

COPYRIGHT NOTICE:

**Alexander J. McNeil, Rüdiger Frey, & Paul Embrechts:
Quantitative Risk Management**

is published by Princeton University Press and copyrighted, © 2005, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

Follow links Class Use and other Permissions. For more information, send email to: permissions@pupress.princeton.edu

1

Risk in Perspective

In this chapter we provide a non-mathematical discussion of various issues that form the background to the rest of the book. In Section 1.1 we begin with the *nature of risk* itself and how risk relates to *randomness*; in the financial context (which includes insurance) we summarize the main kinds of risks encountered and explain what it means to *measure* and *manage* such risks.

A brief *history* of financial risk management, or at least some of the main ideas that are used in modern practice, is given in Section 1.2, including a summary of the process leading to the Basel Accords. Section 1.3 gives an idea of the new *regulatory framework* that is emerging in the financial and insurance industries.

In Section 1.4 we take a step back and attempt to address the fundamental question of why we might want to measure and manage risk at all. Finally, in Section 1.5, we turn explicitly to quantitative risk management (QRM) and set out our own views concerning the *nature of this discipline* and the challenge it poses. This section in particular should give more insight into why we have chosen to address the particular methodological topics in this book.

1.1 Risk

The Concise Oxford English Dictionary defines risk as “hazard, a chance of bad consequences, loss or exposure to mischance”. In a discussion with students taking a course on financial risk management, ingredients which typically enter are events, decisions, consequences and uncertainty. Mostly only the downside of risk is mentioned, rarely a possible upside, i.e. the potential for a gain. For financial risks, the subject of this book, we might arrive at a definition such as “any event or action that may adversely affect an organization’s ability to achieve its objectives and execute its strategies” or, alternatively, “the quantifiable likelihood of loss or less-than-expected returns”. But while these capture some of the elements of risk, no single one-sentence definition is entirely satisfactory in all contexts.

1.1.1 Risk and Randomness

Independently of any context, risk relates strongly to uncertainty, and hence to the notion of randomness. Randomness has eluded a clear, workable definition for many centuries; it was not until 1933 that the Russian mathematician A. N. Kolmogorov gave an axiomatic definition of randomness and probability (see Kolmogorov 1933). This definition and its accompanying theory, though not without their controversial

aspects, now provide the lingua franca for discourses on risk and uncertainty, such as this book.

In Kolmogorov's language a probabilistic model is described by a triplet (Ω, \mathcal{F}, P) . An element ω of Ω represents a realization of an experiment, in economics often referred to as a state of nature. The statement "the probability that an event A occurs" is denoted (and in Kolmogorov's axiomatic system defined) as $P(A)$, where A is an element of \mathcal{F} , the set of all events. P denotes the probability measure. For the less mathematically trained reader it suffices to accept that Kolmogorov's system translates our intuition about randomness into a concise, axiomatic language and clear rules.

Consider the following examples: an investor who holds stock in a particular company; an insurance company that has sold an insurance policy; an individual who decides to convert a fixed-rate mortgage into a variable one. All of these situations have something important in common: the investor holds today an asset with an uncertain future value. This is very clear in the case of the stock. For the insurance company, the policy sold may or may not be triggered by the underlying event covered. In the case of a mortgage, our decision today to enter into this refinancing agreement will change (for better or for worse) the future repayments. So randomness plays a crucial role in the valuation of current products held by the investor, the insurance company or the home owner.

To model these situations a mathematician would now define a one-period risky position (or simply risk) X to be a function on the probability space (Ω, \mathcal{F}, P) ; this function is called a *random variable*. We leave for the moment the range of X (i.e. its possible values) unspecified. Most of the modelling of a risky position X concerns its *distribution function* $F_X(x) = P(X \leq x)$, the probability that by the end of the period under consideration, the value of the risk X is less than or equal to a given number x . Several risky positions would then be denoted by a random vector (X_1, \dots, X_d) , also written in bold face as \mathbf{X} ; time can be introduced, leading to the notion of random (or so-called stochastic) processes, usually written (\mathbf{X}_t) . Throughout this book we will encounter many such processes, which serve as essential building blocks in the mathematical description of risk.

We therefore expect the reader to be at ease with basic notation, terminology and results from elementary *probability and statistics*, the branch of mathematics dealing with *stochastic* models and their application to the real world. The word "stochastic" is derived from the Greek "Stochazesthai", the art of guessing, or "Stochastikos", meaning skilled at aiming, "stochos" being a target. In discussing stochastic methods for risk management we hope to emphasize the skill aspect rather than the guesswork.

1.1.2 Financial Risk

In this book we discuss risk in the context of finance and insurance (although many of the tools introduced are applicable well beyond this context). We start by giving a brief overview of the main risk types encountered in the financial industry.

In banking, the best known type of risk is probably *market risk*, the risk of a change in the value of a financial position due to changes in the value of the underlying

components on which that position depends, such as stock and bond prices, exchange rates, commodity prices, etc. The next important category is *credit risk*, the risk of not receiving promised repayments on outstanding investments such as loans and bonds, because of the “default” of the borrower. A further risk category that has received a lot of recent attention is *operational risk*, the risk of losses resulting from inadequate or failed internal processes, people and systems, or from external events.

The boundaries of these three risk categories are not always clearly defined, nor do they form an exhaustive list of the full range of possible risks affecting a financial institution. There are notions of risk which surface in nearly all categories such as *liquidity* and *model risk*. The latter is the risk associated with using a mis-specified (inappropriate) model for measuring risk. Think, for instance, of using the Black–Scholes model for pricing an exotic option in circumstances where the basic Black–Scholes model assumptions on the underlying securities (such as the assumption of normally distributed returns) are violated. It may be argued that model risk is always present to some degree. Liquidity risk could be roughly defined as the risk stemming from the lack of marketability of an investment that cannot be bought or sold quickly enough to prevent or minimize a loss. Liquidity can be thought of as “oxygen for a healthy market”; we need it to survive but most of the time we are not aware of its presence. Its absence, however, is mostly recognized immediately, with often disastrous consequences.

The concepts, techniques and tools we will introduce in the following chapters mainly apply to the three basic categories of market, credit and operational risk. We should stress that the only viable way forward for a successful handling of financial risk consists of a *holistic* approach, i.e. an integrated approach taking all types of risk and their interactions into account. Whereas this is a clear goal, current models do not yet allow for a fully satisfactory platform.

As well as banks, the insurance industry has a long-standing relationship with risk. It is no coincidence that the Institute of Actuaries and the Faculty of Actuaries use the following definition of the actuarial profession.

Actuaries are respected professionals whose innovative approach to making business successful is matched by a responsibility to the public interest. Actuaries identify solutions to financial problems. They manage assets and liabilities by analysing past events, assessing the present risk involved and modelling what could happen in the future.

An additional risk category entering through insurance is *underwriting risk*, the risk inherent in insurance policies sold. Examples of risk factors that play a role here are changing patterns of natural catastrophes, changes in demographic tables underlying (long-dated) life products, or changing customer behaviour (such as prepayment patterns).

1.1.3 Measurement and Management

Much of this book is concerned with techniques for the measurement of risk, an activity which is part of the process of managing risk, as we attempt to clarify in this section.

Risk measurement. Suppose we hold a portfolio consisting of d underlying investments with respective weights w_1, \dots, w_d so that the change in value of the portfolio over a given holding period (the so-called P&L, or profit and loss) can be written as $X = \sum_{i=1}^d w_i X_i$, where X_i denotes the change in value of the i th investment. Measuring the risk of this portfolio essentially consists of determining its distribution function $F_X(x) = P(X \leq x)$, or functionals describing this distribution function such as its mean, variance or 99th percentile.

In order to achieve this, we need a properly calibrated *joint* model for the underlying random vector of investments (X_1, \dots, X_d) . We will consider this problem in more detail in Chapter 2. At this point it suffices to understand that risk measurement is essentially a statistical issue; based on historical observations and given a specific model, a statistical estimate of the distribution of the change in value of a position, or one of its functionals, is calculated. As we shall see later, and this is indeed a main theme throughout the book, this is by no means an easy task with a unique solution.

It should be clear from the outset that good risk measurement is a must. Increasingly, banking clients demand objective and detailed information on products bought and banks can face legal action when this information is found wanting. For any product sold, a proper quantification of the underlying risks needs to be explicitly made, allowing the client to decide whether or not the product on offer corresponds to his or her risk appetite.

Risk management. In a very general answer to the question of what risk management is about, Kloman (1990) writes that:

To many analysts, politicians, and academics it is the management of environmental and nuclear risks, those technology-generated macro-risks that appear to threaten our existence. To bankers and financial officers it is the sophisticated use of such techniques as currency hedging and interest-rate swaps. To insurance buyers or sellers it is coordination of insurable risks and the reduction of insurance costs. To hospital administrators it may mean “quality assurance”. To safety professionals it is reducing accidents and injuries. In summary, risk management is *a discipline for living with the possibility that future events may cause adverse effects.*

The last phrase in particular (the italics are ours) captures the general essence of risk management, although for a financial institution one can perhaps go further. A bank’s attitude to risk is not passive and defensive; a bank actively and willingly takes on risk, because it seeks a return and this does not come without risk. Indeed risk management can be seen as the core competence of an insurance company or a bank. By using its expertise, market position and capital structure, a financial institution can manage risks by repackaging them and transferring them to markets in customized ways.

Managing the risk is thus related to preserving the flow of profit and to techniques like *asset liability management* (ALM), which might be defined as managing a financial institution so as to earn an adequate return on funds invested, and to maintain

a comfortable surplus of assets beyond liabilities. In Section 1.4 we discuss these corporate finance issues in more depth from a shareholder's point of view.

1.2 A Brief History of Risk Management

In this section we treat the historical development of risk management by sketching some of the innovations and some of the events that have shaped modern risk management for the financial industry. We also describe the more recent development of regulation in that industry, which has to some extent been prompted by a number of recent disasters.

1.2.1 From Babylon to Wall Street

Although risk management has been described as “one of the most important innovations of the 20th century” by Steinherr (1998) and most of the story we tell is relatively modern, some concepts that are used in modern risk management, in particular derivatives, have been around for longer. In our discussion we stress the example of financial derivatives, as these brought the need for increased banking regulation very much to the fore.

The ancient world to the twentieth century. A derivative is a financial instrument derived from an underlying asset, such as an option, future or swap. For example, a European call option with strike K and maturity T gives the holder the right, but not the obligation, to obtain from the seller at maturity the underlying security for a price of K ; a European put option gives the holder the right to dispose of the underlying at a price K .

Dunbar (2000) interprets a passage in the Code of Hammurabi from Babylon of 1800 BC as being early evidence of the use of the option concept to provide financial cover in the event of crop failure. A very explicit mention of options appears in Amsterdam towards the end of the seventeenth century and is beautifully narrated by Joseph de la Vega in his 1688 *Confusión de Confusiones*, a discussion between a lawyer, a trader and a philosopher observing the activity on the Beurs of Amsterdam. Their discussion contains what we now recognize as European call and put options, and a description of their use for investment as well as for risk management, and even the notion of short selling. In an excellent recent translation (de la Vega 1966) we read:

If I may explain “opsies” [further, I would say that] through the payment of the premiums, one hands over values in order to safeguard one's stock or to obtain a profit. One uses them as sails for a happy voyage during a beneficent conjuncture and as an anchor of security in a storm.

After this, de la Vega continues with some explicit examples that would not be out of place in any modern finance course on the topic.

Financial derivatives in general, and options in particular, are not so new. Moreover, they appear here as instruments to manage risk, “anchors of security in a storm”, rather than the inventions of the capitalist devil, the “wild beasts of finance” (Steinherr 1998), that many now believe them to be.

Academic innovation in the twentieth century. While the use of risk-management ideas such as derivatives can be traced further back, it was not until the late twentieth century that a theory of valuation for derivatives was developed. This can be seen as perhaps the most important milestone in an age of academic developments in the general area of quantifying and managing financial risk.

Before the 1950s the desirability of an investment was mainly equated to its return. In his ground-breaking publication of 1952, Harry Markowitz laid the foundation of the theory of portfolio selection by mapping the desirability of an investment onto a risk–return diagram, where risk was measured using standard deviation (see Markowitz 1952, 1959). Through the notion of an *efficient frontier* the portfolio manager could optimize the return for a given risk level. The following decades saw an explosive growth in risk-management methodology, including such ideas as the Sharpe ratio, the Capital Asset Pricing Model (CAPM) and Arbitrage Pricing Theory (APT). Numerous extensions and refinements followed, which are now taught in any MBA course on finance.

The famous Black–Scholes–Merton formula for the price of a European call option appeared in 1973 (see Black and Scholes 1973). The importance of this formula was underscored in 1997, when the Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel was awarded to Robert Merton and Myron Scholes (Fisher Black had died some years earlier) “for a new method to determine the value of derivatives”.

Growth of markets in the twentieth century. The methodology developed for the rational pricing and hedging of financial derivatives changed finance. The *Wizards of Wall Street* (i.e. the mathematical specialists conversant in the new methodology) have had a significant impact on the development of financial markets over the last few decades. Not only did the new option-pricing formula work, it transformed the market. When the Chicago Options Exchange first opened in 1973, less than a thousand options were traded on the first day. By 1995, over a million options were changing hands each day with current nominal values outstanding in the derivatives markets in the tens of trillions. So great was the role played by the Black–Scholes–Merton formula in the growth of the new options market that, when the American stock-market crashed in 1978, the influential business magazine *Forbes* put the blame squarely onto that one formula. Scholes himself has said that it was not so much the formula that was to blame, but rather that market traders had not become sufficiently sophisticated in using it.

Along with academic innovation, technological developments (mainly on the information–technology (IT) side) also laid the foundations for an explosive growth in the volume of new risk-management and investment products. This development was further aided by worldwide deregulation in the 1980s. Important additional factors contributing to an increased demand for risk-management skills and products were the oil crises of the 1970s and the 1970 abolition of the Bretton–Woods system of fixed exchange rates. Both energy prices and foreign exchange risk became highly volatile risk factors and customers required products to hedge them. The

1933 Glass–Steagall Act—passed in the US in the aftermath of the 1929 Depression to prohibit commercial banks from underwriting insurance and most kinds of securities—indirectly paved the way for the emergence of investment banks, hungry for new business. Glass–Steagall was replaced in 1999 by the Financial Services Act, which repealed many of the former’s key provisions. Today many more companies are able to trade and use modern risk-management products.

Disasters of the 1990s. In January 1992, the president of the New York Federal Reserve, E. Gerald Corrigan, speaking at the Annual Mid-Winter Meeting of the New York State Bankers Association, said:

You had all better take a very, very hard look at off-balance-sheet activities. The growth and complexity of [these] activities and the nature of the credit settlement risk they entail should give us cause for concern. . . . I hope this sounds like a warning, because it is. Off-balance-sheet activities [i.e. derivatives] have a role, but they must be managed and controlled carefully and they must be understood by top management as well as by traders and rocket scientists.

Corrigan was referring to the growing volume of derivatives on banking books and the way they were accounted for.

Many of us recall the headline “Barings forced to cease trading” in the Financial Times on 26 February 1995. A loss of £700 million ruined the oldest merchant banking group in the UK (established in 1761). Besides numerous operational errors (violating every qualitative guideline in the risk-management handbook), the final straw leading to the downfall of Barings was a so-called straddle position on the Nikkei held by the bank’s Singapore-based trader Nick Leeson. A straddle is a short position in a call and a put with the same strike—such a position allows for a gain if the underlying (in this case the Nikkei index) does not move too far up or down. There is, however, considerable loss potential if the index moves down (or up) by a large amount, and this is precisely what happened when the Kobe earthquake occurred.

About three years later, on 17 September 1998, *The Observer* newspaper, referring to the downfall of Long-Term Capital Management (LTCM), summarized the mood of the times when it wrote:

last week, free market economy died. Twenty five years of intellectual bullying by the University of Chicago has come to a close.

The article continued:

the derivatives markets are a rarefied world. They are peopled with individuals with an extraordinary grasp of mathematics—“a strange collection of Greeks, misfits and rocket scientists” as one observer put it last week.

And referring to the Black–Scholes formula, the article asked:

is this really the key to future wealth? Win big, lose bigger.

There were other important cases which led to a widespread discussion of the need for increased regulation: the Herstatt Bank case in 1974, Metallgesellschaft in 1993 or Orange County in 1994. See Notes and Comments below for further reading on the above.

The main reason for the general public's mistrust of these modern tools of finance is their perceived triggering effect for crashes and bubbles. Derivatives have without doubt played a role in some spectacular cases and as a consequence are looked upon with a much more careful regulatory eye. However, they are by now so much part of Wall Street (or any financial institution) that serious risk management without these tools would be unthinkable.

Thus it is imperative that mathematicians take a serious interest in derivatives and the risks they generate. Who has not yet considered a prepayment option on a mortgage or a change from a fixed-interest-rate agreement to a variable one, or vice versa (a so-called swap)? Moreover, many life insurance products now have options embedded.

1.2.2 *The Road to Regulation*

There is no doubt that regulation goes back a long way, at least to the time of the Venetian banks and the early insurance enterprises sprouting in London's coffee shops in the eighteenth century. In those days one would rely to a large extent on self-regulation or local regulation, but rules were there. However, key developments leading to the present regulatory risk-management framework are very much a twentieth century story.

Much of the regulatory drive originated from the Basel Committee of Banking Supervision. This committee was established by the Central-Bank Governors of the Group of Ten (G-10) at the end of 1974. The Group of Ten is made up (oddly) of eleven industrial countries which consult and cooperate on economic, monetary and financial matters. The Basel Committee does not possess any formal supranational supervising authority, and hence its conclusions do not have legal force. Rather, it formulates broad supervisory standards and guidelines and recommends statements of best practice in the expectation that individual authorities will take steps to implement them through detailed arrangements—statutory or otherwise—which are best suited to their own national system. The summary below is brief. Interested readers can consult, for example, Crouhy, Galai and Mark (2001) for further details, and should also see Notes and Comments below.

The first Basel Accord. The first Basel Accord of 1988 on Banking Supervision (Basel I) took an important step towards an international minimum capital standard. Its main emphasis was on credit risk, by then clearly the most important source of risk in the banking industry. In hindsight, however, the first Basel Accord took an approach which was fairly coarse and measured risk in an insufficiently differentiated way. Also the treatment of derivatives was considered unsatisfactory.

The birth of VaR. In 1993 the G-30 (an influential international body consisting of senior representatives of the private and public sectors and academia) published a

seminal report addressing for the first time so-called off-balance-sheet products, like derivatives, in a systematic way. Around the same time, the banking industry clearly saw the need for a proper risk management of these new products. At JPMorgan, for instance, the famous Weatherstone 4.15 report asked for a one-day, one-page summary of the bank's market risk to be delivered to the chief executive officer (CEO) in the late afternoon (hence the "4.15"). Value-at-Risk (VaR) as a market risk measure was born and RiskMetrics set an industry-wide standard.

In a highly dynamic world with round-the-clock market activity, the need for instant market valuation of trading positions (known as *marking-to-market*) became a necessity. Moreover, in markets where so many positions (both long and short) were written on the same underlyings, managing risks based on simple aggregation of nominal positions became unsatisfactory. Banks pushed to be allowed to consider *netting* effects, i.e. the compensation of long versus short positions on the same underlying.

In 1996 the important Amendment to Basel I prescribed a so-called *standardized* model for market risk, but at the same time allowed the bigger (more sophisticated) banks to opt for an *internal*, VaR-based model (i.e. a model developed in house). Legal implementation was to be achieved by the year 2000. The coarseness problem for credit risk remained unresolved and banks continued to claim that they were not given enough incentives to diversify credit portfolios and that the regulatory capital rules currently in place were far too risk insensitive. Because of overcharging on the regulatory capital side of certain credit positions, banks started shifting business away from certain market segments that they perceived as offering a less attractive risk–return profile.

The second Basel Accord. By 2001 a consultative process for a new Basel Accord (Basel II) had been initiated; this process is being concluded as this book goes to press. The main theme is credit risk, where the aim is that banks can use a finer, more risk-sensitive approach to assessing the risk of their credit portfolios. Banks opting for a more advanced, so-called *internal-ratings-based* approach are allowed to use internal and/or external credit-rating systems wherever appropriate. The second important theme of Basel II is the consideration of operational risk as a new risk class.

Current discussions imply an implementation date of 2007, but there remains an ongoing debate on specific details. Industry is participating in several Quantitative Impact Studies in order to gauge the risk-capital consequences of the new accord. In Section 1.3.1 we will come back to some issues concerning this accord.

Parallel developments in insurance regulation. It should be stressed that most of the above regulatory changes concern the banking world. We are also witnessing increasing regulatory pressure on the insurance side, coupled with a drive to combine the two regulatory frameworks, either institutionally or methodologically. As an example, the Joint Forum on Financial Conglomerates (Joint Forum) was established in early 1996 under the aegis of the Basel Committee on Banking Supervision, the International Organization of Securities Commissions (IOSCO) and the

International Association of Insurance Supervisors (IAIS) to take forward the work of the so-called Tripartite Group, whose report was released in July 1995. The Joint Forum is comprised of an equal number of senior bank, insurance and securities supervisors representing each supervisory constituency.

The process is underway in many countries. For instance, in the UK the Financial Services Authority (FSA) is stepping up its supervision across a wide range of financial and insurance businesses. The same is happening in the US under the guidance of the Securities and Exchange Commission (SEC) and the Fed. In Switzerland, discussions are underway between the Bundesamt für Privatversicherungen (BPV) and the Eidgenössische Bankenkommision (EBK) concerning a joint supervisory office. In Section 1.3.2 we will discuss some of the current, insurance-related solvency issues.

1.3 The New Regulatory Framework

This section is intended to describe in more detail the framework that has emerged from the Basel II discussions and the parallel developments in the insurance world.

1.3.1 *Basel II*

On 26 June 2004 the G-10 central-bank governors and heads of supervision endorsed the publication of the revised capital framework. The following statement is taken from this release.

The Basel II Framework sets out the details for adopting more risk-sensitive minimum capital requirements [Pillar 1] for banking organizations. The new framework reinforces these risk-sensitive requirements by laying out principles for banks to assess the adequacy of their capital and for supervisors to review such assessments to ensure banks have adequate capital to support their risks [Pillar 2]. It also seeks to strengthen market discipline by enhancing transparency in banks' financial reporting [Pillar 3]. The text that has been released today reflects the results of extensive consultations with supervisors and bankers worldwide. It will serve as the basis for national rule-making and approval processes to continue and for banking organizations to complete their preparations for the new Framework's implementation.

The three-pillar concept. As is apparent from the above quote, a key conceptual change within the Basel II framework is the introduction of the *three-pillar concept*. Through this concept, the Basel Committee aims to achieve a more holistic approach to risk management that focuses on the interaction between the different risk categories; at the same time the three-pillar concept clearly signals the existing difference between quantifiable and non-quantifiable risks.

Under *Pillar 1* banks are required to calculate a *minimum capital charge*, referred to as regulatory capital, with the aim of bringing the quantification of this minimal capital more in line with the banks' economic loss potential. Under the Basel II framework there will be a capital charge for credit risk, market risk and, for the first

time, operational risk. Whereas the treatment of market risk is unchanged relative to the 1996 Amendment of the Basel I Capital Accord, the capital charge for credit risk has been revised substantially. In computing the capital charge for credit risk and operational risk banks may choose between three approaches of increasing risk sensitivity and complexity; some details are discussed below.

It is further recognized that any quantitative approach to risk management should be embedded in a well-functioning corporate governance structure. Thus best-practice risk management imposes clear constraints on the organization of the institution, i.e. the board of directors, management, employees, internal and external audit processes. In particular, the board of directors assumes the ultimate responsibility for oversight of the risk landscape and the formulation of the company's risk appetite. This is where *Pillar 2* enters. Through this important pillar, also referred to as the *supervisory review process*, local regulators review the various checks and balances put into place. This pillar recognizes the necessity of an effective overview of the banks' internal assessments of their overall risk and ensures that management is exercising sound judgement and has set aside adequate capital for the various risks.

Finally, in order to fulfil its promise that increased regulation will also diminish systemic risk, clear reporting guidelines on risks carried by financial institutions are called for. *Pillar 3* seeks to establish *market discipline* through a better public disclosure of risk measures and other information relevant to risk management. In particular, banks will have to offer greater insight into the adequacy of their capitalization.

The capital charge for market risk. As discussed in Section 1.2.2, in the aftermath of the Basel I proposals in the early 1990s, there was a general interest in improving the measurement of market risk, particularly where derivative products were concerned. This was addressed in detail in the 1996 Amendment to Basel I, which prescribed standardized market risk models but also allowed more sophisticated banks to opt for *internal VaR* models. In Chapter 2 we shall give a detailed discussion of the calculation of VaR. For the moment it suffices to know that, for instance, a 10-day VaR at 99% of \$20 million means that our market portfolio will incur a loss of \$20 million *or more* with probability 1% by the end of a 10-day holding period, if the composition remains fixed over this period. The choice of the holding period (10 days) and the confidence level (99%) lies in the hands of the regulators when VaR is used for the calculation of regulatory capital. As a consequence of these regulations, we have witnessed a quantum leap in the prominence of quantitative risk modelling throughout all echelons of financial institutions.

Credit risk from Basel I to II. In a banking context, by far the oldest risk type to be regulated is credit risk. As mentioned in Section 1.2.2, Basel I handled this type of risk in a rather coarse way. Under Basel I and II the credit risk of a portfolio is assessed as the sum of *risk-weighted assets*, that is the sum of notional exposures weighted by a coefficient reflecting the creditworthiness of the counterparty (the risk weight). In Basel I, creditworthiness is split into three crude categories: governments,

regulated banks and others. For instance, under Basel I, the risk-capital charge for a loan to a corporate borrower is five times higher than for a loan to an OECD bank. Also, the risk weight for all corporate borrowers is identical, independent of their credit-rating category.

Due to its coarseness, the implementation of Basel I is extremely simple. But with the establishment of more detailed credit risk databases, the improvement of analytic models, and the rapid growth in the market for credit derivatives, banks have pressed regulators to come up with more risk-specific capital-adequacy guidelines. This is the main content of the new Basel II proposals, where banks will be allowed to choose between *standardized* approaches or more advanced *internal-ratings-based* (IRB) approaches for handling credit risk. The final choice will, however, also depend on the size and complexity of the bank, with the larger, international banks having to go for the more advanced models.

Already the banks opting for the standardized approach can differentiate better among the various credit risks in their portfolio, since under the Basel II framework the risk sensitivity of the available risk weights has been increased substantially. Under the more advanced IRB approach, a bank's *internal* assessment of the riskiness of a credit exposure is used as an input to the risk-capital calculation. The overall capital charge is then computed by aggregating the internal inputs using formulas specified by the Basel Committee. While this allows for increased risk sensitivity in the IRB capital charge compared with the standardized approach, portfolio and diversification effects are not taken into account; this would require the use of fully internal models as in the market risk case. This issue is currently being debated in the risk community, and it is widely expected that in the longer term a revised version of the Basel II Capital Accord allowing for the use of fully internal models will come into effect. In Chapter 8, certain aspects of the regulatory treatment of credit risk will be discussed in more detail.

Opening the door to operational risk. A basic premise for Basel II was that, whereas the new regulatory framework would enable banks to reduce their credit risk capital charge through internal credit risk models, the overall size of regulatory capital throughout the industry should stay unchanged under the new rules. This opened the door for the new risk category of operational risk, which we discuss in more depth in Section 10.1. Recall that Basel II defines operational risk as the risk of losses resulting from inadequate or failed internal processes, people and systems or from external events. The introduction of this new risk class has led to heated discussions among the various stakeholders. Whereas everyone agrees that risks like human risk (e.g. incompetence, fraud), process risk (e.g. model, transaction and operational control risk) and technology risk (e.g. system failure, programming error) are important, much disagreement exists on how far one should (or can) go towards quantifying such risks. This becomes particularly difficult when the financially more important risks like fraud and litigation are taken into account. Nobody doubts the importance of operational risk for the financial and insurance sector, but much less agreement exists on how to measure this risk.

The Cooke ratio. A crude measure of capitalization is the well-known *Cooke ratio*, which specifies that capital should be at least 8% of the risk-weighted assets of a company. The precise definition of risk capital is rather complex, involving various tiers of differing liquidity and legal character, and is very much related to existing accounting standards. For more detail see, for example, Crouhy, Galai and Mark (2001).

Some criticism. The benefits arising from the regulation of financial services are not generally in doubt. Customer-protection acts, basic corporate governance, clear guidelines on fair and comparable accounting rules, the ongoing pressure for transparent customer and shareholder information on solvency, capital- and risk-management issues are all positive developments. Despite these positive points, the specific proposals of Basel II have also elicited criticism; issues that have been raised include the following.

- The *cost factor* of setting up a well-functioning risk-management system compliant with the present regulatory framework is significant, especially (in relative terms) for smaller institutions.
- So-called *risk-management herding* can take place, whereby institutions following similar (perhaps VaR-based) rules may all be running for the same exit in times of crises, consequently destabilizing an already precarious situation even further. This herding phenomenon has been suggested in connection with the 1987 crash and the events surrounding the 1998 LTCM crisis. On a related note, the *procyclical* effects of financial regulation, whereby capital requirements may rise in times of recession and fall in times of expansion, may contribute negatively to the availability of liquidity in moments where the latter is most needed.
- Regulation could lead to *overconfidence* in the quality of statistical risk measures and tools.

Several critical discussions have taken place questioning to what extent the crocodile of regulatory risk management is eating its own tail. In an article of 12 June 1999, the *Economist* wrote that “attempts to measure and put a price on risk in financial markets may actually be making them riskier”; on the first page of the article, entitled “The price of uncertainty”, the proverbial crocodile appeared. The reader should be aware that there are several aspects to the overall regulatory side of risk management which warrant further discussion. As so often, “the truth” of what constitutes good and proper supervision will no doubt be somewhere between the more extreme views. The Basel process has the very laudable aspect that constructive criticism is taken seriously.

1.3.2 Solvency 2

In this section we take a brief look at regulatory developments regarding risk management in the insurance sector. We concentrate on the current solvency discussion, also referred to as *Solvency 2*. The following statement, made by the EU Insurance

Solvency Sub-Committee (2001), focuses on the differences between the Basel II and Solvency 2 frameworks.

The difference between the two prudential regimes goes further in that their actual objectives differ. The prudential objective of the Basel Accord is to reinforce the soundness and stability of the international banking system. To that end, the initial Basel Accord and the draft New Accord are directed primarily at banks that are internationally active. The draft New Accord attaches particular importance to the self-regulating mechanisms of a market where practitioners are dependent on one another. In the insurance sector, the purpose of prudential supervision is to protect policyholders against the risk of (isolated) bankruptcy facing every insurance company. The systematic risk, assuming that it exists in the insurance sector, has not been deemed to be of sufficient concern to warrant minimum harmonisation of prudential supervisory regimes at international level; nor has it been the driving force behind European harmonisation in this field.

More so than in the case of banking regulation, the regulatory framework for insurance companies has a strong local flavour where many local statutory rules prevail. The various solvency committees in EU member countries and beyond are trying to come up with some global principles which would be binding on a larger geographical scale. We discuss some of the more recent developments below.

From Solvency 1 to 2. The first EU non-life and life directives on solvency margins appeared around 1970. The latter was defined as an extra capital buffer against unforeseen events such as higher than expected claims levels or unfavourable investment results. In 1997, the Müller report appeared under the heading “Solvency of insurance undertakings”—this led to a review of the solvency rules and initiated the Solvency 1 project, which was completed in 2002 and came into force in 2004. Meanwhile, Solvency 2 was initiated in 2001 with the publication of the influential Sharma report—the detailed technical rules of Solvency 2 are currently being worked out.

Solvency 1 was a rather coarse framework calling for a minimum guarantee fund (minimal capital required) of €3 million, and a solvency margin consisting of 16–18% of non-life premiums together with 4% of the technical provisions for life. This led to a single, robust system which is easy to understand and inexpensive to monitor. However, on the negative side, it is mainly volume based and not explicitly risk based; issues like guarantees, embedded options and proper matching of assets and liabilities were largely neglected in many countries. These and further shortcomings will be addressed in Solvency 2.

At the heart of Solvency 2 lies a risk-oriented assessment of overall solvency, honouring the three-pillar concept from Basel II (see the previous section). Insurers are encouraged to measure and manage their risks based on internal models. Consistency between Solvency 2 (Insurance) and Basel II (Banking) is adhered to as much as possible. The new framework should allow for an efficient supervision of

insurance groups (holdings) and financial conglomerates (bank-assurance). From the start, an increased harmonization of supervisory methodology between the different legislative entities was envisaged, based on a wide international cooperation with actuarial, financial and accounting bodies.

Without entering into the specifics of the framework, the following points related to Pillar 1 should be mentioned. In principle, *all* risks are to be analysed including underwriting, credit, market, operational (corresponding to internal operational risk under Basel II), liquidity and event risk (corresponding to external operational risk under Basel II). Strong emphasis is put on the modelling of interdependencies and a detailed analysis of stress tests. The system should be as much as possible *principle based* rather than *rules based* and should lead to *prudent regulation* which focuses on the total balance sheet, handling assets and liabilities in a single common framework.

The final decision on solvency is based on a two-tier procedure. This involves setting a first safety barrier at the level of the so-called *target capital* based on risk-sensitive, market-consistent valuation; breaches of this early-warning level would trigger regulatory intervention. The second and final tier is the *minimal capital level* calculated with the old Solvency 1 rules. It is interesting to note that in the definition of target capital, the *expected shortfall* for a holding period is used as a risk measure rather than Value-at-Risk, reflecting actuaries' experience with skewed and heavy-tailed pay-off functions; this alternative risk measure will be defined in Section 2.2.4. The reader interested in finding out more about the ongoing developments in insurance regulation will find relevant references in Notes and Comments.

1.4 Why Manage Financial Risk?

An important issue that we have barely dealt with concerns the reasons why we should invest in QRM in the first place. This question can be posed from various perspectives, including those of the customer of a financial institution, its shareholders, management, board of directors, regulators, politicians, or the public at large. Each of these stakeholders may have a different answer, and, at the end of the day, an equilibrium between the various interests will have to be found. In this section, we will focus on some of the players involved and give a selective account of some of the issues. It is not our aim, nor do we have the competence, to give a full treatment of this important subject.

1.4.1 A Societal View

Modern society relies on the smooth functioning of banking and insurance systems and has a collective interest in the stability of such systems. The regulatory process culminating in Basel II has been strongly motivated by the fear of systemic risk, i.e. the danger that problems in a single financial institution may spill over and, in extreme situations, disrupt the normal functioning of the entire financial system. Consider the following remarks made by Alan Greenspan before the Council on Foreign Relations in Washington, DC, on 19 November 2002 (Greenspan 2002).

Today, I would like to share with you some of the evolving international financial issues that have so engaged us at the Federal Reserve over the

past year. I, particularly, have been focusing on innovations in the management of risk and some of the implications of those innovations for our global economic and financial system. . . . The development of our paradigms for containing risk has emphasized dispersion of risk to those willing, and presumably able, to bear it. If risk is properly dispersed, shocks to the overall economic systems will be better absorbed and less likely to create cascading failures that could threaten financial stability.

In the face of such spillover scenarios, society views risk management positively and entrusts regulators with the task of forging the framework that will safeguard its interests. Consider the debate surrounding the use and misuse of derivatives. Regulation serves to reduce the risk of the misuse of these products, but at the same time recognizes their societal value in the global financial system. Perhaps contrary to the popular view, derivatives should be seen as instruments that serve to enhance stability of the system rather than undermine it, as argued by Greenspan in the same address.

Financial derivatives, more generally, have grown at a phenomenal pace over the past fifteen years. Conceptual advances in pricing options and other complex financial products, along with improvements in computer and telecommunications technologies, have significantly lowered the costs of, and expanded the opportunities for, hedging risks that were not readily deflected in earlier decades. Moreover, the counterparty credit risk associated with the use of derivative instruments has been mitigated by legally enforceable netting and through the growing use of collateral agreements. These increasingly complex financial instruments have especially contributed, particularly over the past couple of stressful years, to the development of a far more flexible, efficient, and resilient financial system than existed just a quarter-century ago.

1.4.2 *The Shareholder's View*

It is widely believed that proper financial risk management can increase the value of a corporation and hence shareholder value. In fact, this is the main reason why corporations which are not subject to regulation by financial supervisory authorities engage in risk-management activities. Understanding the relationship between shareholder value and financial risk management also has important implications for the design of risk-management (RM) systems. Questions to be answered include the following.

- When does RM increase the value of a firm, and which risks should be managed?
- How should RM concerns factor into investment policy and capital budgeting?

There is a rather extensive corporate finance literature on the issue of “corporate risk management and shareholder value”. We briefly discuss some of the main arguments. In this way we hope to alert the reader to the fact that there is more to RM than

the mainly technical questions related to the implementation of RM strategies dealt with in the core of this book.

The first thing to note is that from a corporate-finance perspective it is by no means obvious that in a world with perfect capital markets RM enhances shareholder value: while *individual* investors are typically risk averse and should therefore manage the risk in their portfolios, it is not clear that RM or risk reduction at the *corporate level*, such as hedging a foreign-currency exposure or holding a certain amount of risk capital, increases the value of a corporation. The rationale for this—at first surprising—observation is simple: if investors have access to perfect capital markets, they can do the RM transactions they deem necessary via their own trading and diversification. The following statement from the chief investment officer of an insurance company exemplifies this line of reasoning: “If our shareholders believe that our investment portfolio is too risky, they should short futures on major stock market indices”.

The potential irrelevance of corporate RM for the value of a corporation is an immediate consequence of the famous *Modigliani–Miller Theorem* (Modigliani and Miller 1958). This result, which marks the beginning of modern corporate finance theory, states that, in an ideal world without taxes, bankruptcy costs and informational asymmetries, and with frictionless and arbitrage-free capital markets, the financial structure of a firm—and hence also its RM decisions—are irrelevant for the firm’s value. Hence, in order to find reasons for corporate RM, one has to “turn the Modigliani–Miller Theorem upside down” and identify situations where RM enhances the value of a firm by deviating from the unrealistically strong assumptions of the theorem. This leads to the following rationales for RM.

- RM can reduce *tax costs*. Under a typical tax regime the amount of tax to be paid by a corporation is a *convex* function of its profits; by reducing the variability in a firm’s cash flow, RM can therefore lead to a higher expected after-tax profit.
- RM can be beneficial, since a company may (and usually will) have better access to capital markets than individual investors.
- RM can increase the firm value in the presence of *bankruptcy costs*, as it makes bankruptcy less likely.
- RM can reduce the impact of *costly external financing* on the firm value, as it facilitates the achievement of optimal investment.

The last two points merit a more detailed discussion. Bankruptcy costs consist of direct bankruptcy costs, such as the cost of lawsuits, and the more important indirect bankruptcy costs. The latter may include liquidation costs, which can be substantial in the case of intangibles like research and development (R&D) and know-how. This is why high R&D spending appears to be positively related to the use of RM techniques. Moreover, increased likelihood of bankruptcy often has a negative effect on key employees, management and customer relations, in particular in areas where a client wants a long-term business relationship. For instance, few customers

would want to enter into a life insurance contract with an insurance company which is known to be close to bankruptcy. On a related note, banks which are close to bankruptcy might be faced with the unpalatable prospect of a bank run, where depositors try to withdraw their money simultaneously. A further discussion of these issues is given in Altman (1993).

It is a “stylized fact of corporate finance” that for a corporation external funds are more costly to obtain than internal funds, an observation which is usually attributed to problems of asymmetric information between the management of a corporation and bond and equity investors. For instance, raising external capital from outsiders by issuing new shares might be costly if the new investors, who have incomplete information about the economic prospects of a firm, interpret the share issue as a sign that the firm is overvalued. This can generate a rationale for RM for the following reason: without RM the increased variability of a company’s cash flow will be translated either into an increased variability of the funds which need to be raised externally or to an increased variability in the amount of investment. With increasing marginal costs of raising external capital and decreasing marginal profits from new investment, this leads to a decrease in (expected) profits. Hence proper RM, which amounts to a smoothing of the cash flow generated by a corporation, can be beneficial. For references to the literature see Notes and Comments.

1.4.3 Economic Capital

As we have just seen, a corporation typically has strong incentives to strictly limit the probability of bankruptcy in order to avoid the associated bankruptcy costs. This is directly linked to the notion of economic capital. In a narrow sense, economic capital is the capital that shareholders should invest in the company in order to limit the probability of default to a given confidence level over a given time horizon. More broadly, economic capital offers a firm-wide language for discussing and pricing risk that is related directly to the principal concerns of management and other key stakeholders, namely institutional solvency and profitability (see Matten 2000). In this broader sense, economic capital represents the emerging best practice for measuring and reporting all kinds of risk across a financial organization.

Economic capital is so called because it measures risk in terms of *economic* realities rather than potentially misleading regulatory or accounting rules; moreover, part of the measurement process involves converting a risk distribution into the amount of *capital* that is required to support the risk, in line with the institution’s target financial strength (e.g. credit rating). Hence the calculation of economic capital is a process that begins with the quantification of the risks that any given company faces over a given time period. These risks include those that are well defined from a regulatory point of view, such as credit, market and operational risks, and also includes other categories like insurance, liquidity, reputational and strategic or business risk. When modelled in detail and aggregated one obtains a value distribution in line with the Merton model for firm valuation as discussed in Chapter 8.

Given such a value distribution, the next step involves the determination of the probability of default (solvency standard) that is acceptable to the institution. The mapping from risk (solvency standard) to capital often uses standard external benchmarks for credit risk. For instance, a firm that capitalizes to Moody's Aa standard over a one-year horizon determines its economic capital as the "cushion" required to keep the firm solvent over a one-year period with 99.97% probability; firms rated Aa by Moody's have historically defaulted with a 0.03% frequency over a one-year horizon (see, for example, Duffie and Singleton 2003, Table 4.2). The choice of horizon must relate to natural capital planning or business cycles, which might mean one year for a bank but typically longer for an insurance company. In the ideal RM set-up, it is economic capital that is used for setting risk limits. Or, as stated in (Drzik, Nakada and Schuermann 1998), economic capital can serve as a common currency for risk limits. That paper also discusses the way in which economic capital (capital you need) can be compared with physical capital (capital you have) and how corporate-finance decisions can be based on this comparison.

We hope that our brief discussion of the economic issues surrounding modern RM has convinced the reader that there is more to RM than the mere statistical computation of risk measures, important though the latter may be. The Notes and Comments provide some references for readers who want to learn more about the economic foundations of RM.

1.5 Quantitative Risk Management

In this first chapter we have tried to place QRM in a larger historical, institutional, and even societal framework, since a study of QRM without a discussion of its proper setting and motivation makes little sense. In the remainder of the book we adopt a somewhat narrower view and treat QRM as a quantitative science using the language of mathematics in general, and probability and statistics in particular.

In this section we describe the challenge that we have attempted to meet in this book and discuss where QRM may lead in the future.

1.5.1 The Nature of the Challenge

We set ourselves the task of defining a new discipline of QRM and our approach to this task has two main strands. On the one hand, we have attempted to put current practice onto a firmer mathematical footing where, for example, concepts like profit-and-loss distributions, risk factors, risk measures, capital allocation and risk aggregation are given formal definitions and a consistent notation. In doing this we have been guided by the consideration of what topics should form the core of a course on QRM for a wide audience of students interested in RM issues; nonetheless, the list is far from complete and will continue to evolve as the discipline matures. On the other hand, the second strand of our endeavour has been to put together material on techniques and tools which go beyond current practice and address some of the deficiencies that have been raised repeatedly by critics. In the following paragraphs we elaborate on some of these issues.

Extremes matter. A very important challenge in QRM, and one that makes it particularly interesting as a field for probability and statistics, is the need to address unexpected, abnormal or extreme outcomes, rather than the expected, normal or average outcomes that are the focus of many classical applications. This is in tune with the regulatory view expressed by Alan Greenspan:

From the point of view of the risk manager, inappropriate use of the normal distribution can lead to an understatement of risk, which must be balanced against the significant advantage of simplification. From the central bank's corner, the consequences are even more serious because we often need to concentrate on the left tail of the distribution in formulating lender-of-last-resort policies. Improving the characterization of the distribution of extreme values is of paramount importance.

Joint Central Bank Research Conference, 1995

The need for a response to this challenge became very clear in the wake of the LTCM case in 1998. John Meriwether, the founder of the hedge fund, clearly learned from this experience of extreme financial turbulence; he is quoted as saying:

With globalisation increasing, you'll see more crises. Our whole focus is on the extremes now—what's the worst that can happen to you in any situation—because we never want to go through that again.

The Wall Street Journal, 21 August 2000

Much space is devoted in our book to models for financial risk factors that go beyond the normal (or Gaussian) model and attempt to capture the related phenomena of heavy tails, volatility and extreme values.

The interdependence and concentration of risks. A further important challenge is presented by the multivariate nature of risk. Whether we look at market risk or credit risk, or overall enterprise-wide risk, we are generally interested in some form of aggregate risk that depends on high-dimensional vectors of underlying risk factors such as individual asset values in market risk, or credit spreads and counterparty default indicators in credit risk.

A particular concern in our multivariate modelling is the phenomenon of dependence between extreme outcomes, when many risk factors move against us simultaneously. Again in connection with the LTCM case we find the following quote in *Business Week*, September 1998.

Extreme, synchronized rises and falls in financial markets occur infrequently but they do occur. The problem with the models is that they did not assign a high enough chance of occurrence to the scenario in which many things go wrong at the same time—the “perfect storm” scenario.

In a perfect storm scenario the risk manager discovers that the diversification he thought he had is illusory; practitioners describe this also as a concentration of risk.

Myron Scholes, a prominent figure in the development of RM, alludes to this in Scholes (2000), where he argues against the regulatory overemphasis of VaR in the face of the more important issue of co-movements in times of market stress:

Over the last number of years, regulators have encouraged financial entities to use portfolio theory to produce dynamic measures of risk. VaR, the product of portfolio theory, is used for short-run, day-to-day profit-and-loss exposures. Now is the time to encourage the BIS and other regulatory bodies to support studies on stress test and concentration methodologies. Planning for crises is more important than VaR analysis. And such new methodologies are the correct response to recent crises in the financial industry.

The problem of scale. A further challenge in QRM is the typical scale of the portfolios under consideration; in the most general case a portfolio may represent the entire position in risky assets of a financial institution. Calibration of detailed multivariate models for all risk factors is a well-nigh impossible task and hence any sensible strategy involves dimension reduction, that is to say the identification of key risk drivers and a concentration on modelling the main features of the overall risk landscape.

In short we are forced to adopt a fairly “broad-brush” approach. Where we use econometric tools, such as models for financial return series, we are content with relatively simple descriptions of individual series which capture the main phenomenon of volatility, and which can be used in a parsimonious multivariate factor model. Similarly, in the context of portfolio credit risk, we are more concerned with finding suitable models for the default dependence of counterparties than with accurately describing the mechanism for the default of an individual, since it is our belief that the former is at least as important as the latter in determining the risk of a large diversified portfolio.

Interdisciplinarity. Another aspect of the challenge of QRM is the fact that ideas and techniques from several existing quantitative disciplines are drawn together. When one considers the ideal education for a quantitative risk manager of the future, then no doubt a combined quantitative skillset should include concepts, techniques and tools from such fields as mathematical finance, statistics, financial econometrics, financial economics and actuarial mathematics. Our choice of topics is strongly guided by a firm belief that the inclusion of modern statistical and econometric techniques and a well-chosen subset of actuarial methodology are essential for the establishment of best-practice QRM. Certainly QRM is not just about financial mathematics and derivative pricing, important though these may be.

Of course, the quantitative risk manager operates in an environment where additional non-quantitative skills are equally important. Communication is certainly the most important skill of all, as a risk professional by definition of his/her duties will have to interact with colleagues with diverse training and background at all levels of the organization. Moreover, a quantitative risk manager has to familiarize him or herself quickly with all-important market practice and institutional details. Finally, a

certain degree of humility will also be required to recognize the role of *quantitative* risk management in a much larger picture.

1.5.2 QRM for the Future

It cannot be denied that the use of QRM in the insurance and banking industry has had an overall positive impact on the development of those industries. However, RM technology is not restricted to the financial-services industry and similar developments are taking place in other sectors of industry. Some of the earliest applications of QRM are to be found in the manufacturing industry, where similar concepts and tools exist under names like reliability or total quality control. Industrial companies have long recognized the risks associated with bringing faulty products to the market. The car manufacturing industry in Japan in particular has been an early driving force in this respect.

More recently, QRM techniques have been adopted in the transport and energy industries, to name but two. In the case of energy there are obvious similarities with financial markets: electrical power is traded on energy exchanges; derivatives contracts are used to hedge future price uncertainty; companies optimize investment portfolios combining energy products with financial products; a current debate in the industry concerns the extent to which existing Basel II methodology can be transferred to the energy sector. However, there are also important dissimilarities due to the specific nature of the industry; most importantly there is the issue of the cost of storage and transport of electricity as an underlying commodity and the necessity of modelling physical networks including the constraints imposed by the existence of national boundaries and quasi-monopolies.

A further exciting area concerns the establishment of markets for environmental emission allowances. For example, the Chicago Climate Futures Exchange (CCFE) currently offers futures contracts on sulphur dioxide emissions. These are traded by industrial companies producing the pollutant in their manufacturing process and force such companies to consider the cost of pollution as a further risk in their risk landscape.

A natural consequence of the evolution of QRM thinking in different industries is an interest in the transfer of risks between industries; this process is known as ART (alternative risk transfer). To date the best examples are of risk transfer between the insurance and banking industries, as illustrated by the establishment in 1992 of catastrophe futures by the Chicago Board of Trade. These came about in the wake of Hurricane Andrew, which caused \$20 billion of insured losses on the East Coast of the US. While this was a considerable event for the insurance industry in relation to overall reinsurance capacity, it represented only a drop in the ocean compared with the daily volumes traded worldwide on financial exchanges. This led to the recognition that losses could be covered in future by the issuance of appropriately structured bonds with coupon streams and principal repayments dependent on the occurrence or non-occurrence of well-defined natural catastrophe events, such as storms and earthquakes.

A speculative view of where these developments may lead is given by Shiller (2003), who argues that the proliferation of RM thinking coupled with the technological sophistication of the twenty-first century will allow any agent in society, from a company to a country to an individual, to apply QRM methodology to the risks they face. In the case of an individual this may be the risk of unemployment, depreciation in the housing market or the investment in the education of children.

Notes and Comments

The language of probability and statistics plays a fundamental role throughout the book and readers are expected to have a good knowledge of these subjects. At the elementary level, Rice (1995) gives a good first introduction to both of these. More advanced texts in probability and stochastic processes are Williams (1991), Resnick (1992) and Rogers and Williams (1994); the full depth of these texts is certainly not required for the understanding of this book, though they provide excellent reading material for the mathematically more sophisticated reader who also has an interest in mathematical finance. Further recommended texts on statistical inference include Casella and Berger (2002), Bickel and Doksum (2001), Davison (2003) and Lindsey (1996).

An excellent text on the history of risk and probability with financial applications in mind is Bernstein (1998). Additional useful material on the history of the subject is to be found in Field (2003).

For the mathematical reader looking to acquire more knowledge of relevant economics we recommend Mas-Colell, Whinston and Green (1995) for microeconomics; Campbell, Lo and MacKinlay (1997) or Gouriéroux and Jasak (2001) for econometrics; and Brealey and Myers (2000) for corporate finance. From the vast literature on options, an entry-level text for the general reader is Hull (1997). At a more mathematical level we like Bingham and Kiesel (1998) and Musiela and Rutkowski (1997). One of the most readable texts on the basic notion of options is Cox and Rubinstein (1985). For a rather extensive list of the kind of animals to be found in the zoological garden of derivatives, see, for example, Haug (1998).

There are several texts on the spectacular losses due to speculative trading and careless use of derivatives. The LTCM case is well documented in Dunbar (2000), Lowenstein (2000) and Jorion (2000), the latter particularly for the technical risk-measurement issues involved. Boyle and Boyle (2001) give a very readable account of the Orange County, Barings and LTCM stories. A useful website on RM, containing a growing collection of industry case studies, is www.erisk.com.

An overview of options embedded in life insurance products is given in Dillmann (2002), guarantees are discussed in detail in Hardy (2003), and Briys and de Varenne (2001) contains an excellent account of RM issues facing the (life) insurance industry.

The historical development of banking regulation is well described in Crouhy, Galai and Mark (2001) and Steinherr (1998). For details of the current rules and regulations coming from the Basel Committee, see its website at www.bis.org/bcbs.

Besides copies of the various accords, one also finds useful working papers, publications and comments written by stakeholders on the various consultative packages. For Solvency 2, many documents are being prepared, and the Web is the best place to start looking; a forthcoming text is Sandström (2005). The complexity of RM methodology in the wake of Basel II is critically addressed by Hawke (2003), in his capacity as US Comptroller of the Currency.

For a very detailed overview of relevant practical issues underlying RM we again strongly recommend Crouhy, Galai and Mark (2001). A text stressing the use of VaR as a risk measure and containing several worked examples is Jorion (2001), who also has a useful teaching manual on the same subject (Jorion 2002a). Insurance-related issues in RM are well presented in Doherty (2000).

For a comprehensive discussion of the management of bank capital given regulatory constraints see Matten (2000). Graham and Rogers (2002) contains a discussion of RM and tax incentives. A formal account of the Modigliani–Miller Theorem and its implication can be found in many textbooks on corporate finance: a standard reference is Brealey and Myers (2000); de Matos (2001) gives a more theoretical account from the perspective of modern financial economics. Both texts also discuss the implications of informational asymmetries between the various stakeholders in a corporation. Formal models looking at RM from a corporate finance angle are to be found in Froot and Stein (1998), Froot, Scharfstein and Stein (1993) and Stulz (1996, 2002). For a specific discussion on corporate finance issues in insurance see Froot (2005) and Hancock, Huber and Koch (2001).

There are several studies on the use of RM techniques for non-financial firms (see, for example, Bodnar, Hyat and Marston 1999; Geman 2005). Two references in the area of reliability of industrial processes are Bedford and Cooke (2001) and Does, Roes and Trip (1999). An interesting edited volume on alternative risk transfer (ARTs) is Shimpi (1999); a detailed study of model risk in the ART context is Schmock (1999). An area we have not mentioned so far in our discussion of QRM in the future is that of real options. A real option is the right, but not the obligation, to take an action (e.g. deferring, expanding, contracting or abandoning) at a predetermined cost called the exercise price for a predetermined period of time—the life of the option. This definition is taken from Copeland and Antikarov (2001). Examples of real options discussed in the latter are the valuation of an internet project and of a pharmaceutical research and development project. A further useful reference is Brennan and Trigeorgis (2000).