

COHERENT STRESS TESTING

A Bayesian Approach to the Analysis of Financial Stress



RICCARDO REBONATO

Coherent Stress Testing

A Bayesian Approach to the Analysis of Financial Stress

Riccardo Rebonato



 **WILEY**

A John Wiley and Sons, Ltd., Publication

This edition first published 2010
© 2010 Riccardo Rebonato

Registered office

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, United Kingdom

For details of our global editorial offices, for customer services and for information about how to apply for permission to reuse the copyright material in this book please see our website at www.wiley.com.

The right of the author to be identified as the author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book. This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Library of Congress Cataloging-in-Publication Data

Rebonato, Riccardo.

Coherent stress testing : a Bayesian approach / Riccardo Rebonato.

p. cm.

ISBN 978-0-470-66601-2

1. Risk management. 2. Probabilities. 3. Bayesian statistical decision theory. I. Title.

HD61.R42 2010

658.15'501519542 – dc22

2010005778

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

ISBN 978-0-470-66601-2

Typeset in 10/12 Times-Roman by Laserwords Private Limited, Chennai, India

Printed in Great Britain by CPI Antony Rowe, Chippenham, Wiltshire

Acknowledgements

It gives me great pleasure to acknowledge the support I received while writing this book.

Dr Keating provided very insightful comments and pointed out some 'loose thinking'. I am very grateful for that.

Many of my colleagues have helped me a lot, by challenging my thoughts, pointing out what did and did not work, and suggesting how I could improve the approach. I am sure that I will forget many, but I certainly extend my thanks to Dr Gary Dunn, Dr Ron Keating, Dr Ronnie Barnes, Mr Daniel Burns, Dr Michael Smith, Mr Paul Fairhurst, Dr Jeremy Broughton, Dr Ed Hayes, Mr Craig Schor, Dr Tom Connor and, for his unflinching criticism, Dr Stephen Laughton. Above all, however, I thank Dr Jan Kwiatkowski, who proposed an earlier version of the Linear Programming approach that I describe in Chapter 10. His contribution has extended well beyond this technical suggestion, as over the years he has become my Bayesian mentor.

I have greatly benefited from discussions that I have had with regulators, at the Boston Fed, at the FSA and at the MSA. I am grateful for the time they spent discussing my ideas on stress testing.

I have presented parts of the material in this book at several conferences in the US and in Europe, and therefore I have received extremely useful and insightful comments from many delegates. The book would be much the poorer without their suggestions: thank you.

At John Wiley & Sons, Ltd, Caitlin Cornish first and then Pete Baker have shown from the start great enthusiasm for the project. I am very grateful for this.

Last but not least, my wife and my parents have given me continuous support and encouragement. In a way, it is to them that I owe my greatest debt of gratitude.

Despite all this help, I am sure that there are still many mistakes and imperfections in the book. I am fully and solely responsible for these.

Contents

Acknowledgements	xi
1 Introduction	1
1.1 Why We Need Stress Testing	1
1.2 Plan of the Book	5
1.3 Suggestions for Further Reading	6
I Data, Models and Reality	7
2 Risk and Uncertainty – or, Why Stress Testing is <i>Not</i> Enough	9
2.1 The Limits of Quantitative Risk Analysis	9
2.2 Risk or Uncertainty?	10
2.3 Suggested Reading	13
3 The Role of Models in Risk Management and Stress Testing	15
3.1 How Did We Get Here?	16
3.2 Statement of the Two Theses of this Chapter	17
3.3 Defence of the First Thesis (Centrality of Models)	18
3.3.1 Models as Indispensable Interpretative Tools	18
3.3.2 The Plurality-of-Models View	21
3.4 Defence of the Second Thesis (Coordination)	23
3.4.1 Traders as Agents	23
3.4.2 Agency Brings About Coordination	25
3.4.3 From Coordination to Positive Feedback	26
3.5 The Role of Stress and Scenario Analysis	27
3.6 Suggestions for Further Reading	29
4 What Kind of Probability Do We Need in Risk Management?	31
4.1 Frequentist versus Subjective Probability	31
4.2 Tail Co-dependence	36
4.3 From Structural Models to Co-dependence	38
4.4 Association or Causation?	39
4.5 Suggestions for Further Reading	42

II	The Probabilistic Tools and Concepts	45
5	Probability with Boolean Variables I: Marginal and Conditional Probabilities	47
5.1	The Set-up and What We are Trying to Achieve	47
5.2	(Marginal) Probabilities	50
5.3	Deterministic Causal Relationship	53
5.4	Conditional Probabilities	55
5.5	Time Ordering and Causation	56
5.6	An Important Consequence: Bayes' Theorem	57
5.7	Independence	58
5.8	Two Worked-Out Examples	59
5.8.1	Dangerous Running	59
5.8.2	Rare and Even More Dangerous Diseases	61
5.9	Marginal and Conditional Probabilities: A Very Important Link	62
5.10	Interpreting and Generalizing the Factors x_i^k	65
5.11	Conditional Probability Maps	67
6	Probability with Boolean Variables II: Joint Probabilities	71
6.1	Conditioning on More Than One Event	71
6.2	Joint Probabilities	73
6.3	A Remark on Notation	75
6.4	From the Joint to the Marginal and the Conditional Probabilities	76
6.5	From the Joint Distribution to Event Correlation	77
6.6	From the Conditional and Marginal to the Joint Probabilities?	83
6.7	Putting Independence to Work	84
6.8	Conditional Independence	86
6.9	Obtaining Joint Probabilities with Conditional Independence	88
6.10	At a Glance	89
6.11	Summary	90
6.12	Suggestions for Further Reading	90
7	Creating Probability Bounds	93
7.1	The Lay of the Land	93
7.2	Bounds on Joint Probabilities	93
7.3	How Tight are these Bounds in Practice?	96
8	Bayesian Nets I: An Introduction	99
8.1	Bayesian Nets: An Informal Definition	99
8.2	Defining the Structure of Bayesian Nets	101
8.3	More About Conditional Independence	104
8.4	What Goes in the Conditional Probability Tables?	106
8.5	Useful Relationships	107
8.6	A Worked-Out Example	109

8.7	A Systematic Approach	111
8.8	What Can We Do with Bayesian Nets?	113
8.8.1	Unravelling the Causal Structure	113
8.8.2	Estimating the Joint Probabilities	114
8.9	Suggestions for Further Reading	115
9	Bayesian Nets II: Constructing Probability Tables	117
9.1	Statement of the Problem	117
9.2	Marginal Probabilities – First Approach	118
9.2.1	Starting from a Fixed Probability	119
9.2.2	Starting from a Fixed Magnitude of the Move	120
9.3	Marginal Probabilities – Second Approach	120
9.4	Handling Events of Different Probability	122
9.5	Conditional Probabilities: A Reasonable Starting Point	123
9.6	Conditional Probabilities: Checks and Constraints	125
9.6.1	Necessary Conditions	125
9.6.2	Triplet Conditions	126
9.6.3	Independence	127
9.6.4	Deterministic Causation	127
9.6.5	Incompatibility of Events	128
9.7	Internal Compatibility of Conditional Probabilities: The Need for a Systematic Approach	129
III	Applications	131
10	Obtaining a Coherent Solution I: Linear Programming	133
10.1	Plan of the Work Ahead	133
10.2	Coherent Solution with Conditional Probabilities Only	135
10.3	The Methodology in Practice: First Pass	141
10.4	The CPU Cost of the Approach	144
10.5	Illustration of the Linear Programming Technique	144
10.6	What Can We Do with this Information?	149
10.6.1	Extracting Information with Conditional Probabilities Only	149
10.6.2	Extracting Information with Conditional and Marginal Probabilities	151
11	Obtaining a Coherent Solution II: Bayesian Nets	155
11.1	Solution with Marginal and n -conditioned Probabilities	156
11.1.1	Generalizing the Results	164
11.2	An ‘Automatic’ Prescription to Build Joint Probabilities	165
11.3	What Can We Do with this Information?	167
11.3.1	Risk-Adjusting Returns	168

IV Making It Work In Practice	171
12 Overcoming Our Cognitive Biases	173
12.1 Cognitive Shortcomings and Bounded Rationality	174
12.1.1 How Pervasive are Cognitive Shortcomings?	175
12.1.2 The Social Context	175
12.1.3 Adaptiveness	176
12.2 Representativeness	178
12.3 Quantification of the Representativeness Bias	181
12.4 Causal/Diagnostic and Positive/Negative Biases	182
12.5 Conclusions	184
12.6 Suggestions for Further Reading	185
13 Selecting and Combining Stress Scenarios	187
13.1 Bottom Up or Top Down?	187
13.2 Relative Strengths and Weaknesses of the Two Approaches	187
13.3 Possible Approaches to a Top-Down Analysis	190
13.4 Sanity Checks	191
13.5 How to Combine Stresses – Handling the Dimensionality Curse	192
13.6 Combining the Macro and Bottom-Up Approaches	194
14 Governance	197
14.1 The Institutional Aspects of Stress Testing	197
14.1.1 Transparency and Ease of Use	197
14.1.2 Challenge by Non-specialists	198
14.1.3 Checks for Completeness	198
14.1.4 Interactions among Different Specialists	199
14.1.5 Auditability of the Process and of the Results	201
14.2 Lines of Criticism	201
14.2.1 The Role of Subjective Inputs	202
14.2.2 The Complexity of the Stress-testing Process	203
Appendix A Simple Introduction to Linear Programming	205
A.1 Plan of the Appendix	205
A.2 Linear Programming – A Refresher	205
A.3 The Simplex Method	208
References	213
Index	217

Chapter 1

Introduction

[Under uncertainty] there is no scientific basis on which to form any calculable probability whatever. We simply don't know. Nevertheless, the necessity for action and for decision compels us as practical men to do our best to overlook this awkward fact and to behave exactly as we should if we had behind us [...] a series of prospective advantages and disadvantages, each multiplied by its appropriate probability waiting to be summed....

Robertson (1936)

1.1 Why We Need Stress Testing

Why a book about stress testing? And why a book about stress testing *now*? Stress testing has been part of the risk manager's toolkit for decades.¹ What justifies the renewed interest from practitioners and regulators² for a risk management tool that, truth be told, has always been the poor relation in the family of analytical techniques to control risk?³ And why has stress testing so far been regarded as a second-class citizen?

Understanding the reason for the renewed interest is simple: the financial crisis of 2007–2008–2009 has shown with painful clarity the limitations of the purely statistical techniques (such as Value at Risk (VaR) or Economic Capital) that were supposed to provide the cornerstones of the financial edifice. In the year and a half starting with July 2007, events of once-in-many-thousand-years rarity kept on occurring with disconcerting regularity. Only a risk manager of Stalinist dogmatism could have lived through these events and 'kept the faith'. Clearly, something more – or, rather, something *different* – had to be done. But what? And what analytical tools should we employ to fix the problem?

¹See, e.g., Sorge (2004) for a review of stress testing methodologies put together before the sub-prime crisis, or Alexander and Sheedy (2008) for a recent approach firmly rooted in the statistical analysis of past data.

Stress testing is dealt with, often in a rather cursory manner, in most standard texts on risk management. I assume in this book that the reader is broadly familiar with current stress-testing practice. By quoting from the work of Aragonés, Blanco and Dowd (2001), I explain in this chapter a common shortcoming of these approaches. Indeed, this shortcoming is one of the motivations of the approach I propose.

²See, e.g., BIS (2009).

³"Those who can, do, and those who can't, do stress testing", a proud risk quant told me less than two years ago.

'Stress testing' has become the stock answer to these questions. But the unease and suspicion with which this technique has been regarded has not melted away. The frog has not been kissed (yet) into a handsome prince. The current attitude seems one of resigned acceptance of a *faute-de-mieux* measure of risk: a far cry from an enthusiastic embrace of a new and powerful analytical tool. Two cheers, the mood seems to be, for stress testing. Can we do better? And why has stress testing been regarded as such an ungainly frog in the first place?

If by stress testing we mean the assessment of very severe financial losses arrived at without heavy reliance on statistical techniques, but by deploying instead a large dose of subjective judgement, some answers to the latter question are not difficult to see. Rather than paraphrasing, I would like to quote extensively from an article by Aragonés, Blanco and Dowd (2001), who put their fingers exactly on the problem:

... traditional stress testing is done on a stand-alone basis, and the results of stress tests are evaluated side-by-side with the results of traditional market risk (or VaR) models. This creates problems for risk managers, who then have to choose which set of risk exposures to 'believe'. [R]isk managers often don't know whether to believe their stress test results, because the stress tests exercises give them no idea of how likely or unlikely stress-test scenarios might be....

And again:

A related problem is that the results of stress tests are difficult to interpret because they give us no idea of the probabilities of the events concerned, and in the absence of such information we often don't know what to do with them. Suppose for instance that stress testing reveals that our firm will go bust under a particular scenario. Should we act on this information? The only answer is that we can't say. If the scenario is very likely, we would be very unwise not to act on it. But if the scenario was extremely unlikely, then it becomes almost irrelevant, because we would not usually expect management to take expensive precautions against events that may be too improbable to worry about. So the extent to which our results matter or not depends on unknown probabilities. As Berkowitz [1999] nicely puts it, this absence of probabilities puts 'stress testing in a statistical purgatory. We have some loss numbers, but who is to say whether we should be concerned about them?'

The result of this state of affairs is not pretty: we are left with

... two sets of separate risk estimates – probabilistic estimates (e.g., such as VaR), and the loss estimates produced by stress tests – and no way of combining them. How can we combine a probabilistic risk estimate with an estimate that such-and-such a loss will occur if such-and-such happens? The answer, of course, is that we can't. We therefore have to work with these estimates more or less independently of each other, and the best we can do is use one set of estimates to check for prospective losses that the other might have underrated or missed...

In modern finance, risk and reward are supposed to be two sides of the same coin. Risk is 'priced' in terms of expected return by assigning probabilities⁴ to outcomes. But when

⁴Note that I did not say 'assigning objective probabilities known and agreed upon by all the market participants'. See Shefrin (2008) for a discussion of asset pricing under heterogeneous beliefs about probabilities. See also the discussion in Chapter 3.

it comes to extreme events, absent of any probabilistic assessment, we don't know how to 'price' the outcomes of stress testing. And if our confidence in assigning a probability to extremely rare events has been terminally shaken by the recent market events,⁵ the state of impasse seems inescapable.

Perhaps there is hope – and it is exactly this ray of hope that this book pursues. First of all, 'probabilistic statement' need not be equated with 'frequentist (i.e., purely data-driven) probabilistic statement'. As I discuss in Chapter 4, there is a different way of looking at probability that takes into account, but is not limited to, the pure analysis of data. I maintain (and I have argued at length elsewhere⁶) that the subjective view of probability is every bit as 'respectable' as the purely-data-driven (frequentist) one. I also believe that it is much better suited to the needs of risk management.

This view, while not mainstream, is not particularly new, especially in the context of financial risk management – see, e.g., Berkowitz (1999). However, the subjective approach brings about an insidious problem. It is all well and good to assign subjective probabilities to stand-alone events. But, if we want to escape from Berkowitz's purgatory, we will have to do more. We will have to combine different stress scenarios, with different subjective probabilities, into an overall coherent, albeit approximate, stress loss number at a given confidence level (or, perhaps, into a whole stress loss distribution). How is one to do that? How is one to provide *subjectively* these co-dependences – and *tail* co-dependences to boot? How is one to ensure that the subjectively-assigned probabilities are reasonable, let alone feasible (i.e., mathematically possible and self-consistent)?

This book offers two routes to escape this purgatorial dilemma. The first is the acknowledgement that the risk manager can only make sense of data on the basis of a model (or of competing models) of reality. A risk manager, for instance, should have a conception of the direction of causation between different events: does a dramatic fall in equity prices 'cause' an increase in equity implied volatilities? Or is it an increase in implied volatility that 'causes' a dramatic fall in equity prices? The answer, at least in this case, may seem obvious. Unfortunately, correlations, and even conditional probabilities, contain no information about the direction of causation. Yet, this information about causation, even if imperfect, is powerful. It is ignored in the frequentist approach at a great loss for the risk manager. Speaking about the sciences in general, Pearl (2009) points out that there is 'no greater impediment to scientific progress than the prevailing practice of focusing all of our mathematical resources on statistical and probabilistic inferences'. I believe that exactly the same applies in the area of quantitative risk management.

If one is prepared to 'stick one's neck out' and make some reasonable assumptions about the direction of the arrow of causation, just like Dante one can begin to glimpse some light filtering through the thick trees of the *selva oscura*. In the case of stress testing, the route to salvation is via the provision of information that is not 'contained in the data'.

Sure enough, even if one can provide this extra information not all is plain sailing. Organizing one's understanding about how the world might work into a coherent and tractable analytical probabilistic framework is not an easy task. Fortunately, if one is prepared to make some reasonable approximations, there are powerful *and intuitive* techniques that can offer great help in building plausible and mathematically self-consistent joint distributions

⁵Of course, there are techniques such as Extreme Value Theory (EVT) that promise to give us a quantitative probabilistic glimpse of extremely rare events. I discuss in Chapter 4 my reservations about this and related techniques.

⁶See Rebonato (2007).

of the stress losses that have been identified. These technical tools (Bayesian networks and Linear Programming) have been well known for a long time, but their application to risk management problems, and to stress testing in particular, has been hesitant at best. This is a pity, because I believe that they are not only powerful and particularly well suited to the problem at hand, but also extremely intuitively appealing. And in Section 2.2 of the next chapter I will highlight how important appeal to intuition can be if the recommendations of the risk managers are to be acted upon (as opposed to ‘confined to a stress report’).

Once we accept that we can, approximately but meaningfully, associate stress events with a probabilistic assessment of their likelihood, the questions that opened this chapter begin to find a compelling answer. We need stress testing, and we need stress testing *now*, because the purely-data-based statistical techniques we have been using have proven unequal to the task when it really mattered. Perhaps the real question should have been instead: ‘How can we do *without* stress testing?’

Of course, there is a lot more to risk management than predicting the probability of losses large and small. But, even if we look at the management of financial risk through the highly reductive prism of analysing the likelihood of losses, there still is no one single goal for the risk manager. For instance, estimating the kind of profit-and-loss variability that can be expected on a weekly or monthly basis has value and importance. Ensuring that a business line or trading desk effectively ‘diversifies’ the revenue stream from other existing lines of activity *under normal market conditions* is also obviously important. So is estimating the income variability or the degree of diversification that can be expected from a portfolio of businesses over a business cycle. And recent events have shown the importance of ensuring that a set of business activities do not endanger the survival of a financial institution even under exceptional market conditions. These are all important goals for a risk manager. But it would be extraordinary if the same analytical tools could allow the risk manager to handle all these problems – problems, that is, whose solution hinges on the estimation of probabilities of events that should occur, on average, from once every few weeks to once in several decades. This is where stress testing comes in. Stress testing picks up the baton from VaR and other data-driven statistical techniques as the time horizons become longer and longer and the risk manager wants to explore the impact of events that are not present in her dataset – or, perhaps, that have never occurred before.⁷

As I explain in Chapter 4, stress testing, by its very nature, can rely much less on a frequentist concept of probability, and almost *has* to interpret probability in a subjective sense. In Bayesian terms, as the time horizon lengthens and the severity of the events increases, the ‘prior’ acquires a greater and greater weight, and the likelihood function a smaller and smaller one.⁸ In my opinion, this is a strength, not a weakness, of stress testing. It is also, however, the aspect of the project I propose that requires most careful handling. Frequentist probability may make little sense when it comes to stress testing, but this does not mean that probability *tout court* has no place in stress testing. If anything, it is stress testing without *any* notion of probability that, as Aragonés, Blanco and Dowd remind us, is of limited use. The challenge taken up in this book is to provide the missing link between

⁷Is Extreme Value Theory not supposed to provide information about the far tails of a distribution? It may, but only if the underlying phenomenon is time-stationary over sufficiently long time periods (decades if we try to estimate once-in-a-decade events). Also, EVT is silent about the causal mechanism that produces the tail events. There is no way of telling whether, conditional on today’s state of the world, the probability of the tail event predicted by the EVT applies. See Chapter 4 and Sections 4.2 to 4.4 in particular.

⁸For readers not familiar with Bayesian analysis, I shall explain what I mean by this in Chapter 4.

stress events and their approximate likelihood – as explained, an essential prerequisite for action⁹ – without inappropriately resorting to purely frequentist methods.

The enterprise I have briefly sketched therefore gives us some hope of bringing stress losses within the same conceptual framework as the more mundane losses analysed by VaR-like techniques. The approach I suggest in this book bridges the gap between the probabilities that a risk manager *can*, with some effort, provide (marginal, and simple conditioned probabilities) and the probabilities that she requires (the joint probabilities). It does so by exploiting to the fullest the risk manager's understanding of the causal links between the individual stress events. By employing the causal, rather than associative, language, it resonates with our intuition and works with, not against, our cognitive grain.

The approach I suggest is therefore intended to give us guidance as to whether and when we should really worry, and to suggest how to act accordingly. It gives, in short, tools to ensure that the stress losses are approximately but consistently 'priced'. Hopefully, all of this might give us a tool for managing financial risk more effectively than we have been able to do so far.

This is what this book is about, and this is why, I think, it is important.

1.2 Plan of the Book

This book is structured in four parts. The first, which contains virtually no equations, puts stress testing and probabilistic assessments of rare financial events in their context. The second part presents the quantitative ideas and techniques required for the task. Here lots of formulae will be found. The third part deals with the quantitative applications of the concepts introduced in Part II. The fourth and last part deals with practical implementation issues, and equations therefore disappear from sight again.

Let me explain in some detail what is covered in these four parts.

After the optimistic note with which I closed the previous section, in Chapter 2 I move swiftly to dampen the reader's enthusiasm, by arguing that stress testing is not the solution to all our risk management problems. In particular, I make the important distinction, too often forgotten, between risk and uncertainty and explain what this entails for stress analysis.

With these caveats out of the way, I argue that the expert knowledge of the risk manager is essential in constructing, using and associating probabilities to stress events. This expert knowledge (and the 'models of reality' that underpin it) constitutes the link between the past data and the possible future outcomes. In Chapter 3 I therefore try to explain the role played by competing interpretative models of reality in helping the risk manager to 'conceive of the unconceivable'. Chapter 3 is therefore intended to put into context the specific suggestions about stress testing that I provide in the rest of the book.

In Chapter 4 I describe the different types of probability (frequentist and subjective) that can be used for risk management, and discuss which 'type of probability' is better suited to different analytical tasks. The chapter closes with an important distinction between associative and causal descriptions. This distinction is at the basis of the efficient elicitation of conditional probabilities, and of the Bayesian-net approach described later in the book.

In Part II I lay the quantitative foundations required for the applications presented in the rest of the book. Some of the concepts are elementary, others are less well-known. In

⁹'Some of the things I vaguely apprehend are, like the end of the world, uninsurable risks, and it's useless to worry about them' – Keynes, quoted in Skidelsky (2009), page 71.

order to give a unified treatment I deal with both the elementary concepts (Chapter 5) and the somewhat-more-advanced ones (Chapter 6) using the same conceptual framework and formalism. Venn diagrams will play a major role throughout the book. Chapter 7 shows how very useful bounds on joint probabilities can be obtained by specifying marginals and (some) singly-conditioned probabilities. Chapter 8 introduces Bayesian nets, and Chapter 9 explains how to build the conditional probability tables required to use them. This concludes the tool-gathering part of the book. (A simple introduction to Linear Programming can be found in Appendix Chapter 15.)

Part III is then devoted to the application of the conceptual tools and techniques presented in Part II. This is achieved by introducing two different possible systematic approaches to stress testing, of different ambition and scope, which are described in Chapters 10 and 11.

Finally, in Part IV I address more practical questions: how we can try to overcome the difficulties and the cognitive biases that stand in the way of providing reasonable conditional probabilities (Chapter 12); how we can structure our chain of stress events (Chapter 13); and how we can embed the suggestions of the book into a viable approach in a real financial institution (as opposed to a classroom exercise). Doing so requires taking into account the reality of its governance structure, its reporting lines and the need for independence of a well-functioning risk-management function (Chapter 14).

I have prepared Parts II and III with exercises. I have done so not because I see this book necessarily as a text for a formal course, but because I firmly believe that, in order to really understand new quantitative techniques, there is no substitute for getting one's hands dirty and actually working out some problems in full.

1.3 Suggestions for Further Reading

Stress testing is the subject of a seemingly endless number of white, grey and variously-coloured consultation papers by the BIS and other international bodies. At the time of writing, the latest such paper I am aware of is BIS (2009), but, no doubt, by the time this book reaches the shelves many new versions will have appeared. Good sources of up-to-date references are the publication sites of the BIS, the IIF and the IMF.

Part I

Data, Models and Reality

Chapter 2

Risk and Uncertainty – or, Why Stress Testing is *Not* Enough

In the introductory chapter I made my pitch as to why stress testing is important and why I believe that the approach I propose can show us the way out of Berkowitz's (1999) purgatory. I don't want to convey the impression, however, that stress testing can be the answer to all our risk management questions. The problem, I think, does not lie with the specific approach I suggest in this book – flawed as this may be – but is of a fundamental nature. To present a balanced picture, I must therefore share two important reservations.

2.1 The Limits of Quantitative Risk Analysis

The first reservation is that the quantitative assessment of risk (and I include stress testing in this category) is an important part of risk management, but it is far from being its beginning and end. Many commentators and risk 'gurus' have stressed the inadequacies of the current quantitative techniques. The point is taken. But even if the best quantitative assessment of risk were available, a lot more would be required to translate this insight into effective risk management. The purpose of analysis is to inform action. Within a complex organization, effective action can only take place in what I call a favourable institutional environment.¹ So, in a favourable institutional environment the output of the quantitative analysis is first escalated, and then understood and challenged by senior management. This is now well accepted. But there is a lot more, and this 'a lot more' has very little to do with quantitative risk analysis. The organizational set-up, for instance, must be such that conflicts of interest are minimized (in the real world they can never be totally eliminated). Or, the agency problems that bedevil any large organization, and financial institutions *in primis*, must be understood and addressed in a satisfactory manner. And again: an effective way must be found to align the interests of the private decision makers of a systemically-relevant institution such as a large

¹ A strong governance structure is a common expression with a similar meaning. I prefer to speak of institutional environment, which encompasses such intangibles as a 'robust risk management culture', because even the strongest governance structure can fail if it does not dovetail constructively with the underlying culture of the organization.