part, risk that can be measured can be insured, avoided, hedged or diversified away. Generally I insist that line risk takers do all the measurement and mitigation they can before I take over the job of managing the residual risk.

Of course, there's room for risk measurement in risk management but less than outsiders tend to think. In addition, it's definitely true that bad risk measurements, as well as inappropriate attempts by inexperienced risk managers to measure non-measurable risks, do a lot more harm in risk management than good risk measurements do good. (I talk about the various components of risk in Chapter 6.)

To see what I mean, consider the graph in Figure 1-1, which shows the distribution of daily returns for the S&P 500 index over the last 50 years.

You have various ways to measure the spread illustrated by this graph. You can compute a standard deviation, a mean absolute deviation, an interquartile range or something else. For that matter, you can just reproduce the graph. However, there's something misleading about representing the data this way: You cannot see the essential risk on this graph, and the risk you think you see is largely irrelevant.



In round terms, the stock market has turned £1 into £100 over the last 50 years. On about 99 days out of 100, the market moved less than 3.5 per cent in either direction. But consider the 80 days on which the market went up more than 3.5 per cent. They're barely visible on the chart, but collectively they caused about a 4,000 per cent increase in wealth. All other days were responsible for about a 150 per cent increase. If you consider the 60 days when the market went down more than 3.5 per cent, they collectively turned £1 into £0.03.

Now the 150 per cent increase from the 99 per cent of normal days isn't insignificant. However, most of the action, especially to a risk manager, happens in the 1 per cent of extreme days, which are nearly invisible. This percentage isn't true just of stock market returns, but also true of many important things in the world.

Consider the risk going forward, which of course is what matters. Suppose that you're considering an investment in stocks with a 1,000-day horizon – about four years of trading days. You expect to get 990 normal days in which

the market moves less than 3.5 per cent. You may get 996 or 987 or even 1,000 such days; but you won't get much different from 990. Also, getting a few days more or less won't matter much because the average return on these days is 0.04 per cent, and no day can make a difference of more than 3.5 per cent. With 990 or so events and limited range, you're highly likely to get something quite close to the expected outcome. Moreover, you have lots and lots of historical data on what happens on normal days, so you're reasonably confident you know what the expected outcome is. There just isn't a lot of risk in 99 per cent of the days, and what risk does exist can be easily handled by front-line risk takers. After all, if they couldn't handle the stuff that happens 99 days out of 100, you'd have noticed long ago.

You also expect to get about five days when the market loses more than 3.5 per cent, plus about five days when the market gains more than 3.5 per cent. However, there's a lot of potential variability around those numbers. You might get 2 or 8 or even 0 or 10 or more of either one. Each one of these days is significant as they average about a 5 per cent move, and may be as large as -28 per cent or +18 per cent. With only a few events, you can get outcomes far away from the mean. Moreover, you have little historical data, you don't really know how big these days can get; and you can't be confident that your front-line risk takers are prepared for them unless you check.



If you take a closer look, you have even more reason to be concerned about a small number of big days. Markets often don't function properly. You may not be able to trade the way you usually do or at all. Financial intermediaries may fail. Trades may be reversed after the fact. Events may trigger investigations and fines. Financial instruments don't move together as they usually do – correlations are different on big days.

Another problem is that the big days in the market can seldom be tied to observable economic events. On normal days, some fraction of stock price movements occurs in discrete jumps after clear news events such as central bank actions or corporate earnings announcements. A lot of unexplainable *noise* (price movements that cannot be easily explained) is evident too (which doesn't stop commentators from jumping in with explanations after the fact), but it's possible to imagine that prices are changing in response to economic news. On many of the biggest days, no news turns up at all, and on others, the extent and timing of the price move is inconsistent with the news the market is supposed to be reacting to.

If that weren't enough, not all the days the stock market makes big moves are abnormal; some are just normal big moves. On the other hand, on some abnormal days, the market behaves strangely but prices don't move a lot by the end of the day, such as the Flash Crash of May 2010 or the Quant Equity Crisis of August 2007. In addition, you need to consider days missing from the graph because the stock market was closed, such as the days after the 9/11 attacks.

The point is that almost everything a risk manager is concerned about is missing from the graph in Figure 1-1, or is nearly invisible on it. Therefore, any measurement of the graph is of only marginal use to a risk manager. Doing sophisticated analytics on the 99 per cent of normal days can be useful to line risk takers, but it's false precision to a risk manager.



Consider Nassim Taleb's example of a casino that can measure the risks of the bets it makes with its customers at the roulette and craps tables. This risk averages out quickly, and a risk manager who focuses on it would be wasting his time. The three biggest losses of one particular casino in one year were:

- ✓ The star performer was mauled by a tiger.
- ✓ The owner's daughter was kidnapped and held for ransom.
- ✓ It was discovered that a long-time, low-level employee, for unexplainable reasons, had been stuffing tax reporting forms in his drawer rather than sending them in to the IRS for years, which resulted in large penalties.

None of these things would have shown up in a graph of profit and loss from table games bets. None of these risks could have been reasonably measured before the fact.



Never confuse risk measurement with risk management. If you can measure it, you probably don't have to manage it.

Calculating risk

People often like to segregate calculated risk from other types of risk. *Calculated risk* covers situations in which you know the possible outcomes and have good estimates of their probabilities. Examples are the risk of rolling a seven while trying to make your point in craps (one chance in six) or the chance of rain tomorrow. The more general *risk* covers situations where you can't even specify all the possible outcomes, such as starting a war or embarking on a course of scientific research, and have no basis to estimate the probabilities of the outcomes you can foresee.

University of Chicago professor Frank Knight famously labelled the calculated risk as *risk* and the second, more general condition, as *uncertainty*. Risk management is about the uncertainty that remains after front-office risk takers – traders, portfolio managers, lending officers and others – make the calculations that are possible. If you can calculate a risk, you almost always want to minimise it, subject to constraints. For example, a portfolio manager may select a portfolio that minimises annual volatility subject to a constraint that the expected annual return be 8 per cent or better.

Planning and plunging

The quotations here about planning and results emphasise a few of the ideas that a risk manager should absorb:

- ✓ 'In preparing for battle I have always found that plans are useless, but planning is indispensable.' General Dwight D. Eisenhower
- ✓ 'Everybody's got plans. . . until they get hit.' Boxer Mike Tyson
- ✓ 'If you wait until the right time to have a child you'll die childless, and I think film-making is very much the same thing. You just have to take the plunge and just start shooting something even if it's bad.' Filmmaker James Cameron
- ✓ 'Plunge, don't plan.' Instruction for commandos
- ✓ 'Earlier theorists aimed to equip the conduct of war with principles, rules, or even

systems, and thus considered only factors that could be mathematically calculated (e.g., numerical superiority; supply; the base; interior lines). All these attempts are objectionable, however, because they aim at fixed values. In war everything is uncertain and variable, intertwined with psychological forces and effects, and the product of a continuous interaction of opposites.' General Carl von Clausewitz

- ✓ 'Plunge boldly into the thick of life, and seize it where you will, it is always interesting.' Philosopher Johann Wolfgang von Goethe

Careful planning is necessary, but don't count on anything ever going to plan, and recognise that success in anything requires risks.



Minimising risk isn't managing risk. This point is important because not many people know it beyond those with extensive day-to-day experience making significant financial decisions from a risk management – as opposed to a portfolio management – perspective.

Financial risk management is based on a different mathematical tradition than the one used in most economics and statistics. The conventional academic analysis of risk uses gambling games as models, and works only if the solution to the simplified game is a good approximation to the solution to the real-world decision. That works pretty well sometimes, and you don't need a risk manager to help you with it. But in other cases it leads to disastrous decisions, even when done properly and carefully. Risk management doesn't assume you know enough about possible outcomes and probabilities to treat decisions like actions in a casino game, and that you instead need to draw on concepts from information theory and other fields to improve your chances of long-term success.

I spare you most of the gory details of the calculations you use to manage risk – or at least segregate them in technical sections with clear warning

signs posted. You don't need to do the maths to understand the ideas. However, you do need to know that maths is an option. In other words, you need to understand that you can bring powerful mathematical tools to bear on incalculable uncertainty just as you can on calculated risk.

In my experience, people who are good at calculations tend to overanalyse the calculated risks and pretend that their models are an approximation to reality, which leads to disastrous risk management. People who aren't good at calculations tend to emphasise the unknown unknowns (in Donald Rumsfeld's famous phrase) – the deficiencies in the data, the un-modelled complexities of the situation and all kinds of other things that cause the calculated risks to be unreliable. This attitude is less problematic than the first, but is far from optimal. Risk managers provide a clear third voice, one that says, 'We may not be able to calculate enough of the risks to be useful, but we can calculate our actions. We may not be able to measure the risk, but we can manage it.'

Regenerating dinosaurs

The movie *Jurassic Park* does a great job of illustrating how risk management differs from conventional approaches to uncertainty. In the book, the point is even clearer. (Author Michael Crichton should be an honorary risk manager for the many insights peppered through his fiction. I consider him the most intellectually stimulating popular fiction writer of the 20th century. He was also an outstandingly successful director and producer for movies and television.) When investors in a park that brings extinct dinosaur species back to life get concerned about the risks of the venture, they demand a report from three experts: a palaeontologist (Sam Neill), a palaeobiologist (Laura Dern) and a 'mathematician with a deplorable excess of personality' (Jeff Goldblum).

A number of movie reviewers remarked on the implausibility of sending a mathematician, especially one calling himself a *chaotician*. But the palaeo-people can only calculate and analyse factors about dinosaurs; they have no particular training in risk and are unlikely to have the kind of life experiences that build risk wisdom. All they can do is double-check the calculations of the palaeo-experts who designed the park (which were probably double- and triple-checked already). Although some people tell you that an extra check is always prudent, I disagree. One person with clear responsibility for a decision is often more reliable than three people who all think someone else will catch any error.

The mathematician doesn't do the careful observation of the other two experts – the palaeontologist who scrutinises the pack dynamics of running gallimimus or the palaeobiologist who sticks her arms into triceratops excrement. However, he correctly predicts disaster, without knowing anything

about dinosaurs, genetics or park security. He understands that evolution is a powerful force powered by risk – far too powerful to be controlled by electric fences. (Evolution is also known as *natural selection of random variation*, and both random and variation are essential risk concepts.) He did not predict the specifics of disaster, only that the imperatives of life would easily win over the calculations of human experts.

Risk managers understand that risk is a powerful force that can be harnessed for great success or that can blast apart the best-laid schemes. Risk is not about laying better schemes; it's about making sure that risk is the wind in your sails, not the approaching hurricane that will swamp your boat. And generally speaking (although certainly not always), experts in specialised fields are bad at recognising risk. Experts usually get paid to take the risk out of decisions – or at least to reduce the risk by making things more predictable. Doing so is certainly worthwhile, but it never works perfectly, so you need risk managers as well. More importantly, experts often get paid to reduce the appearance of risk, not risk itself. And most important of all, reflexively taking the risk out of decisions eliminates opportunities as well as dangers.

Adding a little maths

As I say, you need no maths to understand this book. However, if you're willing to dip your toe into mathematical waters, you can get a deeper understanding of risk management more quickly. Feel free to skip this section if you're not interested in the maths at all.




Suppose someone offers you a proposal that has a 50 per cent chance of a +20 per cent return and a 50 per cent chance of a –18 per cent return. A standard approach in economics for analysing this choice begins by asking how much happier a 20 per cent increase in wealth would make you and how much unhappier an 18 per cent decrease in wealth would make you. Because the probabilities are equal, you take this gamble if the happiness increase from 20 per cent is greater than the happiness decrease from –18 per cent. With certain qualifications, this approach can be reasonable for front-office risk takers, and it's the usual approach in academic portfolio management (although economists prefer to speak about abstract *utility* rather than practical happiness). In this book, I refer to this approach as the portfolio management approach.

Most non-economists would find such a gamble too risky for 100 per cent of their wealth, but the risk gets more attractive if it can be repeated many times. With many repetitions, this gamble seems like being the casino – statistically certain to win in the long run due to a built-in edge.

The chart in Figure 1-2 shows a random simulation of 20 risk takers who repeat this bet 250 times, starting with initial wealth of 1. The solid black

curve shows the growth of wealth at the expected rate of 1 per cent per bet (maths alert: 50 per cent probability times 20 per cent plus 50 per cent probability times -18 per cent equals 10 per cent - 9 per cent = 1 per cent expected growth of wealth) and the 20 other lines show individual paths.

Figure 1-2:
Charting
growth in
wealth.



Most paths go quickly to near zero. A few soar up far beyond the expected one per cent rate for a while, but all eventually crash. If you run the simulation longer, all paths would become indistinguishable from zero. To a risk manager, this bet is terrible – one that leads to certain disaster. The more times you repeat it, the worse it gets, not the better. Your psychology, your risk appetite, has nothing to do with it. This bet is worse than just losing all your money quickly because the paths that soar attract imitators and cause all kinds of foolish overreactions.

The problem is simple. If you win half your bets, you lose money. If you win 20 per cent, you turn £1.00 into £1.20. If you then lose 18 per cent, your £1.20 falls to £0.984. (The order doesn't matter. If you first lose 18 per cent to turn £1.00 into £0.82, then a 20 per cent win turns £0.82 to the same £0.984.) Every pair of win and loss costs you 0.6 per cent of your wealth. In the long run, you're virtually certain to have nearly 50 per cent wins and losses, so you're virtually certain to wipe out your wealth.

How does the median 0.3 per cent loss per bet square with the expected 1 per cent return? It's absolutely true that your expected wealth increases 1 per cent each time you repeat this bet, but in the long run this fact results from a microscopic probability of winning an astronomical amount of money. You're virtually certain to be broke, but theoretically have enough chance of winning far more money than exists in the universe that your expected value is positive.

This example is oversimplified, of course. With real risks, you never know the exact probabilities and outcomes. You don't repeat them an infinite number of times, and the results are not independent of each other. You don't bet constant fractions of your wealth each time. I use the example only to make the point that you can ask two different questions about any risk:

- ✓ The *line risk taker*, the person making risk decisions, asks some version of, 'Will I be happier on average, or will the organisation be better off on average, if I take this specific bet once?'
- ✓ The risk manager asks, 'Will a long-term strategy of taking this kind of bet lead with average luck to exponential growth or to disaster?'

The answers to these two questions are independent. Some risks increase average utility if taken once but can't be accepted as part of a systematic strategy that leads to success, and some risks fit perfectly into systematic strategies but are unattractive as individual propositions. The only risks worth taking are the ones that make sense on their own and as steps in the long-term strategy. That's why you need both line risk takers to ensure the first, and risk managers to ensure the second.

I emphasize that this is a practical result discovered by experience, not a theoretical one. The mathematical example was invented to illustrate the idea; it's not the source of it. Quantitative risk managers learned that it was possible to analyse real risk-taking histories of real risk takers without assuming anything about probabilities or future possibilities or risk preferences and determine accurately whether they were on paths to riches or ruin. First they learned with their own risk taking, often from bitter experience, and then they learned it was possible to prove their contentions to risk takers, even when markets were in the peaks of success or the depths of slumps. This was the birth of the modern field of quantitative risk management.

Working with Financial Risk

The topic of this book is financial risk. Financial risk is created by people. It can represent natural risk: for example, an insurance company writing hurricane insurance or a venture capital investor taking on some of the economic

risk of a start-up company. But most financial risk is entirely contained within the financial system, such as a futures trader making zero-sum bets with other futures traders or a government bond portfolio manager speculating on changes in interest rates.

Even when financial risk represents physical risk, it represents the virtual version, not the real thing. An insurance company writes checks after a hurricane; it isn't pinned under a fallen tree without fresh water available. A venture capitalist writes off his investment if the company fails; he doesn't fire people and auction off the office furniture.

The idea of converting physical and economic risk to virtual form and trading it is revolutionary. It allows people to shed their excess concentrated risk, such as that their company will fail or their house will burn down, and to take on opportunities in a diversified portfolio of other people's risks. Done properly, the good risk – the innovation, the opportunity, the creation – is maximised for everyone's benefit, and the bad risk – the ruin, the disaster, the catastrophe – is diversified into broadly shared affordable losses. This diversification is not always done properly, unfortunately, but it sometimes is. And whether the finance is done well or badly, it deals with risk, and risk must be managed.

Managing financial risk

In managing financial risk, you need to distinguish between the risk of the financial product – the stuff that's bought, repackaged and sold – from the risk of running a financial business.



A printing company has a contract with the government to print money. On one day it prints a billion pounds worth of bills to send to the government, and earns £100,000 for the job. If you ask the CEO how much money the company made, the answer is £100,000, not one billion pounds. If the CEO forgets this distinction and starts spending the money his company prints for the government, he goes to jail.

The distinction between types of risk is easiest to see with the manager of an S&P 500 index fund. The manager doesn't make judgements about securities, he just promises to take investors' money and use it to buy the 500 stocks that make up the index. (Investing in an index is slightly more complicated than this, but that doesn't matter for this example.)

One risk, of course, is whether the S&P 500 basket of stocks goes up or down. However, this risk isn't to the index fund manager. He sells this risk to his investors. His investors want it. This is like the billion pounds of bills the printing company printed for the government.

The index fund management company has a risk manager. The risk manager doesn't spend time thinking about how risky the S&P 500 stocks are. That's not his job. He is, however, concerned with the liquidity of the S&P 500 stocks because the index fund needs to trade in order to honour new subscriptions and redemptions. He worries a lot about valuation because errors may result in underpayments or overpayments. He pays attention to counterparty risks, such as what happens if a dealer fails to honour a trade, a custodian goes suddenly bankrupt or a stock-lending counterparty is unable to return borrowed shares. A host of other risks are present as well. The point is that the risk manager's concern is that the management company does what it promises – deliver the risk of the S&P 500 to its investors – not whether the risk of the S&P 500 is a good or bad risk.

Investing with a mutual fund company that picks and chooses among stocks in an attempt to beat the S&P 500 is a bit more complicated. Now the company is selling a more complicated risk, a combination of S&P 500 risk plus the risk of the portfolio manager's outperformance or underperformance. The company's risk manager has all the concerns of the index company's risk manager, plus the risk that the portfolio manager's stock-picking skill isn't properly represented in the portfolio. This scenario can happen, for example, if the manager makes unintentionally concentrated bets, changes strategy from what the prospectus promises, engages in *chasing* (doubling bets to offset past losses rather than allocating funds in sober calculation about the future) or *window dressing* (making trades just before reporting dates so the portfolio looks good in the report) or manages with an eye toward gaining more assets rather than delivering the best possible performance to existing investors. But the risk manager's job ends with making sure that the fund delivers the manager's best efforts to beat the S&P 500 within the terms of regulation and the prospectus. Fund investors choose to be exposed to S&P 500 stocks and to the fund manager; the wisdom of this choice isn't the risk manager's concern.

With other financial businesses such as dealers, banks and insurance companies, the risk situation is even more complicated. Despite the complexity, you must keep separate the risk that is the company's product from the risks that the company incurs in buying and selling its product.

Check out Chapter 15 for a discussion of the range of risk.

Working in financial institutions

The modern practice of financial risk management was developed in the late 1980s and early 1990s. It sprang up on *prop trading desks*, places where traders make financial bets for the benefit of the firm as opposed to executing trades

on behalf of customers or for the convenience of customers. At that time, no one outside the prop trading desks knew or cared about these developments.

The 1990s saw the spread of modern financial risk management to all trading businesses of large financial institutions, creating what became known as a *middle office* between front-office risk takers and back-office support personnel.



The tumultuous financial events of the last 20 years (really no more tumultuous than any 20-year period but getting much wider attention) led to two strong ideas:

- ✓ Disasters occur when risk managers aren't sufficiently independent of risk takers.

This idea led to walling off of the middle office so that these risk managers report only to other middle-office risk managers up to the level of the chief risk officer (CRO). The CRO reports directly to the CEO and board. No one in the firm can bypass the chain of command to direct the actions of any middle-office risk manager without going through the CRO.

- ✓ Every stakeholder needs voluminous reports about every aspect of risk.

This requirement has led to the creation of gigantic back-office risk management organisations that dwarfed the size of middle-office and front-office risk management.

Crunching data in the back office

Most of the available risk jobs are in back-office risk management, compiling reports, building IT (information technology) systems, scrubbing data, checking limits, auditing results and similar functions. Generally, back-office risk people learn their financial risk management on the job. They're hired for programming or auditing or legal skills or for their general business information skills.

Although back-office jobs lack the pay and glamour of front-office positions, they tend to offer better quality of life, more stable careers and advancement based on doing the job well rather than politics or luck. Back-office risk reporting, for example, tends to be more interesting than other back-office jobs because it stitches together information from all parts of the organisation – everything affects risk. Moreover, risk is about reality rather than abstractions. In my experience, back-office risk offers more opportunity to move to middle- or front-office than other back-office jobs, but in most organisations that opportunity is limited even in the risk department.

Front-office risk management works directly with traders and portfolio managers. This is the best-developed part of risk management. The front office

is the place with the highest pay and most day-to-day excitement. It used to be the place that all risk managers got their start – you moved from trader or portfolio manager to front-office risk manager to middle-office risk manager. That's still the best career path, but is rare these days. Most front-office risk managers start in other front-office roles, and never leave the front office.

Making the most of the middle office

Middle-office risk management is where the overall risk policies and methodologies are set, where front-office risk decisions are aggregated and where back-office risk reports are analysed and interpreted for other departments. The middle office is smaller than either front-office or back-office risk groups. It is, however, my main focus in this book. The reason is that all stakeholders, and in particular all risk managers, have to understand risk from the perspective of the middle office.



The middle office is where everything gets put together and communicated to the world. The back office does all the work, and the front office takes all the risks, but the middle office is where the risk is managed.

Communicating Risk

I'm often asked what the most important job of a risk manager is. The answer is simple but unexpected to most people: The risk manager's most crucial task is to communicate a single vision of risk to all stakeholders: equity holders, creditors, customers, executives, regulators, employees, trading counterparties – everyone. It's nice if that single vision happens to be accurate, but unfortunately you can only do your best in that respect. What you can promise is that the vision is the same for everyone.



Note that I don't say that the risk manager convinces everyone of the same vision of risk. People will always disagree about what the risk is. What they should not disagree about is what vision of risk is driving firm decisions.

Suppose that an entrepreneur lays out a proposed project. A lot of people take a look, and most have no interest. But some people are optimistic enough to lend money for the venture. Others are even more optimistic and are willing to put money in for a share of any profits after the lenders are paid. Some people want to work in the project for salary, or for equity options. Some people want to sign up to be suppliers to the project, or customers of it. The government probably gets into the act with various regulatory and tax interests. These people disagree by necessity; otherwise they would all be vying for the same role.

How do you keep the process honest? That is, how do you prevent the entrepreneur from telling creditors that the project will be run for maximum safety of repayment, telling the equity buyers that the project will be run for maximum upside, and the government that it will be run for social benefit? How do you prevent the owner from promising the same money to employees, suppliers and customers?

If you knew exactly how the project would turn out, an accountant could audit the projected books to make sure that each dollar went to exactly one place. However, given that many future scenarios are possible, that solution isn't practical. Instead of an accountant, you need a risk manager to lay out the range of possible futures in a form that balances simplicity (so stakeholders can understand it) with detail (so it captures the important contingencies and decisions). Each stakeholder makes an informed decision to participate based on a consistent promise of how the project will be run.

Of course, the actual outcome of the project will differ from all the risk manager's projections, perhaps in crucial ways. Some stakeholders may prosper while others suffer. After the fact, it's impossible to say whether the outcomes were fair or not. However, as long as everyone had the same risk information going in and as long as the project was run consistent with the promises made, then the responsibility for any gain or loss rests with the stakeholders' choices, and is fair in that sense. (I talk about communication in Chapter 18.)

This isn't to say that communication is the only duty of a financial risk manager. You can do things to make risk taking more productive and successful. You can encourage good risk (innovation, opportunity, experimentation, creativity, attractive bets) and discourage bad risk (carelessness, recklessness, unnecessary danger, chasing unattractive bets). You can build a positive risk culture, and gain consensus behind a shrewd risk strategy. But consistent risk communication is job one.