

In life we simply cannot avoid the presence of risk. However, we tend to avoid its potential impact because, on the whole, we do a good job of risk management; we wear a bicycle helmet when cycling, we fasten our seat belts in a moving car, we use gloves when handling corrosive substances, etc. In the world of financial investments the universally held view is that the more risk we take the more we stand to gain but, just as importantly, the greater the chance we will lose. The task of the financial risk manager is to be aware of the presence of risk, to understand how it can damage a potential investment and, most of all, to be able to reduce the exposure to it in order to avert a potential disaster. It is the aim of this book to develop the mathematical tools which can be used to manage and control risks that are inherent in the financial market. We will be guided by two basic principles. Firstly, we shall endeavour to ensure that, on average, a financial investment provides a healthy return rate for a tolerable amount of risk. Secondly, we shall be prepared for rare market events whose impact could trigger a potentially catastrophic loss. The purpose of this chapter is to shed light on both the day-to-day issues and also the big challenges that a typical risk manager is likely to face, thus it serves as aperitif to stimulate the mathematical journey ahead.

## 1.1 BASIC CHALLENGES IN RISK MANAGEMENT

We open our discussion by considering a seemingly simple problem. Assume that we are armed with a wealth of  $\$W$  and we decide to invest this today, at time  $t$ , in a single financial asset for a period of  $\tau$  days into the future. The value of the asset today is known and denoted by  $S(t)$  but its future value  $S(t + \tau)$  is uncertain. We think of our asset being a simple market product such as a share in a stock, an amount of foreign currency or the ownership of a bond or some other commodity. In this situation there are two possible strategies:

- The holding strategy.  
If we believe the asset price will rise then we simply buy it today and sell it in the future at the (hopefully) higher price. In which case we make a profit from the purchase.
- The short-selling strategy.  
If we believe the asset price will fall then we can profit out of this situation by employing the strategy of short selling. This is summarized as follows.

$t$	$t + \tau$
Borrow the asset today	Buy the asset for $S(t + \tau)$
and	and
sell it immediately	return it to the lender
to receive $S(t)$	

If, as we suspected, the value of the asset falls (i.e., if  $S(t + \tau) < S(t)$ ) then we have made a profit.

The risk profiles of the two strategies are very different. The asset price can never drop below zero but, theoretically, it can grow without bound. For the holding strategy this means potentially unlimited profits and a bounded loss. However, for short selling the reverse is true and there is a potential for unlimited losses. In view of this one finds that, in practice, the process of pure short selling is supplemented by certain restrictions and safeguards.

We now suppose that we choose to invest our  $\$W$  in a collection of  $n$  risky assets denoted by  $\{S_1, \dots, S_n\}$ . Our strategy is simply to invest a fraction of our wealth,  $w_i$  say, in asset  $S_i$  for  $i = 1, \dots, n$ . We shall assume that short selling is allowed and so some of the  $w_i$  may be negative. This scenario leads us to our next challenging problem:

### The Portfolio Problem

*How can we choose an optimal set of weights  $\{w_1, \dots, w_n\}$ , so that our overall investment is likely to yield a promising return with minimal risk?*

Occasionally in mathematics one finds that seemingly complex problems have the most elegant and rewarding solutions. The portfolio problem above is such an example and it is the perfect starting point for our mathematical journey through risk management. The problem itself was solved in the early 1950s by Harry Markowitz (1952) in his PhD studies. The route that Markowitz took to derive his famous solution is as follows:

- Establish a formula for the random return rate for the portfolio, denoted by  $r_p$ , as a function of its weights  $w_1, \dots, w_n$ .
- Use basic probability theory to derive expressions for the expected return  $\mu_p$  (a measure of potential reward) and the volatility  $\sigma_p$  (a measure of risk) of  $r_p$ .
- We now search for the weights  $w_1^*, \dots, w_n^*$  that provide a desired level of expected return while ensuring that the risk involved is as small as possible.

The mathematical tools needed to attack this problem are developed in Chapters 2–4 and its full solution is delivered in Chapter 5; this will represent our first major landmark result.

Before Markowitz's theory emerged most investment decisions were made on the basis of gut instinct or simple advice such as *don't put all of your eggs in one basket*, there was little in the way of quantitative analysis. Markowitz gave investment theory a scientific footing and, in Chapter 6, we will discover some intriguing consequences of his pioneering work. Indeed these discoveries subsequently inspired many other researchers to investigate more deeply the relationship between the value of an asset and its perceived riskiness. This is a tough problem and one that is made even more difficult by the fact that asset prices do not always move of their own accord. More often than not we find that asset prices are related to each other. Strong price fluctuations in one asset will influence the movements of another and vice versa, we say they possess a correlation structure. This leads us to address the following.

### The Modelling Challenge

*How can we accurately model the way the price of a risky asset evolves through time?*

### The Correlation Challenge

*How can we accurately model the correlation structure of a collection of many risky assets?*

In the early 1960s Markowitz encouraged a PhD student, William Sharpe, to investigate these problems. To do this Sharpe imagined a world where all investors build their portfolios with Markowitz weights and, in this setting, he developed the famous Capital Asset Pricing Model (CAPM) Sharpe (1964). Chapter 7 of this book is devoted to the mathematical derivation of this model. We shall demonstrate some of its practical uses and its consequences, including the following intriguing discovery:

$$\begin{aligned} & \textit{Asset prices are related to each other through} \\ & \textit{their responses to a } \mathbf{single} \textit{ risk factor.} \end{aligned} \tag{1.1}$$

This revelation tells us that, in the Markowitz world, a single known risk factor can be viewed as the main driving force behind the movements and co-movements of all our risky assets. This conclusion is a remarkable one and, not surprisingly, it fuelled much debate amongst financial economists. Indeed, a great deal of empirical work has been done over the years to test the validity of the CAPM and its underlying assumptions.

In the 1970s a more cavalier approach to the development of financial risk models was taken. Specifically, inspired by the CAPM, the following more general situation was considered:

$$\begin{aligned} & \textit{Asset prices are related to each other through} \\ & \textit{their responses to } \mathbf{several} \textit{ risk factors.} \end{aligned} \tag{1.2}$$

In response to the above hypothesis a more general class of risk model was proposed, the so-called linear factor model. We will examine this popular approach in greater detail in Chapter 8 of this book. The most appealing feature of the linear factor model is the fact that there is a great deal of flexibility in the choice and composition of the driving factors. This flexibility leads us to an important practical risk management challenge.

### **The Factor Selection Challenge**

*How do we choose the number and nature of the driving risk factors?*

We shall conclude Chapter 8 by describing how principal components analysis, a famous dimension-reduction tool from multivariate statistics, can be used to deliver a useful scientific solution to this challenge.

## **1.2 VALUE AT RISK**

In the late 1980s fund managers and traders with complicated risk positions looked increasingly to a new breed of so-called derivative products as a means of dampening their risk profiles. Derivatives are literally products that are derived from simpler assets like those we have already encountered (i.e., stocks and shares, commodities, foreign currencies and bonds). When used correctly derivatives are able to protect those who hold them against risk; they can be viewed as a kind of insurance policy. However, as their popularity began to rise it became clear that the misuse of these products can have devastating consequences. Indeed, throughout the mid-1990s a whole host of derivatives-related disasters finally led to a much needed shake-up in the way banks were regulated. New tighter controls were imposed on financial institutions and consequently the industry as a whole had to rethink its approach to

risk management. In the present day all financial institutions have dedicated research teams of applied scientists who employ sophisticated mathematical and statistical methods to quantify and control exposure to risk. The risk-management revolution was initiated in the early 1990s when the famous Basel committee (on banking supervision) began a consultation process which, essentially, set about addressing the following important questions.

### **Ensuring Against Large Losses**

*How can investment banks measure their exposure to unfavourable and unanticipated movements of the basic financial assets?*

*How can they use this measure to determine their capital adequacy requirements?*

In order to attack this problem the committee proposed that each investment bank should divide its market positions into two books, the trading book and the bank book. The trading book, as its name suggests, contains all products that are used as part of an active day-to-day trading strategy (e.g., investment portfolios and derivatives would belong in the trading book). In contrast, the bank book consists of positions that are held over a much longer time horizon such as long-term loans.

The Basel committee directed its attention on the trading book and investigated how its riskiness could be quantified. The value of each product in the trading book has a price which can be discovered on the market (provided there is enough liquidity). The prices of these products in the future however are unknown, and thus, even though we may know the value of the trading book today, its value tomorrow or at any time in the future is unknown. When market conditions are calm one would hope that the trading portfolio would report a daily profit or at least only a mild, manageable loss. However, we cannot control market conditions and history dictates that, once in a while, we can expect a financial storm where an increase in market volatility can wipe away significant value from a financial product. In view of this a natural question to ask could be the following:

*What is the largest loss the trading book is likely to suffer 99 out of every 100 days?*

The answer to this question is known as the Value at Risk (VaR) for the trading book at the 99% confidence level; obviously the same question can be posed for other confidence levels, e.g., 95% represents the maximum likely loss in 95 out of 100 days. The idea of measuring the VaR of a portfolio is popular with practitioners; it represents a potential monetary loss and, in that respect, it is concise, practical and easy to understand. In 1996, the Basel committee added their own support to the VaR concept by proposing that banks could use VaR as a measure of its trading book's exposure to market risk. The final Basel committee report is viewed as pioneering for two reasons:

1. It endorses that investment banks can use their own internal models to calculate VaR estimates.
2. It provides all investment banks with a universal formula which they can use to calculate their own capital adequacy requirements; the formula is based upon the bank's own VaR estimates.

Value at Risk is widely regarded as one of the key milestones in the new risk-management revolution. However, the simplicity of the VaR concept disguises the complexity involved

in its measurement. For instance, before a single computation takes place we need to ensure that we have access to all relevant financial data, both historical and real time. Thus, a typical financial institution faces the following significant task:

### **The IT Challenge**

*Construct an IT system with the following functionality:*

- *Real-time position data for all products in the trading portfolio are gathered and correctly mapped to the risk calculation engine.*
- *A database that is dynamically populated with historical prices at regular intervals (e.g., daily prices) is accessible.*

This IT challenge is enormous, especially for multinational investment banks whose trading portfolio consists of products that span the global markets. Not surprisingly most investment banks choose to hand these data management projects over to one of the many IT consultancy firms with specialized skills in database architecture.

The VaR concept can be viewed as the trigger for a new approach to risk management; indeed, it marks the starting point of an exciting area of science where academic progress and real-world applications are in constant exchange. We consider the VaR calculation challenge in two parts.

### **The VaR Calculation Challenge**

*For a given confidence level  $\alpha \in (0, 1)$  how can we measure the corresponding Value at Risk for a portfolio which consists entirely of:*

1. *Basic financial assets such as stocks and shares, commodities, foreign currencies and bonds.*
2. *More complex derivatives products.*

We take up the first part of the VaR challenge in Chapter 9 where we examine its mathematical properties. We shall discover some of VaR's enticing features, however we also reveal some unfortunate problems. We endeavour to correct these problems by investigating alternative risk measures, and ask whether such candidates can be viewed as serious competitors to VaR.

In Chapter 10 we turn to the practical calculation of VaR and its associated challenges. As a starting point, we propose a basic model which assumes that the random variable representing the daily portfolio loss is normally distributed. In particular, for this simplified framework, we will show how we can derive neat closed-form solutions to almost all of the crucial VaR-related challenges.

The second part of the VaR challenge involves an additional level of complexity as we now allow derivative products to be included in the portfolio. In order to attack this problem we need some advanced results from probability theory and statistics and these are developed in Chapters 11 and 12. At the simplest level we can invest in a single derivative whose value depends upon the price of its underlying asset. Mathematically we say that the derivative price is a function of the asset price and write

$$\text{derivative price at } t = f(S(t)),$$

where  $f$  is some non-linear function. In order to examine the potential profit/loss associated with the derivative we need to determine, as accurately as possible, the form of  $f$ . This leads us to our next challenge:

### **Derivative Pricing**

*For a given derivative how can we determine the relationship between its value and the level of the underlying asset?*

Derivative pricing is a branch of mathematical finance in its own right and there are a whole host of excellent textbooks written on this subject (e.g., Higham (2004), Joshi (2005), Neftci (1996) and Wilmott, Howison and Dewynne (1995)). However, in Chapter 13 we provide a self-contained derivation of the celebrated Black–Scholes option pricing model for the simplest plain European options. This model dates back to the early 1970s and yet its impact on the development of modern mathematical finance cannot be overstated; a great deal of the pioneering work on derivative pricing can be viewed as an extension or an innovation of the original Black–Scholes model.

We will not pursue derivative pricing in any further depth, but will simply assume that a calculation engine exists and is able to deliver a price for any derivative we encounter. In this situation we are able to tackle the problem of computing VaR estimates for a portfolio of derivatives. In a deliberate effort to reduce the computational burden of this problem we shall investigate the possibility of providing a closed-form solution. We remark that this problem is difficult for at least two reasons:

1. The number of underlying assets (upon which the derivatives are written) can be very large, i.e., the problem is a high-dimensional one.
2. Even if we understand the probabilistic nature of a particular asset it is much harder to predict how a non-linear function (i.e., a derivative) of it will behave.

In the late 1990s Britten-Jones and Schaeffer (1999) tackled the above issues and proposed the following recipe:

- Step 1. Dimension reduction.  
A linear factor model is proposed as a model for the changes in the underlying asset returns. It is assumed that the number of factors is much smaller than the number of assets and thus the size (dimension) of the problem is greatly reduced.
- Step 2. Probabilistic assumptions.  
Some simplifying assumptions are proposed for the probabilistic laws that govern the random nature of the risk factors.
- Step 3. Function approximations.  
A local approximation of the non-linear derivative function is made.

In Chapter 14 we will develop the above steps in detail and show how the local approximations to the derivatives can be used to provide closed-form expressions for so-called non-linear Value at Risk.

## **1.3 FURTHER CHALLENGES IN RISK MANAGEMENT**

The early attempts to calculate VaR were made in the mid-1990s and, during this time, the main priority for most practitioners was to establish a straightforward solution that could

be implemented with ease. As a result these early attempts were based upon rather simple assumptions regarding the random behaviour of the financial losses/returns. Towards the end of the 1990s almost all financial institutions took advantage of the rapid advances in information technology, where faster computing speed coupled with increased data storage enabled teams of quantitative analysts to perform deeper scientific investigations. A particularly important example is to use historical data to help gain an insight into the characteristic properties of the underlying financial variables; indeed, this becomes the focus of our next challenge:

### **Statistical Investigation**

*Using realized price data, perform a statistical investigation to determine the key empirical properties of asset losses/returns.*

In Chapters 15–18 we develop the statistical tools and techniques needed to tackle this problem. Then, in Chapter 19, we put these tools into action and conduct a scientific investigation whose aim is to pin down the key statistical properties that characterize the true nature of financial losses/returns. These properties are commonly referred to as the stylized facts and they serve as a guide for the development of new and improved risk models; a successful mathematical model should capture as many (if not all) of these properties as possible.

One particular result of our investigation is the observation that extreme values tend to occur more often than some of the basic models would predict, with large losses occurring more often than large profits. In relation to this we also discover that the future volatility of a basic financial asset is closely related to its past. This is an important observation because it implies that when an asset experiences a period of high volatility the likelihood of an extreme swing is increased; unfortunately, the swing can be downward as well as upward. These observations lead us to one of the central questions that all risk managers must address:

### **The Volatility Challenge**

*How can we construct a time-dependent volatility model which accurately captures the stylized facts of financial losses?*

The topic of volatility modelling is so large that it can also be regarded as a branch of mathematical finance in its own right, indeed there are several textbooks devoted to this topic (e.g., Gouriéroux (1997), Poon (2005) and Taylor (2007)). We take on the volatility challenge in Chapter 20 where, rather than provide a bite-size review of the many different approaches, we present the mathematical story of one of the most popular, the so-called GARCH family of models. GARCH models have found a wide range of applications in financial modelling because, as we shall discover, they have the ability to capture almost all of the stylized facts and, what's more, they are also fairly simple to implement. The GARCH modelling framework is also extremely flexible; once one understands the basic model, it is then possible to introduce extensions designed to enhance its performance. This is reflected in the vast range of innovative GARCH-type volatility models on the market.

In order to motivate the next important challenge we recall that our Value at Risk measure, as we know it, is designed to cope with those unanticipated events which typically occur two



or three times in a year. Unfortunately, however, experience has shown that financial markets can also be exposed to tornado-like events such as terrorist attacks, political instabilities and natural disasters. These events have the potential to wipe billions off the value of global stock markets. Thus, one of the new challenges of mathematical risk management is to develop a methodology to cater for such extreme events. In this respect we face a new challenge:

### **The Challenge of Quantifying Losses Due to Rare Events**

*How do we assign appropriate probabilities to potential extreme movements of a financial asset?*

We tackle this problem in Chapter 21 where we appeal to extreme value theory (EVT), a branch of probability theory that is concerned with describing the statistical properties of extreme events. EVT has applications in many areas of science and engineering. In particular, hydrologists have successfully used EVT to help predict the likelihood and size of potentially damaging floods, the hydrologist then uses these findings to estimate the optimal height of a dam which is to be constructed to protect against such floods. In finance the application of EVT is much the same; the risk manager uses the theory to model the likelihood and size of a portfolio loss due to a financial storm, he can then use this data to determine the size of the buffer fund which is designed to absorb such losses.

A common situation in finance (and other branches of science and engineering) is that closed-form solutions to real-world problems only tend to be available in the simplest of cases. For instance, the fair price of a plain European option can be derived analytically, however numerical methods are needed for most non-standard options. We find this in risk management too, for instance if we (erroneously) assume that the portfolio loss random variable is normally distributed then we can derive expressions for almost any risk measure, however if a more sophisticated risk model is used then we must turn to numerical techniques. This leads us to our next challenge:

### **Numerical Methods for Risk Quantification**

*How can we develop numerical techniques to compute the Value at Risk for a financial portfolio?*

In order to address this problem we must study the mathematical ideas behind one of the most crucial numerical tools in risk management – the ability to perform numerical simulations. We take on this challenge in Chapter 22 where we demonstrate how simulation techniques can be used to deliver estimates of financial risk measures such as Value at Risk. In particular we describe how to design a simulation algorithm whose purpose is to generate a range of potential future prices for each asset and/or derivative in the portfolio, these are then combined to produce a simulated value of the portfolio. Then, as more and more simulated values are generated, a clearer picture of the crucial statistical properties of the portfolio loss random variable emerges and, as a result, estimates for VaR (and other risk measures) can easily be derived. The success of the method lies in the specification of the algorithm. It may depend upon the past price history of the portfolio (historical simulation) or it may depend upon some mathematical model that is calibrated to real market prices (Monte Carlo simulation).



Obviously, from a practitioner's perspective, an accurate closed-form expression for VaR is highly desirable. Indeed, in the late 1990s several alternative VaR methodologies were proposed, each delivering closed-form solutions while attempting to simultaneously capture the true statistical properties of the loss random variable. In Chapter 23 we shall present two of the most commonly used methods and by doing so we bring the story of VaR calculation methods to a close.

At the end of the day the model that is finally selected to compute the VaR of the trading portfolio is of particular importance. The resulting VaR calculations will determine the size of the institution's buffer fund, the amount of regulatory capital it must set aside to help absorb unexpected losses. The buffer fund cannot be used for investment purposes, it is off-limits and its size must adhere to the regulator's formula. If we choose a model that consistently overestimates the true VaR then we will be overcommitting funds that could otherwise be used to generate profits. On the other hand, if we choose a model that consistently underestimates the true VaR then it will be punished; the regulator will revise its formula so that the size of the buffer fund is increased, i.e., the regulator penalizes a substandard VaR model. In view of these two influences we must select the most accurate calculation method to suit the characteristics of our trading portfolio, i.e., we must face the following challenge:

### **Verification of Risk Models**

*How can we scientifically test the performance of a particular VaR model?*

We address this question in Chapter 24, the final chapter of the book. Specifically we develop Christoffersen's testing methodology Christoffersen (1998) that dates back to the late 1990s. The idea here is to appeal to our earlier review of statistical testing (Chapter 18) and use this to propose a certain test statistic whose value is dependent upon the past performance of the VaR model. The test statistic itself can be viewed as a random quantity which obeys certain known probability laws and we can use this fact to construct a decision rule that determines whether the model should be accepted or rejected.

