Measuring Market Risk

Kevin Dowd

JOHN WILEY & SONS, LTD

Measuring Market Risk

Measuring Market Risk

Kevin Dowd

JOHN WILEY & SONS, LTD

Published 2002 John Wiley & Sons Ltd,
The Atrium, Southern Gate, Chichester,
West Sussex PO19 8SQ, England

Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): cs-books@wiley.co.uk
Visit our Home Page on www.wileyeurope.com or www.wiley.com

Other Wiley Editorial Offices

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA

Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany

John Wiley & Sons Australia Ltd, 33 Park Road, Milton, Queensland 4064, Australia

John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark, Singapore 129809

John Wiley & Sons Canada Ltd, 22 Worcester Road, Etobicoke, Ontario, Canada M9W 1L1

Library of Congress Cataloging-in-Publication Data

Dowd, Kevin.
Measuring market risk / Kevin Dowd.
p. cm. — (Wiley finance series)
Includes bibliographical references and index.
ISBN 0-471-52174-4 (alk. paper)
1. Financial futures. 2. Risk management. I. Title. II. Series.
HG6024.3 .D683 2002
332.63′2042—dc21 2002071367

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN 0-471-52174-4

Typeset in 10/12pt Times by TechBooks, New Delhi, India
Printed and bound in Great Britain by TJ International, Padstow, Cornwall, UK
This book is printed on acid-free paper responsibly manufactured from sustainable forestry in which at least two trees are planted for each one used for paper production.

Wiley Finance Series

Brand Assets
Tony Tollington

Swaps and other Derivatives
Richard Flavell

An Introduction to Capital Markets: Products, Strategies Participants
Andrew Chisholm

Asset Management: Equities Demystified
Shanta Acharya

Currency Strategy: A Practitioner's Guide to Currency Trading, Hedging and Forecasting
Callum Henderson

Hedge Funds: Myths and Limits
Francois-Serge Lhabitant

The Manager's Concise Guide to Risk
Jihad S Nader

Securities Operations: A Guide to Trade and Position Management
Michael Simmons

Modelling, Measuring and Hedging Operational Risk
Marcelo Cruz

Monte Carlo Methods in Finance
Peter Jäckel

Building and Using Dynamic Interest Rate Models
Ken Kortanek and Vladimir Medvedev

Structured Equity Derivatives: The Definitive Guide to Exotic Options and Structured Notes
Harry Kat

Advanced Modelling in Finance Using Excel and VBA
Mary Jackson and Mike Staunton

Operational Risk: Measurement and Modelling
Jack King

Advance Credit Risk Analysis: Financial Approaches and Mathematical Models to Assess, Price and Manage Credit Risk
Didier Cossin and Hugues Pirotte

Dictionary of Financial Engineering
John F. Marshall

Pricing Financial Derivatives: The Finite Difference Method
Domingo A Tavella and Curt Randall

Interest Rate Modelling
Jessica James and Nick Webber

Handbook of Hybrid Instruments: Convertible Bonds, Preferred Shares, Lyons, ELKS, DECS and Other Mandatory Convertible Notes
Izzy Nelken (ed)

Options on Foreign Exchange, Revised Edition
David F DeRosa

Volatility and Correlation in the Pricing of Equity, FX and Interest-Rate Options
Riccardo Rebonato

Risk Management and Analysis vol. 1: Measuring and Modelling Financial Risk
Carol Alexander (ed)

Risk Management and Analysis vol. 2: New Markets and Products
Carol Alexander (ed)

Implementing Value at Risk
Philip Best

Implementing Derivatives Models
Les Clewlow and Chris Strickland

Interest-Rate Option Models: Understanding, Analysing and Using Models for Exotic Interest-Rate Options (second edition)
Riccardo Rebonato

Contents

Preface

> You are responsible for managing your company's foreign exchange positions. Your boss, or your boss's boss, has been reading about derivatives losses suffered by other companies, and wants to know if the same thing could happen to his company. That is, he wants to know just how much market risk the company is taking. What do you say?
>
> You could start by listing and describing the company's positions, but this isn't likely to be helpful unless there are only a handful. Even then, it helps only if your superiors understand all of the positions and instruments, and the risks inherent in each. Or you could talk about the portfolio's sensitivities, i.e., how much the value of the portfolio changes when various underlying market rates or prices change, and perhaps option delta's and gamma's. However, you are unlikely to win favour with your superiors by putting them to sleep. Even if you are confident in your ability to explain these in English, you still have no natural way to net the risk of your short position in Deutsche marks against the long position in Dutch guilders. . . . You could simply assure your superiors that you never speculate but rather use derivatives only to hedge, but they understand that this statement is vacuous. They know that the word 'hedge' is so ill-defined and flexible that virtually any transaction can be characterized as a hedge. So what do you say? (Linsmeier and Pearson (1996, p.1))

The obvious answer, 'The most we can lose is . . . ' is also clearly unsatisfactory, because the most we can possibly lose is everything, and we would hope that the board already knows that. Consequently, Linsmeier and Pearson continue, "Perhaps the best answer starts: 'The value at risk is . . . '".

So what is value at risk? Value at risk (VaR) is our maximum likely loss over some target period — the most we expect to lose over that period, at a specified probability level. It says that on 95 days out of 100, say, the most we can expect to lose is \$10 million or whatever. This is a good answer to the problem posed by Linsmeier and Pearson. The board or other recipients specify their probability level — 95%, 99% and so on — and the risk manager can tell them the maximum they can lose at that probability level. The recipients can also specify the horizon period — the next day, the next week, month, quarter, etc. — and again the risk manager can tell them the maximum amount they stand to lose over that horizon period. Indeed, the recipients can specify any combination of probability and horizon period, and the risk manager can give them the VaR applicable to that probability and horizon period.

We then have to face the problem of how to measure the VaR. This is a tricky question, and the answer is very involved and takes up much of this book. The short answer is, therefore, to read this book or others like it.

However, before we get too involved with VaR, we also have to face another issue. Is a VaR measure the best we can do? The answer is no. There are alternatives to VaR, and at least one of these — the so-called expected tail loss (ETL) or expected shortfall — is demonstrably superior. The ETL is the loss we can expect to make if we get a loss in excess of VaR. Consequently, I would take issue with Linsmeier and Pearson's answer. 'The VaR is ...' is generally a reasonable answer, but it is not the best one. A better answer would be to tell the board the ETL — or better still, show them curves or surfaces plotting the ETL against probability and horizon period. Risk managers who use VaR as their preferred risk measure should really be using ETL instead. VaR is already passé.

But if ETL is superior to VaR, why both with VaR measurement? This is a good question, and also a controversial one. Part of the answer is that there will be a need to measure VaR for as long as there is a demand for VaR itself: if someone wants the number, then someone has to measure it, and whether they should want the number in the first place is another matter. In this respect VaR is a lot like the infamous beta. People still want beta numbers, regardless of the well-documented problems of the Capital Asset Pricing Model on whose validity the beta risk measure depends. A purist might say they shouldn't, but the fact is that they do. So the business of estimating betas goes on, even though the CAPM is now widely discredited. The same goes for VaR: a purist would say that VaR is inferior to ETL, but people still want VaR numbers and so the business of VaR estimation goes on. However, there is also a second, more satisfying, reason to continue to estimate VaR: we often need VaR estimates to be able to estimate ETL. We don't have many formulas for ETL and, as a result, we would often be unable to estimate ETL if we had to rely on ETL formulas alone. Fortunately, it turns out that we can always estimate the ETL if we can estimate the VaR. The reason is that the VaR is a quantile and, if we can estimate the quantile, we can easily estimate the ETL — because the ETL itself is just a quantile average.

INTENDED READERSHIP

This book provides an overview of the state of the art in VaR and ETL estimation. Given the size and rate of growth of this literature, it is impossible to cover the field comprehensively, and no book in this area can credibly claim to do so, even one like this that focuses on risk measurement and does not really try to grapple with the much broader field of market risk management. Within the sub-field of market risk measurement, the coverage of the literature provided here — with a little under 400 references — is fairly extensive, but can only provide, at best, a rather subjective view of the main highlights of the literature.

The book is aimed at three main audiences. The first consists of practitioners in risk measurement and management — those who are developing or already using VaR and related risk systems. The second audience consists of students in MBA, MA, MSc and professional programmes in finance, financial engineering, risk management and related subjects, for whom the book can be used as a textbook. The third audience consists of PhD students and academics working on risk measurement issues in their research. Inevitably, the level at which the material is pitched must vary considerably, from basic (e.g., in Chapters 1 and 2) to advanced (e.g., the simulation methods in Chapter 6). Beginners will therefore find some of it heavy going, although they should get something out of it by skipping over difficult parts and trying to get an overall feel for the material. For their part, advanced readers will find a lot of familiar material, but many of them should, I hope, find some material here to engage them.

To get the most out of the book requires a basic knowledge of computing and spreadsheets, statistics (including some familiarity with moments and density/distribution functions), mathematics (including basic matrix algebra) and some prior knowledge of finance, most especially derivatives and fixed-income theory. Most practitioners and academics should have relatively little difficulty with it, but for students this material is best taught after they have already done their quantitative methods, derivatives, fixed-income and other 'building block' courses.

USING THIS BOOK

This book is divided into two parts — the chapters that discuss risk measurement, presupposing that the reader has the technical tools (i.e., the statistical, programming and other skills) to follow the discussion, and the toolkit at the end, which explains the main tools needed to understand market risk measurement. This division separates the material dealing with risk measurement *per se* from the material dealing with the technical tools needed to carry out risk measurement. This helps to simplify the discussion and should make the book much easier to read: instead of going back and forth between technique and risk measurement, as many books do, we can read the technical material first; once we have the tools under our belt, we can then focus on the risk measurement without having to pause occasionally to re-tool.

I would suggest that the reader begin with the technical material — the tools at the end — and make sure that this material is adequately digested. Once that is done, the reader will be equipped to follow the risk measurement material without needing to take any technical breaks. My advice to those who might use the book for teaching purposes is the same: first cover the tools, and then do the risk measurement. However, much of the chapter material can, I hope, be followed without too much difficulty by readers who don't cover the tools first; but some of those who read the book in this way will occasionally find themselves having to pause to tool up.

In teaching market risk material over the last few years, it has also become clear to me that one cannot teach this material effectively — and students cannot really absorb it — if one teaches only at an abstract level. Of course, it is important to have lectures to convey the conceptual material, but risk measurement is not a purely abstract subject, and in my experience students only really grasp the material when they start playing with it — when they start working out VaR figures for themselves on a spreadsheet, when they have exercises and assignments to do, and so on. When teaching, it is therefore important to balance lecture-style delivery with practical sessions in which the students use computers to solve illustrative risk measurement problems.

If the book is to be read and used practically, readers also need to use appropriate spreadsheets or other software to carry out estimations for themselves. Again, my teaching and supervision experience is that the use of software is critical in learning this material, and we can only ever claim to understand something when we have actually measured it. The software and risk material are also intimately related, and the good risk measurer knows that risk measurement always boils down to some spreadsheet or other computer function. In fact, much of the action in this area boils down to software issues — comparing alternative software routines, finding errors, improving accuracy and speed, and so forth. Any risk measurement book should come with at least some indication of how risk measurement routines can be implemented on a computer.

It is better still for such books to come with their own software, and this book comes with a CD that contains 150 risk measurement and related functions in MATLAB and a manual explaining their use.[1] My advice to users is to print out the manual and go through

[1]MATLAB is a registered trademark of The MathWorks, Inc. For more information on MATLAB, please visit their website, www.mathworks.com., or contact The MathWorks, Inc., 3 Apple Hill Drive, Natick, MA 01760-2098, USA.

the functions on a computer, and then keep the manual to hand for later reference.[2] The examples and figures in the book are produced using this software, and readers should be able to reproduce them for themselves. Readers are very welcome to contact me with any feedback; however, I would ask any who do so to bear in mind that because of time pressures I cannot guarantee a reply. Nonetheless, I will keep the toolkit and the manual up-to-date on my website (www.nottingham.ac.uk/∼lizkd) and readers are welcome to download updates from there.

In writing this software, I should explain that I chose MATLAB mainly because it is both powerful and user-friendly, unlike its obvious alternatives (VBA, which is neither powerful nor particularly user-friendly, or the C or S languages, which are certainly not user-friendly). I also chose MATLAB in part because it produces very nice graphics, and a good graph or chart is often an essential tool for risk measurement. Unfortunately, the downside of MATLAB is that many users of the book will not be familiar with it or will not have ready access to it, and I can only advise such readers to think seriously about going through the expense and/or effort to get it.[3]

In explaining risk measurement throughout this book, I have tried to focus on the underlying ideas rather than on programming code: understanding the ideas is much more important, and the coding itself is mere implementation. My advice to risk measurers is that they should aim to get to the level where they can easily write their own code once they know what they are trying to do. However, for those who want it, the code I use is easily accessible — one simply opens up MATLAB, goes into the Measuring Market Risk (MMR) Toolbox, and opens the relevant function. The reader who wants the code should therefore refer directly to the program coding rather than search around in the text: I have tried to keep the text itself free of such detail to focus on more important conceptual issues.

The MMR Toolbox also has many other functions besides those used to produce the examples or figures in the text. I have tried to produce a fairly extensive set of software functions that would cover all the obvious VaR or ETL measurement problems, as well as some of the more advanced ones. Users — such as students doing their dissertations, academics doing their research, and practitioners working on practical applications — might find some of these functions useful, and they are welcome to make whatever use of these functions they wish. However, before anyone takes the MMR functions too seriously, they should appreciate that I am not a programmer, and anyone who uses these functions must do so at his or her own risk. As always in risk measurement, we should keep our wits about us and not be too trusting of the software we use or the results we get.

[2] The user should copy the Measuring Market Risk folder into his or her MATLAB works folder and activate the path to the Measuring Market Risk folder thus created (so MATLAB knows the folder is there). The functions were written in MATLAB 6.0 and most of the MMR functions should work if the user has the Statistics Toolbox as well as the basic MATLAB 6.0 or later software installed on their machine. However, a small number of MMR functions draw on functions in other MATLAB toolboxes (e.g., such as the Garch Toolbox), so users with only the Statistics Toolbox will find that the occasional MMR function does not work on their machine.

[3] When I first started working on this book, I initially tried writing the software functions in VBA to take advantage of the fact that almost everyone has access to Excel; unfortunately, I ran into too many problems and eventually had to give up. Had I not done so, I would still be struggling with VBA code even now, and this book would never have seen the light of day. So, whilst I sympathise with those who might feel pressured to learn MATLAB or some other advanced language and obtain the relevant software, I don't see any practical alternative: if you want software, Excel/VBA is just not up to the job — although it can be useful for many simpler tasks and for teaching at a basic level.

However, for those addicted to Excel, the enclosed CD also includes a number of Excel workbooks to illustrate some basic risk measurement functions in Excel. Most of these are not especially powerful, but they give an idea of how one might go about risk measurement using Excel. I should add, too, that some of these were written by Peter Urbani, and I would like to thank Peter for allowing me to include them here.

OUTLINE OF THE BOOK

As mentioned earlier, the book is divided into the chapters proper and the toolkit at the end that deals with the technical issues underlying (or the tools needed for) market risk measurement. It might be helpful to give a brief overview of these so readers know what to expect.

The Chapters

The first chapter provides a brief overview of recent developments in risk measurement — market risk measurement especially — to put VaR and ETL in their proper context. Chapter 2 then looks at different measures of financial risk. We begin here with the traditional mean–variance framework. This framework is very convenient and provides the underpinning for modern portfolio theory, but it is also limited in its applicability because it has difficulty handling skewness (or asymmetry) and 'fat tails' (or fatter than normal tails) in our P/L or return probability density functions. We then consider VaR and ETL as risk measures, and compare them to traditional risk measures and to each other.

Having established what our basic risk measures actually are, Chapter 3 has a first run through the issues involved in estimating them. We cover three main sets of issues here:

- Preliminary data issues — how to handle data in profit/loss (or P/L) form, rate of return form, etc.
- How to estimate VaR based on alternative sets of assumptions about the distribution of our data and how our VaR estimation procedure depends on the assumptions we make.
- How to estimate ETL — and, in particular, how we can always approximate ETL by taking it as an average of 'tail VaRs' or losses exceeding VaR.

Chapter 3 is followed by an appendix dealing with the important subject of mapping — the process of describing the positions we hold in terms of combinations of standard building blocks. We would use mapping to cut down on the dimensionality of our portfolio, or deal with possible problems caused by having closely correlated risk factors or missing data. Mapping enables us to estimate market risk in situations that would otherwise be very demanding or even impossible.

Chapter 4 then takes a closer look at non-parametric VaR and ETL estimation. Non-parametric approaches are those in which we estimate VaR or ETL making minimal assumptions about the distribution of P/L or returns: we let the P/L data speak for themselves as much as possible. There are various non-parametric approaches, and the most popular is historical simulation (HS), which is conceptually simple, easy to implement, widely used and has a fairly good track record. We can also carry out non-parametric estimation using non-parametric density approaches (see Tool No. 5) and principal components and factor analysis methods (see Tool No. 6); the latter methods are sometimes useful when dealing with high-dimensionality problems (i.e., when dealing with portfolios with very large numbers of risk factors). As a general rule, non-parametric methods work fairly well if market conditions remain reasonably stable, and they are capable of considerable refinement and improvement. However, they can be unreliable if market conditions change, their results are totally dependent on the data set, and their estimates of VaR and ETL are subject to distortions from one-off events and ghost effects.

Chapter 5 looks more closely at parametric approaches, the essence of which is that we fit probability curves to the data and then infer the VaR or ETL from the fitted curve. Parametric approaches are more powerful than non-parametric ones, because they make use of additional information contained in the assumed probability density function. They are also easy to use, because they give rise to straightforward formulas for VaR and sometimes ETL, but are vulnerable to error if the assumed density function does not adequately fit the data. The chapter discusses parametric VaR and ETL at two different levels — at the portfolio level, where we are dealing with portfolio P/L or returns, and assume that the underlying distribution is normal, Student *t*, extreme value or whatever and at the sub-portfolio or individual position level, where we deal with the P/L or returns to individual positions and assume that these are multivariate normal, elliptical, etc., and where we look at both correlation- and copula-based methods of obtaining portfolio VaR and ETL from position-level data. This chapter is followed by appendices dealing with the use of delta–gamma and related approximations to deal with non-linear risks (e.g., such as those arising from options), and with analytical solutions for the VaR of options positions.

Chapter 6 examines how we can estimate VaR and ETL using simulation (or random number) methods. These methods are very powerful and flexible, and can be applied to many different types of VaR or ETL estimation problem. Simulation methods can be highly effective for many problems that are too complicated or too messy for analytical or algorithmic approaches, and they are particularly good at handling complications like path-dependency, non-linearity and optionality. Amongst the many possible applications of simulation methods are to estimate the VaR or ETL of options positions and fixed-income positions, including those in interest-rate derivatives, as well as the VaR or ETL of credit-related positions (e.g., in default-risky bonds, credit derivatives, etc.), and of insurance and pension-fund portfolios. We can also use simulation methods for other purposes — for example, to estimate VaR or ETL in the context of dynamic portfolio management strategies. However, simulation methods are less easy to use than some alternatives, usually require a lot of calculations, and can have difficulty dealing with early-exercise features.

Chapter 7 looks at tree (or lattice or grid) methods for VaR and ETL estimation. These are numerical methods in which the evolution of a random variable over time is modelled in terms of a binomial or trinomial tree process or in terms of a set of finite difference equations. These methods have had a limited impact on risk estimation so far, but are well suited to certain types of risk estimation problem, particularly those involving instruments with early-exercise features. They are also fairly straightforward to program and are faster than some simulation methods, but we need to be careful about their accuracy, and they are only suited to low-dimensional problems.

Chapter 8 considers risk addition and decomposition — how changing our portfolio alters our risk, and how we can decompose our portfolio risk into constituent or component risks. We are concerned here with:

- *Incremental risks.* These are the changes in risk when a factor changes — for example, how VaR changes when we add a new position to our portfolio.
- *Component risks.* These are the component or constituent risks that make up a certain total risk — if we have a portfolio made up of particular positions, the portfolio VaR can be broken down into components that tell us how much each position contributes to the overall portfolio VaR.

Both these (and their ETL equivalents) are extremely useful measures in portfolio risk management: amongst other uses, they give us new methods of identifying sources of risk, finding natural hedges, defining risk limits, reporting risks and improving portfolio allocations.

Chapter 9 examines liquidity issues and how they affect market risk measurement. Liquidity issues affect market risk measurement not just through their impact on our standard measures of market risk, VaR and ETL, but also because effective market risk management involves an ability to measure and manage liquidity risk itself. The chapter considers the nature of market liquidity and illiquidity, and their associated costs and risks, and then considers how we might take account of these factors to estimate VaR and ETL in illiquid or partially liquid markets. Furthermore, since liquidity is important in itself and because liquidity problems are particularly prominent in market crises, we also need to consider two other aspects of liquidity risk measurement — the estimation of liquidity at risk (i.e., the liquidity equivalent to value at risk), and the estimation of crisis-related liquidity risks.

Chapter 10 deals with backtesting — the application of quantitative, typically statistical, methods to determine whether a model's risk estimates are consistent with the assumptions on which the model is based or to rank models against each other. To backtest a model, we first assemble a suitable data set — we have to 'clean' accounting data, etc. — and it is good practice to produce a backtest chart showing how P/L compares to measured risk over time. After this preliminary data analysis, we can proceed to a formal backtest. The main classes of backtest procedure are:

- Statistical approaches based on the frequency of losses exceeding VaR.
- Statistical approaches based on the sizes of losses exceeding VaR.
- Forecast evaluation methods, in which we score a model's forecasting performance in terms of a forecast error loss function.

Each of these classes of backtest comes in alternative forms, and it is generally advisable to run a number of them to get a broad feel for the performance of the model. We can also backtest models at the position level as well as at the portfolio level, and using simulation or bootstrap data as well as 'real' data. Ideally, 'good' models should backtest well and 'bad' models should backtest poorly, but in practice results are often much less clear: in this game, separating the sheep from the goats is often much harder than many imagine.

Chapter 11 examines stress testing — 'what if' procedures that attempt to gauge the vulnerability of our portfolio to hypothetical events. Stress testing is particularly good for quantifying what we might lose in crisis situations where 'normal' market relationships break down and VaR or ETL risk measures can be very misleading. VaR and ETL are good on the probability side, but poor on the 'what if' side, whereas stress tests are good for 'what if' questions and poor on probability questions. Stress testing is therefore good where VaR and ETL are weak, and vice versa. As well as helping to quantify our exposure to bad states, the results of stress testing can be a useful guide to management decision-making and help highlight weaknesses (e.g., questionable assumptions, etc.) in our risk management procedures.

The final chapter considers the subject of model risk — the risk of error in our risk estimates due to inadequacies in our risk measurement models. The use of any model always entails exposure to model risk of some form or another, and practitioners often overlook this exposure because it is out of sight and because most of those who use models have a tendency to end up 'believing' them. We therefore need to understand what model risk is, where and how it arises, how to measure it, and what its possible consequences might be. Interested parties such as risk practitioners and their managers also need to understand what they can do to combat it. The problem of model risk never goes away, but we can learn to live with it.

The MMR Toolkit

We now consider the Measuring Market Risk Toolkit, which consists of 11 different 'tools', each of which is useful for risk measurement purposes. Tool No. 1 deals with how we can estimate the standard errors of quantile estimates. Quantiles (e.g., such as VaR) give us the quantity values associated with specified probabilities. We can easily obtain quantile estimates using parametric or non-parametric methods, but we also want to be able to estimate the precision of our quantile estimators, which can be important when estimating confidence intervals for our VaR.

Tool No. 2 deals with the use of the theory of order statistics for estimating VaR and ETL. Order statistics are ordered observations — the biggest observation, the second biggest observation, etc. — and the theory of order statistics enables us to predict the distribution of each ordered observation. This is very useful because the VaR itself is an order statistic — for example, with 100 P/L observations, we might take the VaR at the 95% confidence level as the sixth largest loss observation. Hence, the theory of order statistics enables us to estimate the whole of the VaR probability density function — and this enables us to estimate confidence intervals for our VaR. Estimating confidence intervals for ETLs is also easy, because there is a one-to-one mapping from the VaR observations to the ETL ones: we can convert the P/L observations into average loss observations, and apply the order statistics approach to the latter to obtain ETL confidence intervals.

Tool No. 3 deals with the Cornish–Fisher expansion, which is useful for estimating VaR and ETL when the underlying distribution is near normal. If our portfolio P/L or return distribution is not normal, we cannot take the VaR to be given by the percentiles of an inverse normal distribution function; however, if the non-normality is not too severe, the Cornish–Fisher expansion gives us an adjustment factor that we can use to correct the normal VaR estimate for non-normality. The Cornish–Fisher adjustment is easy to apply and enables us to retain the easiness of the normal approach to VaR in some circumstances where the normality assumption itself does not hold.

Tool No. 4 deals with bootstrap procedures. These methods enable us to sample repeatedly from a given set of data, and they are useful because they give a reliable and easy way of estimating confidence intervals for any parameters of interest, including VaRs and ETLs.

Tool No. 5 discusses the subject of non-parametric density estimation: how we can best represent and extract the most information from a data set without imposing parametric assumptions on the data. This topic covers the use and usefulness of histograms and related methods (e.g., naïve and kernel estimators) as ways of representing our data, and how we can use these to estimate VaR.

Tool No. 6 covers principal components analysis and factor analysis, which are alternative methods of gaining insight into the properties of a data set. They are helpful in risk measurement because they can provide a simpler representation of the processes that generate a given data set, which then enables us to reduce the dimensionality of our data and so reduce the number of variance–covariance parameters that we need to estimate. Such methods can be very useful when we have large-dimension problems (e.g., variance–covariance matrices with hundreds of different instruments), but they can also be useful for cleaning data and developing data mapping systems.

The next tool deals with fat-tailed distributions. It is important to consider fat-tailed distributions because most financial returns are fat-tailed and because the failure to allow for fat tails can lead to major underestimates of VaR and ETL. We consider five different ways of representing fat tails: stable Lévy distributions, sometimes known as α-stable or stable Paretian distributions; Student *t*-distributions; mixture-of-normal distributions; jump diffusion distributions; and distributions with truncated Lévy flight. Unfortunately, with the partial exception of the Student *t*, these distributions are not nearly as tractable as the normal distribution, and they each tend to bring their own particular baggage. But that's the way it is in risk measurement: fat tails are a real problem.

Tool No. 8 deals with extreme value theory (EVT) and its applications in financial risk management. EVT is a branch of statistics tailor-made to deal with problems posed by extreme or rare events — and in particular, the problems posed by estimating extreme quantiles and associated probabilities that go well beyond our sample range. The key to EVT is a theorem — the extreme value theorem — that tells us what the distribution of extreme values should look like, at least asymptotically. This theorem and various associated results tell us what we should be estimating, and also give us some guidance on estimation and inference issues.

Tool No. 9 then deals with Monte Carlo and related simulation methods. These methods can be used to price derivatives, estimate their hedge ratios, and solve risk measurement problems of almost any degree of complexity. The idea is to simulate repeatedly the random processes governing the prices or returns of the financial instruments we are interested in. If we take enough simulations, the simulated distribution of portfolio values will converge to the portfolio's unknown 'true' distribution, and we can use the simulated distribution of end-period portfolio values to infer the VaR or ETL.

Tool No. 10 discusses the forecasting of volatilities, covariances and correlations. This is one of the most important subjects in modern risk measurement, and is critical to derivatives pricing, hedging, and VaR and ETL estimation. The focus of our discussion is the estimation of volatilities, in which we go through each of four main approaches to this problem: historical estimation, exponentially weighted moving average (EWMA) estimation, GARCH estimation, and implied volatility estimation. The treatment of covariances and correlations parallels that of volatilities, and is followed by a brief discussion of the issues involved with the estimation of variance–covariance and correlation matrices.

Finally, Tool No. 11 deals with the often misunderstood issue of dependency between risky variables. The most common way of representing dependency is by means of the linear correlation coefficient, but this is only appropriate in limited circumstances (i.e., to be precise, when the risky variables are elliptically distributed, which includes their being normally distributed as a special case). In more general circumstances, we should represent dependency in terms of copulas, which are functions that combine the marginal distributions of different variables to produce a multivariate distribution function that takes account of their dependency structure. There are many different copulas, and we need to choose a copula function appropriate for the problem at hand. We then consider how to estimate copulas, and how to use copulas to estimate VaR.

Acknowledgements

It is real pleasure to acknowledge those who have contributed in one way or another to this book. To begin with, I should like to thank Barry Schachter for his excellent website, www.gloriamundi.org, which was my primary source of research material. I thank Naomi Fernandes and the The MathWorks, Inc., for making MATLAB available to me through their authors' program. I thank Christian Bauer, David Blake, Carlos Blanco, Andrew Cairns, Marc de Ceuster, Jon Danielsson, Kostas Giannopoulos, Paul Glasserman, Glyn Holton, Imad Moosa, and Paul Stefiszyn for their valuable comments on parts of the draft manuscript and/or other contributions, I thank Mark Garman for permission to include Figures 8.2 and 8.3, and Peter Urbani for allowing me to include some of his Excel software with the CD. I also thank the Wiley team — Sam Hartley, Sarah Lewis, Carole Millett and, especially, Sam Whittaker — for many helpful inputs. I should also like to thank participants in the Dutch National Bank's Capital Markets Program and seminar participants at the Office of the Superintendent of Financial Institutions in Canada for allowing me to test out many of these ideas on them, and for their feedback.

In addition, I would like to thank my colleagues and students at the Centre for Risk and Insurance Studies (CRIS) and also in the rest of Nottingham University Business School, for their support and feedback. I also thank many friends for their encouragement and support over the years: particularly Mark Billings, Dave Campbell, David and Yabing Fisher, Ian Gow, Duncan Kitchin, Anneliese Osterspey, Dave and Frances Owen, Sheila Richardson, Stan and Dorothy Syznkaruk, and Basil and Margaret Zafiriou. Finally, as always, my greatest debts are to my family — to my mother, Maureen, my brothers Brian and Victor, and most of all, to my wife Mahjabeen and my daughters Raadhiyah and Safiah — for their love and unfailing support, and their patience. I would therefore like to dedicate this book to Mahjabeen and the girls. I realise of course that other authors' families get readable books dedicated to them, and all I have to offer is another soporific statistical tome. But maybe next time I will take their suggestion and write a novel instead. On second thoughts, perhaps not.

1
The Risk Measurement Revolution

Financial risk is the prospect of financial loss — or gain — due to unforeseen changes in underlying risk factors. In this book we are concerned with the measurement of one particular form of financial risk — namely, market risk, or the risk of loss (or gain) arising from unexpected changes in market prices (e.g., such as security prices) or market rates (e.g., such as interest or exchange rates). Market risks, in turn, can be classified into interest rate risks, equity risks, exchange rate risks, commodity price risks, and so on, depending on whether the risk factor is an interest rate, a stock price, or whatever. Market risks can also be distinguished from other forms of financial risk, most especially credit risk (or the risk of loss arising from the failure of a counterparty to make a promised payment) and operational risk (or the risk of loss arising from the failures of internal systems or the people who operate in them).

The theory and practice of risk management — and, included within that, risk measurement — have developed enormously since the pioneering work of Harry Markowitz in the 1950s. The theory has developed to the point where risk management/measurement is now regarded as a distinct sub-field of the theory of finance, and one that is increasingly taught as a separate subject in the more advanced master's and MBA programmes in finance. The subject has attracted a huge amount of intellectual energy, not just from finance specialists but also from specialists in other disciplines who are attracted to it — as illustrated by the large number of ivy league theoretical physics PhDs who now go into finance research, attracted not just by the high salaries but also by the challenging intellectual problems it poses.

1.1 CONTRIBUTORY FACTORS

1.1.1 A Volatile Environment

One factor behind the rapid development of risk management was the high level of instability in the economic environment within which firms operated. A volatile environment exposes firms to greater financial risk, and therefore provides an incentive for firms to find new and better ways of managing this risk. The volatility of the economic environment is reflected in various factors:

- *Stock market volatility*. Stock markets have always been volatile, but sometimes extremely so: for example, on October 19, 1987, the Dow Jones fell 23% and in the process knocked off over $1 trillion in equity capital; and from July 21 through August 31, 1998, the Dow Jones lost 18% of its value. Other western stock markets have experienced similar falls, and some Asian ones have experienced much worse ones (e.g., the South Korean stock market lost over half of its value during 1997).
- *Exchange rate volatility*. Exchange rates have been volatile ever since the breakdown of the Bretton Woods system of fixed exchange rates in the early 1970s. Occasional exchange rate crises have also led to sudden and significant exchange rate changes, including — among many others — the ERM devaluations of September 1992, the problems of the peso in

1994, the East Asian currency problems of 1997–98, the rouble crisis of 1998 and Brazil in 1999.
- *Interest rate volatility*. There have been major fluctuations in interest rates, with their attendant effects on funding costs, corporate cash flows and asset values. For example, the Fed Funds rate, a good indicator of short-term market rates in the US, approximately doubled over 1994.
- *Commodity market volatility*. Commodity markets are notoriously volatile, and commodity prices often go through long periods of apparent stability and then suddenly jump by enormous amounts: for instance, in 1990, the price of West Texas Intermediate crude oil rose from a little over $15 a barrel to around $40 a barrel. Some commodity prices (e.g., electricity prices) also show extremely pronounced day-to-day and even hour-to-hour volatility.

1.1.2 Growth in Trading Activity

Another factor contributing to the transformation of risk management is the huge increase in trading activity since the late 1960s. The average number of shares traded per day in the New York Stock Exchange has grown from about 3.5m in 1970 to around 100m in 2000; and turnover in foreign exchange markets has grown from about a billion dollars a day in 1965 to $1,210 billion in April 2001.[1] There have been massive increases in the range of instruments traded over the past two or three decades, and trading volumes in these new instruments have also grown very rapidly. New instruments have been developed in offshore markets and, more recently, in the newly emerging financial markets of Eastern Europe, China, Latin America, Russia, and elsewhere. New instruments have also arisen for assets that were previously illiquid, such as consumer loans, commercial and industrial bank loans, mortgages, mortgage-based securities and similar assets, and these markets have grown very considerably since the early 1980s.

There has also been a phenomenal growth of derivatives activity. Until 1972 the only derivatives traded were certain commodity futures and various forwards and over-the-counter (OTC) options. The Chicago Mercantile Exchange then started trading foreign currency futures contracts in 1972, and in 1973 the Chicago Board Options Exchange started trading equity call options. Interest-rate futures were introduced in 1975, and a large number of other financial derivatives contracts were introduced in the following years: swaps and exotics (e.g., swaptions, futures on interest rate swaps, etc.) then took off in the 1980s, and catastrophe, credit, electricity and weather derivatives in the 1990s. From negligible amounts in the early 1970s, the daily notional amounts turned over in derivatives contracts grew to nearly $2,800 billion by April 2001.[2] However, this figure is misleading, because notional values give relatively little indication of what derivatives contracts are really worth. The true size of derivatives trading is better represented by the replacement cost of outstanding derivatives contracts, and these are probably no more than 4% or 5% of the notional amounts involved. If we measure size by replacement cost rather than notional principals, the size of the daily turnover in the derivatives market in 2001 was therefore around $126 billion — which is still not an inconsiderable amount.

[1] The latter figure is from Bank for International Settlements (2001, p. 1).
[2] Bank for International Settlements (2001, p. 9).

1.1.3 Advances in Information Technology

A third contributing factor to the development of risk management was the rapid advance in the state of information technology. Improvements in IT have made possible huge increases in both computational power and the speed with which calculations can be carried out. Improvements in computing power mean that new techniques can be used (e.g., such as computer-intensive simulation techniques) to enable us to tackle more difficult calculation problems. Improvements in calculation speed then help make these techniques useful in real time, where it is often essential to get answers quickly.

This technological progress has led to IT costs falling by about 25–30% a year over the past 30 years or so. To quote Guldimann:

> Most people know that technology costs have dropped rapidly over the years but few realise how steep and continuous the fall has been, particularly in hardware and data transmission. In 1965, for example, the cost of storing one megabyte of data (approximately the equivalent of the content of a typical edition of the *Wall Street Journal*) in random access memory was about $100,000. Today it is about $20. By 2005, it will probably be less than $1.
>
> The cost of transmitting electronic data has come down even more dramatically. In 1975, it cost about $10,000 to send a megabyte of data from New York to Tokyo. Today, it is about $5. By 2005, it is expected to be about $0.01. And the cost of the processor needed to handle 1 million instructions a second has declined from about $1 million in 1965 to $1.50 today. By 2005, it is expected to drop to a few cents. (All figures have been adjusted for inflation.) (Guldimann (1996, p. 17))

Improvements in computing power, increases in computing speed, and reductions in computing costs have thus come together to transform the technology available for risk management. Decision-makers are no longer tied down to the simple ‘back of the envelope’ techniques that they had to use earlier when they lacked the means to carry out more complex calculations. They can now use sophisticated algorithms programmed into computers to carry out real-time calculations that were not possible before. The ability to carry out such calculations then creates a whole new range of risk measurement and risk management possibilities.

1.2 RISK MEASUREMENT BEFORE VaR

To understand recent developments in risk measurement, we need first to appreciate the more traditional risk measurement tools.

1.2.1 Gap Analysis

One common approach was (and, in fact, still is) gap analysis, which was initially developed by financial institutions to give a simple, albeit crude, idea of interest-rate risk exposure.[3] Gap analysis starts with the choice of an appropriate horizon period — 1 year, or whatever. We then determine how much of our asset or liability portfolio will re-price within this period, and the amounts involved give us our rate-sensitive assets and rate-sensitive liabilities. The gap is the difference between these, and our interest-rate exposure is taken to be the change in net interest income that occurs in response to a change in interest rates. This in turn is assumed to be equal

[3] For more on gap analysis, see, e.g., Sinkey (1992, ch. 12).

to the gap times the interest-rate change:

$$\Delta NII = (GAP)\Delta r \tag{1.1}$$

where ΔNII is the change in net interest income and Δr is the change in interest rates.

Gap analysis is fairly simple to carry out, but has its limitations: it only applies to on-balance sheet interest-rate risk, and even then only crudely; it looks at the impact of interest rates on income, rather than on asset or liability values; and results can be sensitive to the choice of horizon period.

1.2.2 Duration Analysis

Another method traditionally used by financial institutions for measuring interest-rate risks is duration analysis.[4] The (Macaulay) duration D of a bond (or any other fixed-income security) can be defined as the weighted average term to maturity of the bond's cash flows, where the weights are the present value of each cash flow relative to the present value of all cash flows:

$$D = \sum_{i=1}^{n} [i \times PVCF_i] \Big/ \sum_{i=1}^{n} PVCF_i \tag{1.2}$$

where $PVCF_i$ is the present value of the period i cash flow, discounted at the appropriate spot period yield. The duration measure is useful because it gives an approximate indication of the sensitivity of a bond price to a change in yield:

$$\%\,\text{Change in bond price} \approx -D \times \Delta y/(1+y) \tag{1.3}$$

where y is the yield and Δy the change in yield. The bigger the duration, the more the bond price changes in response to a change in yield. The duration approach is very convenient because duration measures are easy to calculate and the duration of a bond portfolio is a simple weighted average of the durations of the individual bonds in that portfolio. It is also better than gap analysis because it looks at changes in asset (or liability) values, rather than just changes in net income.

However, duration approaches have similar limitations to gap analysis: they ignore risks other than interest-rate risk; they are crude,[5] and even with various refinements to improve accuracy,[6] duration-based approaches are still inaccurate relative to more recent approaches to fixed-income analysis (e.g., such as HJM models). Moreover, the main reason for using duration approaches in the past — their (comparative) ease of calculation — is no longer of much significance, since more sophisticated models can now be programmed into microcomputers to give their users more accurate answers rapidly.

[4] For more on duration approaches, see, e.g., Fabozzi (1993, ch. 11 and 12) or Tuckman (1995, ch. 11–13).

[5] They are crude because they only take a first-order approximation to the change in the bond price, and because they implicitly presuppose that any changes in the yield curve are parallel ones (i.e., all yields across the maturity spectrum change by the same amount). Duration-based hedges are therefore inaccurate against yield changes that involve shifts in the slope of the yield curve.

[6] There are two standard refinements. (1) We can take a second-order rather than a first-order approximation to the bond price change. The second-order term — known as convexity — is related to the change in duration as yield changes, and this duration–convexity approach gives us a better approximation to the bond price change as the yield changes. (For more on this approach, see, e.g., Fabozzi (1993, ch. 12) or Tuckman (1995, ch. 11).) However, the duration–convexity approach generally only gives modest improvements in accuracy. (2) An alternative refinement is to use key rate durations: if we are concerned about shifts in the yield curve, we can construct separate duration measures for yields of specified maturities (e.g., short-term and long-term yields); these would give us estimates of our exposure to changes in these specific yields and allow us to accommodate non-parallel shifts in the yield curve. For more on this key rate duration approach, see Ho (1992) or Tuckman (1995, ch. 13).

1.2.3 Scenario Analysis

A third approach is scenario analysis (or 'what if' analysis), in which we set out different scenarios and investigate what we stand to gain or lose under them. To carry out scenario analysis, we select a set of scenarios — or paths describing how relevant variables (e.g., stock prices, interest rates, exchange rates, etc.) might evolve over a horizon period. We then postulate the cash flows and/or accounting values of assets and liabilities as they would develop under each scenario, and use the results to come to a view about our exposure.

Scenario analysis is not easy to carry out. A lot hinges on our ability to identify the 'right' scenarios, and there are relatively few rules to guide us when selecting them. We need to ensure that the scenarios we examine are reasonable and do not involve contradictory or excessively implausible assumptions, and we need to think through the interrelationships between the variables involved.[7] We also want to make sure, as best we can, that we have all the main scenarios covered. Scenario analysis also tells us nothing about the likelihood of different scenarios, so we need to use our judgement when assessing the practical significance of different scenarios. In the final analysis, the results of scenario analyses are highly subjective and depend to a very large extent on the skill or otherwise of the analyst.

1.2.4 Portfolio Theory

A somewhat different approach to risk measurement is provided by portfolio theory.[8] Portfolio theory starts from the premise that investors choose between portfolios on the basis of their expected return, on the one hand, and the standard deviation (or variance) of their return, on the other.[9] The standard deviation of the portfolio return can be regarded as a measure of the portfolio's risk. Other things being equal, an investor wants a portfolio whose return has a high expected value and a low standard deviation. These objectives imply that the investor should choose a portfolio that maximises expected return for any given portfolio standard deviation or, alternatively, minimises standard deviation for any given expected return. A portfolio that meets these conditions is efficient, and a rational investor will always choose an efficient portfolio. When faced with an investment decision, the investor must therefore determine the set of efficient portfolios and rule out the rest. Some efficient portfolios will have more risk than others, but the more risky ones will also have higher expected returns. Faced with the set of efficient portfolios, the investor then chooses one particular portfolio on the basis of his or her own preferred trade-off between risk and expected return. An investor who is very averse to risk will choose a safe portfolio with a low standard deviation and a low expected return, and an

[7] We will often want to examine scenarios that take correlations into account as well (e.g., correlations between interest-rate and exchange-rate risks), but in doing so, we need to bear in mind that correlations often change, and sometimes do so at the most awkward times (e.g., during a market crash). Hence, it is often good practice to base scenarios on relatively conservative assumptions that allow for correlations to move against us.

[8] The origin of portfolio theory is usually traced to the work of Markowitz (1952, 1959). Later scholars then developed the Capital Asset Pricing Model (CAPM) from the basic Markowitz framework. However, I believe the CAPM — which I interpret to be portfolio theory combined with the assumptions that everyone is identical and that the market is in equilibrium — was an unhelpful digression and that the current discredit into which it has fallen is justified. (For the reasons behind this view, I strongly recommend Frankfurter's withering assessment of the rise and fall of the CAPM empire (Frankfurter (1995)).) That said, in going over the wreckage of the CAPM, it is also important not to lose sight of the tremendous insights provided by portfolio theory (i.e., *à la* Markowitz). I therefore see the way forward as building on portfolio theory (and, indeed, I believe that much of what is good in the VaR literature does exactly that) whilst throwing out the CAPM.

[9] This framework is often known as the mean–variance framework, because it implicitly presupposes that the mean and variance (or standard deviation) of the return are sufficient to guide investors' decisions. In other words, investors are assumed not to need information about higher order moments of the return probability density function, such as the skewness or kurtosis coefficients.

investor who is less risk averse will choose a more risky portfolio with a higher expected return.

One of the key insights of portfolio theory is that the risk of any individual asset is not the standard deviation of the return to that asset, but rather the extent to which that asset contributes to overall portfolio risk. An asset might be very risky (i.e., have a high standard deviation) when considered on its own, and yet have a return that correlates with the returns to other assets in our portfolio in such a way that acquiring the new asset adds nothing to the overall portfolio standard deviation. Acquiring the new asset would then be riskless, even though the asset held on its own would still be risky. The moral of the story is that the extent to which a new asset contributes to portfolio risk depends on the correlation or covariance of its return with the returns to the other assets in our portfolio — or, if one prefers, the beta, which is equal to the covariance between the return to asset i and the return to the portfolio, r_p, divided by the variance of the portfolio return. The lower the correlation, other things being equal, the less the asset contributes to overall risk. Indeed, if the correlation is sufficiently negative, it will offset existing risks and lower the portfolio standard deviation.

Portfolio theory provides a useful framework for handling multiple risks and taking account of how those risks interact with each other. It is therefore of obvious use to — and is in fact widely used by — portfolio managers, mutual fund managers and other investors. However, it tends to run into problems over data. The risk-free return and the expected market return are not too difficult to estimate, but estimating the betas is often more problematic. Each beta is specific not only to the individual asset to which it belongs, but also to our current portfolio. To estimate a beta coefficient properly, we need data on the returns to the new asset and the returns to all our existing assets, and we need a sufficiently long data set to make our statistical estimation techniques reliable. The beta also depends on our existing portfolio and we should, in theory, re-estimate all our betas every time our portfolio changes. Using the portfolio approach can require a considerable amount of data and a substantial amount of ongoing work.

In practice users often wish to avoid this burden, and in any case they sometimes lack the data to estimate the betas accurately in the first place. Practitioners are then tempted to seek a short-cut, and work with betas estimated against some hypothetical market portfolio. This leads them to talk about *the* beta for an asset, as if the asset had only a single beta. However, this short-cut only gives us good answers if the beta estimated against the hypothetical market portfolio is close to the beta estimated against the portfolio we actually hold, and in practice we seldom know whether it is.[10] If the two portfolios are sufficiently different, the ‘true’ beta (i.e., the beta measured against our actual portfolio) might be very different from the hypothetical beta we are using.[11]

[10] There are also other problems. (1) If we wish to use this short-cut, we have relatively little firm guidance on what the hypothetical portfolio should be. In practice, investors usually use some ‘obvious’ portfolio such as the basket of shares behind a stock index, but we never really know whether this is a good proxy for the CAPM market portfolio or not. It is probably not. (2) Even if we pick a good proxy for the CAPM market portfolio, it is still very doubtful that *any* such portfolio will give us good results (see, e.g., Markowitz (1992, p. 684)). If we wish to use proxy risk estimates, there is a good argument that we should abandon single-factor models in favour of multi-factor models that can mop up more systematic risks. This leads us to the arbitrage pricing theory (APT) of Ross (1976). However, the APT has its own problems: we can’t easily identify the risk factors, and even if we did identify them, we still don’t know whether the APT will give us a good proxy for the systematic risk we are trying to proxy.

[11] We can also measure risk using statistical approaches applied to equity, FX, commodity and other risks, as well as interest-rate risks. The idea is that we postulate a measurable relationship between the exposure variable we are interested in (e.g., the loss/gain on our bond or FX portfolio or whatever) and the factors that we think influence that loss or gain. We then estimate the parameters of this relationship by an appropriate econometric technique, and the parameter estimates give us an idea of our risk exposures. This approach is limited by the availability of data (i.e., we need enough data to estimate the relevant parameters) and by linearity assumptions, and there can be problems caused by misspecification and instability in estimated statistical relationships.