

WILEY FINANCE

Financial Risk Management

SECOND EDITION

*A Practitioner's Guide to
Managing Market and Credit Risk*

+ WEBSITE

STEVEN ALLEN

Financial Risk Management

Founded in 1807, John Wiley & Sons is the oldest independent publishing company in the United States. With offices in North America, Europe, Australia, and Asia, Wiley is globally committed to developing and marketing print and electronic products and services for our customers' professional and personal knowledge and understanding.

The Wiley Finance series contains books written specifically for finance and investment professionals as well as sophisticated individual investors and their financial advisors. Book topics range from portfolio management to e-commerce, risk management, financial engineering, valuation, and financial instrument analysis, as well as much more.

For a list of available titles, visit our Web site at www.WileyFinance.com.

Financial Risk Management

*A Practitioner's Guide to Managing
Market and Credit Risk*

Second Edition

STEVEN ALLEN



John Wiley & Sons, Inc.

Cover image: John Wiley & Sons, Inc.

Cover design: © Tom Fewster / iStockphoto, © samxmeg / iStockphoto

Copyright © 2013 by Steven Allen. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Published simultaneously in Canada.

The First Edition of this book was published in 2003 by John Wiley & Sons, Inc.

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, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the Web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Allen, Steven, 1945–

Financial risk management [electronic resource]: a practitioner's guide to managing market and credit risk / Steven Allen. — 2nd ed.

1 online resource.

Includes bibliographical references and index.

Description based on print version record and CIP data provided by publisher; resource not viewed.

ISBN 978-1-118-17545-3 (cloth); 978-1-118-22652-0 (ebk.); ISBN 978-1-118-23164-7 (ebk.); ISBN 978-1-118-26473-7 (ebk.)

1. Financial risk management. 2. Finance. I. Title.

HD61

658.15'5—dc23

2012029614

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

*To Caroline
For all the ways she has helped bring
this project to fruition
And for much, much more*

Financial Risk Management

Contents

Foreword	xvii
Preface	xix
Acknowledgments	xxiii
About the Author	xxvii
CHAPTER 1	
Introduction	1
1.1 Lessons from a Crisis	1
1.2 Financial Risk and Actuarial Risk	2
1.3 Simulation and Subjective Judgment	4
CHAPTER 2	
Institutional Background	7
2.1 Moral Hazard—Insiders and Outsiders	7
2.2 Ponzi Schemes	17
2.3 Adverse Selection	19
2.4 The Winner’s Curse	21
2.5 Market Making versus Position Taking	24
CHAPTER 3	
Operational Risk	29
3.1 Operations Risk	31
3.1.1 The Risk of Fraud	31
3.1.2 The Risk of Nondeliberate Incorrect Information	35
3.1.3 Disaster Risk	36
3.1.4 Personnel Risk	36
3.2 Legal Risk	37
3.2.1 The Risk of Unenforceable Contracts	37
3.2.2 The Risk of Illegal Actions	40
3.3 Reputational Risk	41
3.4 Accounting Risk	42

3.5	Funding Liquidity Risk	42
3.6	Enterprise Risk	44
3.7	Identification of Risks	44
3.8	Operational Risk Capital	45
CHAPTER 4		
Financial Disasters		49
4.1	Disasters Due to Misleading Reporting	49
4.1.1	Chase Manhattan Bank/Drysdale Securities	52
4.1.2	Kidder Peabody	53
4.1.3	Barings Bank	55
4.1.4	Allied Irish Bank (AIB)	57
4.1.5	Union Bank of Switzerland (UBS)	59
4.1.6	Société Générale	61
4.1.7	Other Cases	66
4.2	Disasters Due to Large Market Moves	68
4.2.1	Long-Term Capital Management (LTCM)	68
4.2.2	Metallgesellschaft (MG)	75
4.3	Disasters Due to the Conduct of Customer Business	77
4.3.1	Bankers Trust (BT)	77
4.3.2	JPMorgan, Citigroup, and Enron	79
4.3.3	Other Cases	80
CHAPTER 5		
The Systemic Disaster of 2007–2008		83
5.1	Overview	83
5.2	The Crisis in CDOs of Subprime Mortgages	85
5.2.1	Subprime Mortgage Originators	86
5.2.2	CDO Creators	88
5.2.3	Rating Agencies	89
5.2.4	Investors	92
5.2.5	Investment Banks	93
5.2.6	Insurers	106
5.3	The Spread of the Crisis	108
5.3.1	Credit Contagion	108
5.3.2	Market Contagion	109
5.4	Lessons from the Crisis for Risk Managers	111
5.4.1	Subprime Mortgage Originators	111
5.4.2	CDO Creators	111
5.4.3	Rating Agencies	111
5.4.4	Investors	111
5.4.5	Investment Banks	112

5.4.6	Insurers	114
5.4.7	Credit Contagion	115
5.4.8	Market Contagion	115
5.5	Lessons from the Crisis for Regulators	115
5.5.1	Mortgage Originators	116
5.5.2	CDO Creators	116
5.5.3	Rating Agencies	117
5.5.4	Investors	118
5.5.5	Investment Banks	118
5.5.6	Insurers	126
5.5.7	Credit Contagion	126
5.5.8	Market Contagion	129
5.6	Broader Lessons from the Crisis	132
CHAPTER 6		
Managing Financial Risk		133
6.1	Risk Measurement	133
6.1.1	General Principles	133
6.1.2	Risk Management of Instruments That Lack Liquidity	144
6.1.3	Market Valuation	147
6.1.4	Valuation Reserves	152
6.1.5	Analysis of Revenue	156
6.1.6	Exposure to Changes in Market Prices	157
6.1.7	Risk Measurement for Position Taking	159
6.2	Risk Control	161
CHAPTER 7		
VaR and Stress Testing		169
7.1	VaR Methodology	170
7.1.1	Simulation of the P&L Distribution	173
7.1.2	Measures of the P&L Distribution	187
7.2	Stress Testing	192
7.2.1	Overview	192
7.2.2	Economic Scenario Stress Tests	193
7.2.3	Stress Tests Relying on Historical Data	197
7.3	Uses of Overall Measures of Firm Position Risk	201
CHAPTER 8		
Model Risk		209
8.1	How Important Is Model Risk?	210
8.2	Model Risk Evaluation and Control	212

8.2.1	Scope of Model Review and Control	213
8.2.2	Roles and Responsibilities for Model Review and Control	214
8.2.3	Model Verification	219
8.2.4	Model Verification of Deal Representation	222
8.2.5	Model Verification of Approximations	223
8.2.6	Model Validation	226
8.2.7	Continuous Review	232
8.2.8	Periodic Review	234
8.3	Liquid Instruments	237
8.4	Illiquid Instruments	241
8.4.1	Choice of Model Validation Approach	241
8.4.2	Choice of Liquid Proxy	243
8.4.3	Design of Monte Carlo Simulation	245
8.4.4	Implications for Marking to Market	247
8.4.5	Implications for Risk Reporting	249
8.5	Trading Models	250
CHAPTER 9		
	Managing Spot Risk	253
9.1	Overview	253
9.2	Foreign Exchange Spot Risk	257
9.3	Equity Spot Risk	258
9.4	Physical Commodities Spot Risk	259
CHAPTER 10		
	Managing Forward Risk	263
10.1	Instruments	270
10.1.1	Direct Borrowing and Lending	270
10.1.2	Repurchase Agreements	271
10.1.3	Forwards	272
10.1.4	Futures Contracts	272
10.1.5	Forward Rate Agreements	274
10.1.6	Interest Rate Swaps	275
10.1.7	Total Return Swaps	276
10.1.8	Asset-Backed Securities	278
10.2	Mathematical Models of Forward Risks	282
10.2.1	Pricing Illiquid Flows by Interpolation	284
10.2.2	Pricing Long-Dated Illiquid Flows by Stack and Roll	291
10.2.3	Flows Representing Promised Deliveries	293
10.2.4	Indexed Flows	295

10.3	Factors Impacting Borrowing Costs	299
10.3.1	The Nature of Borrowing Demand	299
10.3.2	The Possibility of Cash-and-Carry Arbitrage	300
10.3.3	The Variability of Storage Costs	301
10.3.4	The Seasonality of Borrowing Costs	302
10.3.5	Borrowing Costs and Forward Prices	303
10.4	Risk Management Reporting and Limits for Forward Risk	304
CHAPTER 11		
	Managing Vanilla Options Risk	311
11.1	Overview of Options Risk Management	313
11.2	The Path Dependence of Dynamic Hedging	318
11.3	A Simulation of Dynamic Hedging	321
11.4	Risk Reporting and Limits	329
11.5	Delta Hedging	344
11.6	Building a Volatility Surface	346
11.6.1	Interpolating between Time Periods	346
11.6.2	Interpolating between Strikes—Smile and Skew	347
11.6.3	Extrapolating Based on Time Period	352
11.7	Summary	355
CHAPTER 12		
	Managing Exotic Options Risk	359
12.1	Single-Payout Options	364
12.1.1	Log Contracts and Variance Swaps	367
12.1.2	Single-Asset Quanto Options	369
12.1.3	Convexity	370
12.1.4	Binary Options	371
12.1.5	Contingent Premium Options	377
12.1.6	Accrual Swaps	378
12.2	Time-Dependent Options	378
12.2.1	Forward-Starting and Cliquet Options	378
12.2.2	Compound Options	379
12.3	Path-Dependent Options	381
12.3.1	Standard Analytic Models for Barriers	383
12.3.2	Dynamic Hedging Models for Barriers	385
12.3.3	Static Hedging Models for Barriers	387
12.3.4	Barrier Options with Rebates, Lookback, and Ladder Options	402
12.3.5	Broader Classes of Path-Dependent Exotics	403
12.4	Correlation-Dependent Options	404

12.4.1	Linear Combinations of Asset Prices	405
12.4.2	Risk Management of Options on Linear Combinations	409
12.4.3	Index Options	413
12.4.4	Options to Exchange One Asset for Another	415
12.4.5	Nonlinear Combinations of Asset Prices	417
12.4.6	Correlation between Price and Exercise	422
12.5	Correlation-Dependent Interest Rate Options	425
12.5.1	Models in Which the Relationship between Forwards Is Treated as Constant	426
12.5.2	Term Structure Models	430
12.5.3	Relationship between Swaption and Cap Prices	437

CHAPTER 13**Credit Risk 445**

13.1	Short-Term Exposure to Changes in Market Prices	446
13.1.1	Credit Instruments	447
13.1.2	Models of Short-Term Credit Exposure	451
13.1.3	Risk Reporting for Market Credit Exposures	456
13.2	Modeling Single-Name Credit Risk	457
13.2.1	Estimating Probability of Default	458
13.2.2	Estimating Loss Given Default	465
13.2.3	Estimating the Amount Owed at Default	468
13.2.4	The Option-Theoretic Approach	471
13.3	Portfolio Credit Risk	479
13.3.1	Estimating Default Correlations	479
13.3.2	Monte Carlo Simulation of Portfolio Credit Risk	482
13.3.3	Computational Alternatives to Full Simulation	486
13.3.4	Risk Management and Reporting for Portfolio Credit Exposures	490
13.4	Risk Management of Multiname Credit Derivatives	493
13.4.1	Multiname Credit Derivatives	493
13.4.2	Modeling of Multiname Credit Derivatives	495
13.4.3	Risk Management and Reporting for Multiname Credit Derivatives	498
13.4.4	CDO Tranches and Systematic Risk	500

CHAPTER 14**Counterparty Credit Risk 505**

14.1	Overview	505
14.2	Exchange-Traded Derivatives	506

14.3	Over-the-Counter Derivatives	512
14.3.1	Overview	512
14.3.2	The Loan-Equivalent Approach	513
14.3.3	The Collateralization Approach	515
14.3.4	The Collateralization Approach— Wrong-Way Risk	521
14.3.5	The Active Management Approach	526
	References	533
	About the Companion Website	547
	Index	553

Foreword

Risk was a lot easier to think about when I was a doctoral student in finance 25 years ago. Back then, risk was measured by the variance of your wealth. Lowering risk meant lowering this variance, which usually had the unfortunate consequence of lowering the average return on your wealth as well.

In those halcyon days, we had only two types of risk, systemic and un-systematic. The latter one could be lowered for free via diversification, while the former one could only be lowered by taking a hit to average return. In that idyllic world, financial risk management meant choosing the variance that maximized expected utility. One merely had to solve an optimization problem. What could be easier?

I started to appreciate that financial risk management might not be so easy when I moved from the West Coast to the East Coast. The New York-based banks started creating whole departments to manage financial risk. Why do you need dozens of people to solve a simple optimization problem? As I talked with the denizens of those departments, I noticed they kept introducing types of risk that were not in my financial lexicon. First there was credit risk, a term that was to be differentiated from market risk, because you can lose money lending whether a market exists or not. Fine, I got that, but then came liquidity risk on top of market and credit risk. Just as I was struggling to integrate these three types of risk, people started worrying about operational risk, basis risk, mortality risk, weather risk, estimation risk, counterparty credit risk, and even the risk that your models for all these risks were wrong. If model risk existed, then you had to concede that even your model for model risk was risky.

Since the proposed solution for all these new risks were new models and since the proposed solution for the model risk of the new models was yet more models, it was no wonder all of those banks had all of those people running around managing all of those risks.

Well, apparently, not quite enough people. As I write these words, the media are having a field day denouncing JPMorgan's roughly \$6 billion loss related to the London whale's ill-fated foray into credit default swaps (CDSs).

As the flag bearer for the TV generation, I can't help but think of reviving a 1970s TV show to star Bruno Iksil as the Six Billion Dollar Man. As eye-popping as these numbers are, they are merely the fourth largest trading loss since the first edition of this book was released. If we ignore Bernie Madoff's

\$50 billion Ponzi scheme, the distinction for the worst trade ever belongs to Howie Hubler, who lost \$9 billion trading CDSs in 2008 for another bank whose name I'd rather not write. However, if you really need to know, then here's a hint. The present occupant of Mr. Hubler's old office presently thinks that risk management is a complicated subject, very complicated indeed, and has to admit that a simple optimization is not the answer. So what is the answer? Well, when the answer to a complicated question is nowhere to be found in the depths of one's soul, then one can always fall back on asking the experts instead. The Danish scientist Niels Bohr, once deemed an expert, said an expert is, "A person that has made every possible mistake within his or her field."

As an expert in the field of derivative securities valuation, I believe I know a fellow expert when I see one. Steve Allen has been teaching courses in risk management at New York University's Courant Institute since 1998. Steve retired from JPMorgan Chase as a managing director in 2004, capping a 35-year career in the finance industry. Given the wide praise for the first edition of this book, the author could have rested on his laurels, comforted by the knowledge that the wisdom of the ages is eternal. Instead, he has taken it upon himself to write a second edition of this timeless book.

Most authors in Steve's enviable situation would have contented themselves with exploiting the crisis to elaborate on some extended version of "I told you so." Instead, Steve has added much in the way of theoretical advances that have arisen out of the necessity of ensuring that history does not repeat itself. These advances in turn raise the increasing degree of specialization we see inside the risk management departments of modern financial institutions and increasingly in the public sector as well. Along with continued progress in the historically vital problem of marking to market of illiquid positions, there is an increasing degree of rigor in the determination of reserves that arise due to model risk, in the limits used to control risk taking, and in the methods used to review models. The necessity of testing every assumption has been made plain by the stress that the crisis has imposed on our fragile financial system. As the aftershocks reverberate around us, we will not know for many years whether the present safeguards will serve their intended purpose. However, the timing for an update to Steve's book could not be better. I truly hope that the current generation of risk managers, whether they be grizzled or green, will take the lessons on the ensuing pages to heart. Our shared financial future depends on it.

Peter Carr, PhD
Managing Director at Morgan Stanley,
Global Head of Market Modeling, and
Executive Director of New York University Courant's
Masters in Mathematical Finance

Preface

This book offers a detailed introduction to the field of risk management as performed at large investment and commercial banks, with an emphasis on the practices of specialist market risk and credit risk departments as well as trading desks. A large portion of these practices is also applicable to smaller institutions that engage in trading or asset management.

The aftermath of the financial crisis of 2007–2008 leaves a good deal of uncertainty as to exactly what the structure of the financial industry will look like going forward. Some of the business currently performed in investment and commercial banks, such as proprietary trading, may move to other institutions, at least in some countries, based on new legislation and new regulations. But in whatever institutional setting this business is conducted, the risk management issues will be similar to those encountered in the past. This book focuses on general lessons as to how the risk of financial institutions can be managed rather than on the specifics of particular regulations.

My aim in this book is to be comprehensive in looking at the activities of risk management specialists as well as trading desks, at the realm of mathematical finance as well as that of the statistical techniques, and, most important, at how these different approaches interact in an integrated risk management process.

This second edition reflects lessons that have been learned from the recent financial crisis of 2007–2008 (for more detail, see Chapters 1 and 5), as well as many new books, articles, and ideas that have appeared since the publication of the first edition in 2003. Chapter 6 on managing market risk, Chapter 7 on value at risk (VaR) and stress testing, Chapter 8 on model risk, and Chapter 13 on credit risk are almost completely rewritten and expanded from the first edition, and a new Chapter 14 on counterparty credit risk is an extensive expansion of a section of the credit risk chapter in the first edition.

The website for this book (www.wiley.com/go/frm2e) will be used to provide both supplementary materials to the text and continuous updates. Supplementary materials will include spreadsheets and computer code that illustrate computations discussed in the text. In addition, there will be classroom aids available only to professors on the Wiley Higher Education website. Updates will include an updated electronic version of the References

section, to allow easy cut-and-paste linking to referenced material on the web. Updates will also include discussion of new developments. For example, at the time this book went to press, there is not yet enough public information about the causes of the large trading losses at JPMorgan's London investment office to allow a discussion of risk management lessons; as more information becomes available, I will place an analysis of risk management lessons from these losses on the website.

This book is divided into three parts: general background to financial risk management, the principles of financial risk management, and the details of financial risk management.

- The general background part (Chapters 1 through 5) gives an institutional framework for understanding how risk arises in financial firms and how it is managed. Without understanding the different roles and motivations of traders, marketers, senior firm managers, corporate risk managers, bondholders, stockholders, and regulators, it is impossible to obtain a full grasp of the reasoning behind much of the machinery of risk management or even why it is necessary to manage risk. In this part, you will encounter key concepts risk managers have borrowed from the theory of insurance (such as moral hazard and adverse selection), decision analysis (such as the winner's curse), finance theory (such as the arbitrage principle), and in one instance even the criminal courts (the Ponzi scheme). Chapter 4 provides discussion of some of the most prominent financial disasters of the past 30 years, and Chapter 5 focuses on the crisis of 2007–2008. These serve as case studies of failures in risk management and will be referenced throughout the book. This part also contains a chapter on operational risk, which is necessary background for many issues that arise in preventing financial disasters and which will be referred to throughout the rest of the book.
- The part on principles of financial risk management (Chapters 6 through 8) first lays out an integrated framework in Chapter 6, and then looks at VaR and stress testing in Chapter 7 and the control of model risk in Chapter 8.
- The part on details of financial risk management (Chapters 9 through 14) applies the principles of the second part to each specific type of financial risk: spot risk in Chapter 9, forward risk in Chapter 10, vanilla options risk in Chapter 11, exotic options risk in Chapter 12, credit risk in Chapter 13, and counterparty credit risk in Chapter 14. As each risk type is discussed, specific references are made to the principles elucidated in Chapters 6 through 8, and a detailed analysis of the models used to price these risks and how these models can be used to measure and control risk is presented.

Since the 1990s, an increased focus on the new technology being developed to measure and control financial risk has resulted in the growth of corporate staff areas manned by risk management professionals. However, this does not imply that financial firms did not manage risks prior to 1990 or that currently all risk management is performed in staff areas. Senior line managers such as trading desk and portfolio managers have always performed a substantial risk management function and continue to do so. In fact, confusion can be caused by the tradition of using the term *risk manager* as a synonym for a senior trader or portfolio manager and as a designation for members of corporate staff areas dealing with risk. Although this book covers risk management techniques that are useful to both line trading managers and corporate staff acting on behalf of the firm's senior management, the needs of these individuals do not completely overlap. I will try to always make a clear distinction between information that is useful to a trading desk and information that is needed by corporate risk managers, and explain how they might intersect.

Books and articles on financial risk management have tended to focus on statistical techniques embodied in measures such as value at risk (VaR). As a result, risk management has been accused of representing a very narrow specialty with limited value, a view that has been colorfully expressed by Nassim Taleb (1997), "There has been growth in the number of 'risk management advisors,' an industry sometimes populated by people with an amateurish knowledge of risk. Using some form of shallow technical skills, these advisors emit pronouncements on such matters as 'risk management' without a true understanding of the distribution. Such inexperience and weakness become more apparent with the value-at-risk fad or the outpouring of books on risk management by authors who never traded a contract" (p. 4).

This book gives a more balanced account of risk management. Less than 20 percent of the material looks at statistical techniques such as VaR. The bulk of the book examines issues such as the proper mark-to-market valuation of trading positions, the determination of necessary reserves against valuation uncertainty, the structuring of limits to control risk taking, and the review of mathematical models and determination of how they can contribute to risk control. This allocation of material mirrors the allocation of effort in the corporate risk management staff areas with which I am familiar. This is reflected in the staffing of these departments. More personnel is drawn from those with experience and expertise in trading and building models to support trading decisions than is drawn from a statistical or academic finance background.

Although many readers may already have a background in the instruments—bonds, stocks, futures, and options—used in the financial markets, I have supplied definitions every time I introduce a term. Terms are italicized in the text at the point they are defined. Any reader feeling the need for a

more thorough introduction to market terminology should find the first nine chapters of Hull (2012) adequate preparation for understanding the material in this book.

My presentation of the material is based both on theory and on how concepts are utilized in industry practice. I have tried to provide many concrete instances of either personal experience or reports I have heard from industry colleagues to illustrate these practices. Where incidents have received sufficient previous public scrutiny or occurred long enough ago that issues of confidentiality are not a concern, I have provided concrete details. In other cases, I have had to preserve the anonymity of my sources by remaining vague about particulars. My preservation of anonymity extends to a liberal degree of randomness in references to gender.

A thorough discussion of how mathematical models are used to measure and control risks must make heavy reference to the mathematics used in creating these models. Since excellent expositions of the mathematics exist, I do not propose to enter into extensive derivations of results that can readily be found elsewhere. Instead, I will concentrate on how these results are used in risk management and how the approximations to reality inevitable in any mathematical abstraction are dealt with in practice. I will provide references to the derivation of results. Wherever possible, I have used Hull (2012) as a reference, since it is the one work that can be found on the shelf of nearly every practitioner in the field of quantitative finance.

Although the material for this book was originally developed for a course taught within a mathematics department, I believe that virtually all of its material will be understandable to students in finance programs and business schools, and to practitioners with a comparable educational background. A key reason for this is that whereas derivatives mathematics often emphasizes the use of more mathematically sophisticated continuous time models, discrete time models are usually more relevant to risk management, since risk management is often concerned with the limits that real market conditions place on mathematical theory.

This book is designed to be used either as a text for a course in risk management or as a resource for self-study or reference for people working in the financial industry. To make the material accessible to as broad an audience as possible, I have tried everywhere to supplement mathematical theory with concrete examples and have supplied spreadsheets on the accompanying website (www.wiley.com/go/frm2e) to illustrate these calculations. Spreadsheets on the website are referenced throughout the text and a summary of all spreadsheets supplied is provided in the “About the Companion Website” section at the back of the book. At the same time, I have tried to make sure that all the mathematical theory that gets used in risk management practice is addressed. For readers who want to pursue the theoretical developments at greater length, a full set of references has been provided.

Acknowledgments

The views expressed in this book are my own, but have been shaped by my experiences in the financial industry. Many of my conclusions about what constitutes best practice in risk management have been based on my observation of and participation in the development of the risk management structure at JPMorgan Chase and its Chemical Bank and Chase Manhattan Bank predecessors.

The greatest influence on my overall view of how financial risk management should be conducted and on many of the specific approaches I advocate has been Lesley Daniels Webster. My close collaboration with Lesley took place over a period of 20 years, during the last 10 of which I reported to her in her position as director of market risk management. I wish to express my appreciation of Lesley's leadership, along with that of Marc Shapiro, Suzanne Hammett, Blythe Masters, and Andy Threadgold, for having established the standards of integrity, openness, thoroughness, and intellectual rigor that have been the hallmarks of this risk management structure.

Throughout most of the period in which I have been involved in these pursuits, Don Layton was the head of trading activities with which we interacted. His recognition of the importance of the risk management function and strong support for a close partnership between risk management and trading and the freedom of communication and information sharing were vital to the development of these best practices.

Through the years, my ideas have benefited from my colleagues at Chemical, Chase, JPMorgan Chase, and in consulting assignments since my retirement from JPMorgan Chase. At JPMorgan Chase and its predecessors, I would particularly like to note the strong contributions that dialogues with Andrew Abrahams, Michel Araten, Bob Benjamin, Paul Bowmar, George Brash, Julia Chislenko, Enrico Della Vecchia, Mike Diniias, Fawaz Habel, Bob Henderson, Jeff Katz, Bobby Magee, Blythe Masters, Mike Rabin, Barry Schachter, Vivian Shelton, Paul Shotton, Andy Threadgold, Mick Waring, and Richard Wise have played in the development of the concepts utilized here. In my consulting assignments, I have gained much from my exchanges of ideas with Rick Grove, Chia-Ling Hsu, Neil Pearson, Bob Selvaggio, Charles Smithson, and other colleagues at Rutter Associates, and Chris Marty and Alexey Panchevka at Bloomberg. In interactions with risk

managers at other firms, I have benefited from my conversations with Ken Abbott, John Breit, Noel Donohoe, and Evan Picoult. Many of the traders I have interacted with through the years have also had a major influence on my views of how risk management should impact decision making on the trading desk and the proper conduct of relationships between traders and risk management specialists. I particularly want to thank Andy Hollings, Simon Lack, Jeff Larsen, Dinsa Mehta, Fraser Partridge, and Don Wilson for providing me with prototypes for how the risk management of trading should be properly conducted and their generosity in sharing their knowledge and insight. I also wish to thank those traders, who shall remain anonymous here, who have provided me equally valuable lessons in risk management practices to avoid.

This book grew out of the risk management course I created as part of the Mathematics in Finance MS program at New York University's Courant Institute of Mathematical Sciences in 1998. For giving me the opportunity to teach and for providing an outstanding institutional setting in which to do it, I want to thank the administration and faculty of Courant, particularly Peter Carr, Neil Chriss, Jonathan Goodman, Bob Kohn, and Petter Kolm, with whom I have participated in the management of the program, and Caroline Thompson, Gabrielle Tobin, and Melissa Vacca, the program administrators. I have gained many insights that have found their way into this book by attending other courses in the program taught by Marco Avelaneda, Jim Gatheral, Bob Kohn, and Nassim Taleb.

Ken Abbott began participating in the risk management course as a guest lecturer, later became my co-teacher of the course, and now has full responsibility for the course with my participation as a guest lecturer. Many of the insights in this book have been learned from Ken or generated as part of the debates and discussions we have held both in and out of the classroom. The students in my risk management course have helped clarify many of the concepts in this book through their probing questions. I particularly want to thank Karim Beguir, who began as my student and has since graduated to become a Fellow of the program and a frequent and valued contributor to the risk management course. Several of his insights are reflected in the second edition of the book. I also wish to thank Otello Padovani and Andrea Raphael, students who became collaborators on research that appears on the website for the book (www.wiley.com/go/frm2e). Mike Fisher has provided greatly appreciated support as my graduate assistant in helping to clarify class assignments that have evolved into exercises in this book.

The detailed comments and suggestions I have received from Neil Chriss on large portions of this manuscript far exceed the norms of either friendship or collegiality. In numerous instances, his efforts have sharpened both the ideas being presented and the clarity of their expression. I also wish

to thank Mich Araten, Peter Carr, Bobby Magee, Barry Schachter, Nassim Taleb, and Bruce Tuckman for reading the text and offering helpful comments. For the second edition, I would like to thank Ken Abbott and Rick Grove for reading new chapters and offering helpful suggestions.

I also wish to extend my thanks to Chuck Epstein for his help in finding a publisher for this book. Bill Falloon, Meg Freeborn, and Michael Kay, my editors at John Wiley & Sons, have offered very useful suggestions at every stage of the editing. At MacAllister Publishing Services, Andy Stone was very helpful as production manager and Jeanne Henning was a thorough and incisive copy editor for the first edition of this book.

The individual to whom both I and this book owe the greatest debt is my wife, Caroline Thompson. The number of ways in which her beneficial influence has been felt surpass my ability to enumerate, but I at least need to attempt a brief sample. It was Caroline who introduced me to Neil Chriss and first planted the idea of my teaching at Courant. She has been a colleague of Neil's, Jonathan Goodman's, and mine in the continued development of the Courant Mathematics in Finance MS program. From the start, she was the strongest voice in favor of basing a book on my risk management course. At frequent bottlenecks, on both the first and second editions, when I have been daunted by an obstacle to my progress that seemed insurmountable, it was Caroline who suggested the approach, organized the material, or suggested the joint effort that overcame the difficulty. She has managed all aspects of the production format, and style of the book, including efforts from such distant ports as Laos, Vietnam, India, and Holland.

About the Author

Steve Allen is a risk management consultant, specializing in risk measurement and valuation with a particular emphasis on illiquid and hard-to-value assets. Until his retirement in 2004, he was Managing Director in charge of risk methodology at JPMorgan Chase, where he was responsible for model validation, risk capital allocation, and the development of new measures of valuation, reserves, and risk for both market and credit risk. Previously, he was in charge of market risk for derivative products at Chase. He has been a key architect of Chase's value-at-risk and stress testing systems. Prior to his work in risk management, Allen was the head of analysis and model building for all Chase trading activities for over ten years. Since 1998, Allen has been associated with the Mathematics in Finance Masters' program at New York University's Courant Institute of Mathematical Sciences. In this program, he has served as Clinical Associate Professor and Deputy Director and has created and taught courses in risk management, derivatives mathematics, and interest rate and credit models. He was a member of the Board of Directors of the International Association of Financial Engineers and continues to serve as co-chair of their Education Committee.

Introduction

1.1 LESSONS FROM A CRISIS

I began the first edition of this book with a reference to an episode of the television series *Seinfeld* in which the character George Costanza gets an assignment from his boss to read a book titled *Risk Management* and then give a report on this topic to other business executives. Costanza finds the book and topic so boring that his only solution is to convince someone else to read it for him and prepare notes. Clearly, my concern at the time was to write about financial risk management in a way that would keep readers from finding the subject dull. I could hardly have imagined then that eight years later Demi Moore would be playing the part of the head of an investment bank's risk management department in a widely released movie, *Margin Call*. Even less could I have imagined the terrible events that placed financial risk management in such a harsh spotlight.

My concern now is that the global financial crisis of 2007–2008 may have led to the conclusion that risk management is an exciting subject whose practitioners and practices cannot be trusted. I have thoroughly reviewed the material I presented in the first edition, and it still seems to me that if the principles I presented, principles that represented industry best practices, had been followed consistently, a disaster of the magnitude we experienced would not have been possible. In particular, the points I made in the first edition about using stress tests in addition to value at risk (VaR) in determining capital adequacy (see the last paragraphs of Section 7.3 in this edition) and the need for substantial reserves and deferred compensation for illiquid positions (see Sections 6.1.4 and 8.4 in this edition) still seem sound. It is tempting to just restate the same principles and urge more diligence in their application, but that appears too close to the sardonic definition of insanity: “doing the same thing and expecting different results.” So I have looked for places where these principles need strengthening (you'll find a summary in Section 5.4). But I have also reworked the organization of

the book to emphasize two core doctrines that I believe are the keys to the understanding and proper practice of financial risk management.

The first core principle is that financial risk management is not just risk management as practiced in financial institutions; it is risk management that makes active use of trading in liquid markets to control risk. Risk management is a discipline that is important to a wide variety of companies, government agencies, and institutions—one need only think of accident prevention at nuclear power plants and public health measures to avoid influenza pandemics to see how critical it can be. While the risk management practiced at investment banks shares some techniques with risk management practiced at a nuclear facility, there remains one vital difference: much of the risk management at investment banks can utilize liquid markets as a key element in risk control; liquid markets are of virtually no use to the nuclear safety engineer.

My expertise is in the techniques of financial risk management, and that is the primary subject of this book. Some risks that financial firms take on cannot be managed using trading in liquid markets. It is vitally important to identify such risks and to be aware of the different risk management approaches that need to be taken for them. Throughout the book I will be highlighting this distinction and also focusing on the differences that degree of available liquidity makes. As shorthand, I will refer to risk that cannot be managed by trading in liquid markets as *actuarial risk*, since it is the type of risk that actuaries at insurance companies have been dealing with for centuries. Even in cases that must be analyzed using the actuarial risk approach, financial risk management techniques can still be useful in isolating the actuarial risk and in identifying market data that can be used as input to actuarial risk calculations. I will address this in greater detail in Section 1.2.

The second core principle is that the quantification of risk management requires simulation guided by both historical data and subjective judgment. This is a common feature of both financial risk and actuarial risk. The time period simulated may vary greatly, from value at risk (VaR) simulations of daily market moves for very liquid positions to simulations spanning decades for actuarial risk. But I will be emphasizing shared characteristics for all of these simulations: the desirability of taking advantage of as much historical data as is relevant, the need to account for nonnormality of statistical distributions, and the necessity of including subjective judgment. More details on these requirements are in Section 1.3.

1.2 FINANCIAL RISK AND ACTUARIAL RISK

The management of financial risk and the management of actuarial risk do share many methodologies, a point that will be emphasized in the next

section. Both rely on probability and statistics to arrive at estimates of the distribution of possible losses. The critical distinction between them is the matter of time. Actuarial risks may not be fully resolved for years, sometimes even decades. By the time the true extent of losses is known, the accumulation of risk may have gone on for years. Financial risks can be eliminated in a relatively short time period by the use of liquid markets. Continuous monitoring of the price at which risk can be liquidated should substantially lower the possibility of excessive accumulation of risk.

Two caveats need to be offered to this relatively benign picture of financial risk. The first is that taking advantage of the shorter time frame of financial risk requires constant vigilance; if you aren't doing a good job of monitoring how large your risks are relative to liquidation costs, you may still acquire more exposure than desired. This will be described in detail in Chapter 6. The second is the need to be certain that what is truly actuarial risk has not been misclassified as financial risk. If this occurs, it is especially dangerous—not only will you have the potential accumulation of risk over years before the extent of losses is known, but in not recognizing the actuarial nature, you would not exercise the caution that the actuarial nature of the risk demands. This will be examined more closely in Sections 6.1.1 and 6.1.2, with techniques for management of actuarial risk in financial firms outlined in Section 8.4. I believe that this dangerous muddling of financial and actuarial risk was a key contributor to the 2007–2008 crisis, as I argue in Section 5.2.5.

Of course, it is only an approximation to view instruments as being liquid or illiquid. The volume of instruments available for trading differs widely by size and readiness of availability. This constitutes the depth of liquidity of a given market. Often a firm will be faced with a choice between the risks of replicating positions more exactly with less liquid instruments or less exactly with more liquid instruments.

One theme of this book will be the trade-off between liquidity risk and basis risk. *Liquidity risk* is the risk that the price at which you buy (or sell) something may be significantly less advantageous than the price you could have achieved under more ideal conditions. *Basis risk* is the risk that occurs when you buy one product and sell another closely related one, and the two prices behave differently. Let's look at an example. Suppose you are holding a large portfolio of stocks that do not trade that frequently and your outlook for stock prices leads to a desire to quickly terminate the position. If you try selling the whole basket quickly, you face significant liquidity risk since your selling may depress the prices at which the stocks trade. An alternative would be to take an offsetting position in a heavily traded stock futures contract, such as the futures contract tied to the Standard & Poor's™ S&P 500 stock index. This lowers the liquidity risk, but it increases the

basis risk since changes in the price of your particular stock basket will probably differ from the price changes in the stock index. Often the only way in which liquidity risk can be reduced is to increase basis risk, and the only way in which basis risk can be reduced is to increase liquidity risk.

The classification of risk as financial risk or actuarial risk is clearly a function of the particular type of risk and not of the institution. Insurance against hurricane damage could be written as a traditional insurance contract by Metropolitan Life or could be the payoff of an innovative new swap contract designed by Morgan Stanley; in either case, it will be the same risk. What is required in either case is analysis of how trading in liquid markets can be used to manage the risk. Certainly commercial banks have historically managed substantial amounts of actuarial risk in their loan portfolios. And insurance companies have managed to create some ability to liquidate insurance risk through the reinsurance market. Even industrial firms have started exploring the possible transformation of some actuarial risk into financial risk through the theory of *real options*. An introduction to real options can be found in Hull (2012, Section 34) and Dixit and Pindyck (1994).

A useful categorization to make in risk management techniques that I will sometimes make use of, following Gumerlock (1999), is to distinguish between risk management through risk aggregation and risk management through risk decomposition. *Risk aggregation* attempts to reduce risk by creating portfolios of less than completely correlated risk, thereby achieving risk reduction through diversification. *Risk decomposition* attempts to reduce a risk that cannot directly be priced in the market by analyzing it into subcomponents, all or some of which can be priced in the market. Actuarial risk can generally be managed only through risk aggregation, whereas financial risk utilizes both techniques. Chapter 7 concentrates on risk aggregation, while Chapter 8 primarily focuses on risk decomposition; Chapter 6 addresses the integration of the two.

1.3 SIMULATION AND SUBJECTIVE JUDGMENT

Nobody can guarantee that all possible future contingencies have been provided for—this is simply beyond human capabilities in a world filled with uncertainty. But it is unacceptable to use that platitude as an excuse for complacency and lack of meaningful effort. It has become an embarrassment to the financial industry to see the number of events that are declared “once in a millennium” occurrences, based on an analysis of historical data, when they seem in fact to take place every few years. At one point I suggested, only half-jokingly, that anyone involved in risk management who used the words *perfect* and *storm* in the same sentence should be permanently

banned from the financial industry. More seriously, everyone involved in risk management needs to be aware that historical data has a limited utility, and that subjective judgment based on experience and careful reasoning must supplement data analysis. The failure of risk managers to apply critical subjective judgment as a check on historical data in the period leading to the crisis of 2007–2008 is addressed in Section 5.2.5.

This by no means implies that historical data should not be utilized. Historical data, at a minimum, supplies a check against intuition and can be used to help form reasoned subjective opinions. But risk managers concerned with protecting a firm against infrequent but plausible outcomes must be ready to employ subjective judgment.

Let us illustrate with a simple example. Suppose you are trying to describe the distribution of a variable for which you have a lot of historical data that strongly supports a normal distribution with a mean of 5 percent and standard deviation of 2 percent. Suppose you suspect that there is a small but nonnegligible possibility that there will be a regime change that will create a very different distribution. Let's say you guess there is a 5 percent chance of this distribution, which you estimate as a normal distribution with a mean of 0 percent and standard deviation of 10 percent.

If all you cared about was the mean of the distribution, this wouldn't have much impact—lowering the mean from 5 percent to 4.72 percent. Even if you were concerned with both mean and standard deviation, it wouldn't have a huge impact: the standard deviation goes up from 2 percent to 3.18 percent, so the Sharpe ratio (the ratio of mean to standard deviation often used in financial analysis) would drop from 2.50 to 1.48. But if you were concerned with how large a loss you could have 1 percent of the time, it would be a change from a gain of 0.33 percent to a loss of 8.70 percent. Exercise 1.1 will allow you to make these and related calculations for yourself using the Excel spreadsheet **MixtureOfNormals** supplied on the book's website.

This illustrates the point that when you are concerned with the tail of the distribution you need to be very concerned with subjective probabilities and not just with objective frequencies. When your primary concern is just the mean—or even the mean and standard deviation, as might be typical for a mutual fund—then your primary focus should be on choosing the most representative historical period and on objective frequencies.

While this example was drawn from financial markets, the conclusions would look very similar if we were discussing an actuarial risk problem like nuclear safety and we were dealing with possible deaths rather than financial losses. The fact that risk managers need to be concerned with managing against extreme outcomes would again dictate that historical frequencies need to be supplemented by informed subjective judgments. This reasoning

is very much in line with the prevailing (but not universal) beliefs among academics in the fields of statistics and decision theory. A good summary of the current state of thinking in this area is to be found in Hammond, Keeney, and Raiffa (1999, Chapter 7). Rebonato (2007) is a thoughtful book-length treatment of these issues from an experienced and respected financial risk manager that reaches conclusions consistent with those presented here (see particularly Chapter 8 of Rebonato).

The importance of extreme events to risk management has two other important consequences. One is that in using historical data it is necessary to pay particular attention to the shape of the tail of the distribution; all calculations must be based on statistics that take into account any nonnormality displayed in the data, including nonnormality of correlations. The second consequence is that all calculations must be carried out using simulation. The interaction of input variables in determining prices and outcomes is complex, and shortcut computations for estimating results work well only for averages; as soon as you are focused on the tails of the distribution, simulation is a necessity for accuracy.

The use of simulation based on both historical data and subjective judgment and taking nonnormality of data into account is a repeated theme throughout this book—in the statement of general principles in Section 6.1.1, applied to more liquid positions throughout Chapter 7, applied to positions involving actuarial risk in Section 8.4, and applied to specific risk management issues throughout Chapters 9 through 14.

EXERCISE

1.1 The Impact of Nonnormal Distributions on Risk

Use the `MixtureOfNormals` spreadsheet to reproduce the risk statistics shown in Section 1.3 (you will not be able to reproduce these results precisely, due to the random element of Monte Carlo simulation, but you should be able to come close). Experiment with raising the probability of the regime change from 5 percent to 10 percent or higher to see the sensitivity of these risk statistics to the probability you assign to an unusual outcome. Experiment with changes in the mean and standard deviation of the normal distribution used for this lower-probability event to see the impact of these changes on the risk statistics.

Institutional Background

A financial firm is, among other things, an institution that employs the talents of a variety of different people, each with her own individual set of talents and motivations. As the size of an institution grows, it becomes more difficult to organize these talents and motivations to permit the achievement of common goals. Even small financial firms, which minimize the complexity of interaction of individuals within the firm, must arrange relationships with lenders, regulators, stockholders, and other stakeholders in the firm's results.

Since financial risk occurs in the context of this interaction between individuals with conflicting agendas, it should not be surprising that corporate risk managers spend a good deal of time thinking about organizational behavior or that their discussions about mathematical models used to control risk often focus on the organizational implications of these models. Indeed, if you take a random sample of the conversations of senior risk managers within a financial firm, you will find as many references to *moral hazard*, *adverse selection*, and *Ponzi scheme* (terms dealing primarily with issues of organizational conflict) as you will find references to *delta*, *standard deviation*, and *stochastic volatility*.

For an understanding of the institutional realities that constitute the framework in which risk is managed, it is best to start with the concept of moral hazard, which lies at the heart of these conflicts.

2.1 MORAL HAZARD—INSIDERS AND OUTSIDERS

The following is a definition of *moral hazard* taken from Kotowitz (1989):

Moral hazard may be defined as actions of economic agents in maximizing their own utility to the detriment of others, in situations where they do not bear the full consequences or, equivalently,

do not enjoy the full benefits of their actions due to uncertainty and incomplete or restricted contracts which prevent the assignment of full damages (benefits) to the agent responsible. . . . Agents may possess informational advantages of hidden actions or hidden information or there may be excessive costs in writing detailed contingent contracts. . . . Commonly analyzed examples of hidden actions are workers' efforts, which cannot be costlessly monitored by employers, and precautions taken by the insured to reduce the probability of accidents and damages due to them, which cannot be costlessly monitored by insurers. . . . Examples of hidden information are expert services—such as physicians, lawyers, repairmen, managers, and politicians.

In the context of financial firm risk, moral hazard most often refers to the conflict between insiders and outsiders based on a double-edged asymmetry. Information is asymmetrical—the insiders possess superior knowledge and experience. The incentives are also asymmetrical—the insiders have a narrower set of incentives than the outsiders have. This theme repeats itself at many levels of the firm.

Let's begin at the most basic level. For any particular group of financial instruments that a firm wants to deal in, whether it consists of stocks, bonds, loans, forwards, or options, the firm needs to employ a group of experts who specialize in this group of instruments. These experts will need to have a thorough knowledge of the instrument that can rival the expertise of the firm's competitors in this segment of the market. Inevitably, their knowledge of the sector will exceed that of other employees of the firm. Even if it didn't start that way, the experience gained by day-to-day dealings in this group of instruments will result in information asymmetry relative to the rest of the firm. This information asymmetry becomes even more pronounced when you consider information relative to the particular positions in those instruments into which the firm has entered. The firm's experts have contracted for these positions and will certainly possess a far more intimate knowledge of them than anyone else inside or outside the firm. A generic name used within financial firms for this group of experts is the *front office*. A large front office may be divided among groups of specialists: those who negotiate transactions with clients of the firm, who are known as *salespeople*, *marketers*, or *structuring*; those who manage the positions resulting from these negotiated transactions, who are known as *traders*, *position managers*, or *risk managers*; and those who produce research, models, or systems supporting the process of decision making, who are known as *researchers* or *technologists*.

However, this group of experts still requires the backing of the rest of the firm in order to be able to generate revenue. Some of this dependence

may be a need to use the firm's offices and equipment; specialists in areas like tax, accounting, law, and transactions processing; and access to the firm's client base. However, these are services that can always be contracted for. The vital need for backing is the firm's ability to absorb potential losses that would result if the transactions do not perform as expected.

A forceful illustration of this dependence is the case of Enron, which in 2001 was a dominant force in trading natural gas and electricity, being a party to about 25 percent of all trades executed in these markets. Enron's experts in trading these products and the web-enabled computer system they had built to allow clients to trade online were widely admired throughout the industry. However, when Enron was forced to declare bankruptcy by a series of financing and accounting improprieties that were largely unrelated to natural gas and electricity trading, their dominance in these markets was lost overnight.

Why? The traders and systems that were so widely admired were still in place. Their reputation may have been damaged somewhat based on speculation that the company's reporting was not honest and its trading operation was perhaps not as successful as had been reported. However, this would hardly have been enough to produce such a large effect. What happened was an unwillingness of trading clients to deal with a counterparty that might not be able to meet its future contractual obligations. Without the backing of the parent firm's balance sheet, its stockholder equity, and its ability to borrow, the trading operation could not continue.

So now we have the incentive asymmetry to set off the information asymmetry. The wider firm, which is less knowledgeable in this set of instruments than the group of front-office experts, must bear the full financial loss if the front office's positions perform badly. The moral hazard consists of the possibility that the front office may be more willing to risk the possibility of large losses in which it will not have to fully share in order to create the possibility of large gains in which it will have a full share. And the rest of the firm may not have sufficient knowledge of the front office's positions, due to the information asymmetry, to be sure that this has not occurred.

What are some possible solutions? Could a firm just purchase an insurance contract against trading losses? This is highly unlikely. An insurance firm would have even greater concerns about moral hazard because it would not have as much access to information as those who are at least within the same firm, even if they are less expert. Could the firm decide to structure the pay of the front office so that it will be the same no matter what profits are made on its transactions, removing the temptation to take excessive risk to generate potential large gains? The firm could, but experience in financial firms strongly suggests the need for upside participation as an incentive to call forth the efforts needed to succeed in a highly competitive environment.

Inevitably, the solution seems to be an ongoing struggle to balance the proper incentive with the proper controls. This is the very heart of the design of a risk management regime. If the firm exercises too little control, the opportunities for moral hazard may prove too great. If it exercises too much control, it may pass up good profit opportunities if those who do not have as much knowledge as the front office make the decisions. To try to achieve the best balance, the firm will employ experts in risk management disciplines such as market risk, credit risk, legal risk, and operations risk. It will set up independent support staff to process the trades and maintain the records of positions and payments (the *back office*); report positions against limits, calculate the daily profit and loss (P&L), and analyze the sources of P&L and risk (the *middle office*); and take responsibility for the accuracy of the firm's books and records (the *finance* function). However, the two-sided asymmetry of information and incentive will always exist, as the personnel in these control and support functions will lack the specialized knowledge that the front office possesses in their set of instruments.

The two-sided asymmetry that exists at this basic level can be replicated at other levels of the organization, depending on the size and complexity of the firm. The informational disadvantage of the manager of fixed-income products relative to the front office for European bonds will be mirrored by the informational disadvantage of the manager of all trading products relative to the manager of fixed-income products and the firm's CEO relative to the manager of all trading products.

Certainly, the two-sided asymmetry will be replicated in the relationship between the management of the firm and those who monitor the firm from the outside. Outside monitors primarily represent three groups—the firm's creditors (lenders and bondholders), the firm's shareholders, and governments. All three of these groups have incentives that differ from the firm's management, as they are exposed to losses based on the firm's performance in which the management will not fully share.

The existence of incentive asymmetry for creditors is reasonably obvious. If the firm does well, the creditors get their money back, but they have no further participation in how well the firm performs; if the firm does very badly and goes bankrupt, the creditors have substantial, possibly even total, loss of the amount lent. By contrast, the firm's shareholders and management have full participation when the firm performs well, but liability in bankruptcy is limited to the amount originally invested. When we examine credit risk in Section 13.2.4, this will be formally modeled as the creditors selling a put option on the value of the firm to the shareholders. Since all options create nonlinear (hence asymmetric) payoffs, we have a clear source of incentive asymmetry for creditors.

It is less clear whether incentive asymmetry exists for shareholders. In principle, their interests are supposed to be exactly aligned with those of the firm's management, and incentives for management based on stock value are used to strengthen this alignment. In practice, it is always possible that management will take more risk than shareholders would be completely comfortable with in the hope of collecting incentive-based compensation in good performance years that does not have to be returned in bad performance years. Kotowitz (1989) quotes Adam Smith from *Wealth of Nations*: "The directors of such companies, however, being managers rather of other people's money than of their own, it cannot well be expected, that they should watch over it with the same anxious vigilance with which the partners in a private company frequently watch over their own."

Government involvement arises from the asymmetric dangers posed to the health of the overall economy by the failure of a financial firm. If an implicit government guarantee is given to rescue large financial firms from bankruptcy (the notion of "too big to fail"), then moral hazard is created through management's knowledge that it can try to create profit opportunities, in which the government has only limited participation through taxes, by taking risks of losses that will need to be fully absorbed by the government. If the government is not willing to prevent the failure of large financial firms, then it will want to place restrictions on the externalities that those firms can create by not having to bear their share of the cost to the overall economy of a firm's potential bankruptcy.

In all three cases of moral hazard involving outside monitors, the information asymmetry is even more severe than when the information asymmetry takes place wholly inside the firm. Senior management and its risk monitors are at least on the premises, are involved in day-to-day business with more junior managers, and can utilize informal measures, such as the rotation of managers through different segments of the firm, to attempt to diffuse both incentives and knowledge. Outside monitors will have only occasional contact with the firm and must rely mostly on formal requirements to obtain cooperation.

Let us look at some of the outside monitors that creditors, shareholders, and governments rely on:

- In addition to their own credit officers, creditors rely on rating agencies such as Moody's Investors Service and Standard & Poor's (S&P) to obtain information about and make judgments on the creditworthiness of borrowers.
- Shareholders and creditors rely on investment analysts working for investment bankers and brokerage firms to obtain information about and make judgments on the future earnings prospects and share values of

firms. Although neither rating agencies nor investment analysts have any official standing with which to force cooperation from the firms they analyze, their influence with lenders and investors in bonds and stocks gives them the leverage to obtain cooperation and access to information.

- Governments can use their regulatory powers to require access to information from financial firms and employ large staffs to conduct examinations of the firms. For example, for the U.S. government, the Federal Reserve System and the Comptroller of the Currency conduct examinations of commercial banks. A similar function is performed by the Securities and Exchange Commission (SEC) for investment banks.
- Creditors, shareholders, and governments all rely on independent accounting firms to conduct audits of the reliability of the financial information disclosures that are required of all publicly held firms.

Over the years, many critical questions have been raised about how truly independent the judgment of these outside monitors really is:

- Credit rating agencies have been accused of being too slow to downgrade ratings in response to adverse changes in a firm's financial condition because their source of revenue comes from the firms whose debt they rate.
- Similarly, independent auditors have been suspected of being too deferential to the firms they monitor since these firms are the ones who pay their audit fees and hire them for consulting services. The fear is that the desire for more revenue will blunt objections to companies choosing accounting methods that cast their results in a favorable light.
- Investment banks have a built-in conflict of interest from competing for the business of the firms whose performance their investment analysts are monitoring. It has long been noted that analysts' buy recommendations far outnumber sell recommendations.
- Accusations have been leveled that government regulatory agencies are more concerned with protecting the interests of the firms being monitored than with protecting the public interest. These charges have particular force when personnel flow freely between employment in the regulatory agencies and in the firms they regulate.

All of these criticisms seemed to be coming to a head in 2002 amid the scandals involving the now-defunct auditing firm of Arthur Andersen, Enron's declaration of bankruptcy only a week after being rated investment grade, and the massive declines in the stock values of technology firms highly touted by investment analysts. Some useful reforms have been undertaken, such as forbidding auditing firms to sell consulting services to firms

they audit and not allowing the bonuses of investment analysts to be tied to investment banking fees collected from clients whose stocks they cover. However, the basic sources of conflict of interest remain, and investors and lenders will continue to need to employ a skeptical filter when utilizing input from outside monitors.

Although the conflicts between insiders and outsiders due to the two-sided asymmetry of moral hazard cannot be eliminated, a frank understanding by both sides can lead to a cooperative relationship. In a cooperative relationship, insiders will acknowledge the need to have outsiders exercise controls and will voluntarily share information and knowledge with outsiders. In a cooperative relationship, outsiders will acknowledge their need to learn from the insiders and will ease controls in response to a track record of openness, although both must recognize the need to always have some level of controls (the ancient folk wisdom states that “I trust my grandmother, but I still cut the cards when she deals”).

A lack of understanding of moral hazard can lead to an uncooperative relationship fueled by mutual resentments between an insider, such as a trader or structurer, with an outsider, such as a corporate risk manager or regulator. An insider who does not understand the purely situational need to have someone less knowledgeable “look over my shoulder” will attribute it to an insulting lack of personal trust, an arrogant assumption of more knowledge than the other possesses, or a simple desire by the outsider to create a job or grab power (which is not to say that some of these motivations do not exist in reality, mixed in with the need to control moral hazard). The insider’s response will then probably be to withhold information, obfuscate, and mislead, which will drive the outsider to even closer scrutiny and more rigid controls, which is clearly a prescription for a vicious circle of escalation. An outsider who lacks an understanding of the situation may defensively try to pretend to have more knowledge than he actually has or may denigrate the knowledge of the insider, which will only exacerbate any suspicions of the process the insider has.

Moral hazard has long been a key concept in the analysis of insurance risks. A typical example would be an insurance company’s concern that an individual who has purchased insurance against auto theft will not exercise as much care in guarding against theft (for example, parking in a garage rather than on the street) as one who has not purchased insurance. If the insurance company could distinguish between individuals who exercise extra care and those who don’t, it could sell separate contracts to the two types of individuals and price the extra losses into just the type sold to those exercising less care. However, the information advantage of an individual monitoring his own degree of care relative to the insurance company’s ability to monitor it makes this prohibitively expensive. So the insurance company

needs to settle for cruder measures, such as establishing a deductible loss that the insured person must pay in the event of theft, thereby aligning the interests of the insured more closely with the insurer.

It has become increasingly common for moral hazard to be cited in analyses of the economics of firms in general, particularly in connection with the impact of the limited liability of shareholders willing to take larger gambles. The shareholders know that if the gamble succeeds, they will avoid bankruptcy and share in the profits, but will suffer no greater loss in a large bankruptcy than in a smaller one. To quote W. S. Gilbert:

*You can't embark on trading too tremendous,
It's strictly fair and based on common sense,
If you succeed, your profits are stupendous,
And if you fail, pop goes your eighteen pence.*
(from Gilbert and Sullivan's Utopia, Limited)

A firm's creditors can exercise some control over their actions and might be able to forbid such gambles, assuming they have sufficient knowledge of the nature of the firm's investments. This is where the informational advantage of the managers over the creditors with respect to the firm's investments comes in.

What sort of actions can we expect from a trader based on the concept of moral hazard? We can certainly expect that the trader may have a different degree of risk aversion than the firm's management, since traders' participation in favorable results exceeds their participation in downside results. Taleb (1997, 66) refers to this as the trader "owning an option on his profits" and states that in such circumstances "it is always optimal to take as much risk as possible. An option is worth the most when volatility is highest." This will probably become even more noticeable if the trader has been having a poor year. Knowing that she is headed toward a minimal bonus and possible dismissal may incline the trader to swing for the fences and take a large risk. The trader knows that if the risk turns out favorably, it might be enough to reverse previous losses and earn a bonus. If it turns out poorly, then "you can't get less than a zero bonus" and "you can't get fired twice." (You can damage your reputation in the industry, but sharing information about a trader's track record between competitor firms cannot be done that efficiently—more information asymmetry.) For this reason, firms may severely cut the trading limits of a trader having a poor year.

Beyond the differences in risk aversion, moral hazard can even result in the perverse behavior (for the firm) of having a trader willing to increase risk

exposure when faced with a lower expected return. Consider the following advice to traders from Taleb (1997, 65):

How aggressive a trader needs to be depends highly on his edge, or expected return from the game:

- *When the edge is positive (the trader has a positive expected return from the game, as is the case with most market makers), it is always best to take the minimum amount of risk and let central limit slowly push the position into profitability. This is the recommended method for market makers to progressively increase the stakes, in proportion to the accumulated profits. In probability terms, it is better to minimize the volatility to cash-in on the drift.*
- *When the edge is negative, it is best to be exposed as little as possible to the negative drift. The operator should optimize by taking as much risk as possible. Betting small would ensure a slow and certain death by letting central limit catch up on him.*

The mathematics and economic incentives that this advice is based on are certainly sound. It is advice that is known to every gambler (or ought to be) and is well founded in statistical theory. When the odds are in your favor, place many small bets; when the odds are against you, place one large bet. Essentially, when the odds are against you, you are attempting to minimize the length of time you are playing against the house since you are paying a tax, in the form of an expected loss, for the privilege of playing.

However, although this makes perfect economic sense from the viewpoint of the individual trader, it is hardly the strategy the firm employing these traders would want to see them follow. The firm, whose P&L will be the sum of the results of many traders, would like to see traders with a negative expected return not take any positions at all rather than have these be the traders taking on the most risk. To the extent the firm's management can figure out which traders have a negative edge, it will restrict their risk taking through limits and the replacement of personnel. However, the individual traders have the information advantage in knowing more than the firm about their expected returns. They also have the asymmetrical incentive to take larger risks in this case, even though doing so will probably hurt the firm. The traders will not derive much benefit from the firm doing well if they do not contribute to that result, but they will benefit if they do increase their risk and win against the odds.

Moral hazard helps to explain the valuation that investors place on the earnings volatility of financial firms. You could argue that firms should

worry just about the expected value and not about volatility, since the market should place a risk premium only on risk that it cannot hedge away (an investor who wants less risk will just take the stock with the highest expected return and diversify by mixing with government bonds). However, empirical evidence shows that the market places a stiff discount on variable trading earnings. The reason may be information asymmetry. It is hard for outsiders to tell whether a firm is taking sound gambles to maximize expected value or is maximizing its insiders' option on one-way bets. Perold (1998) states:

I view financial intermediaries as being special in several ways: First, these firms are in credit-sensitive businesses, meaning that their customers are strongly risk-averse with respect to issuer default on contractually promised payoffs. (For example, policyholders are averse to having their insurance claims be subject to the economic performance of the issuing firm, and strictly prefer to do business with a highly rated insurer.) The creditworthiness of the intermediary is crucial to its ability to write many types of contracts, and contract guarantees feature importantly in its capital structure.

Second, financial firms are opaque to outsiders. They tend to be in businesses that depend vitally on proprietary financial technology and that cannot be operated transparently. In addition, the balance sheets of financial firms tend to be very liquid, and are subject to rapid change. Financial firms, thus, are difficult to monitor, and bear significant deadweight costs of capital. Guarantors face costs related to adverse selection and moral hazard. . . .

Third, financial firms are also internally opaque. Information tends to be private at the business unit level, or even at the level of individual employees such as traders. Efficient management of these firms thus involves significant use of performance-related compensation to mitigate against monitoring difficulty.

Moral hazard can create a battleground over information between insiders and outsiders. Insiders are fearful that any information obtained by outsiders will be used as a tool to tighten controls over insiders' actions. Insiders can be expected to have an inherent bias against tighter controls, partly because narrowing the range of actions available leads to suboptimal solutions and partly because incentive asymmetry makes riskier action more rewarding to insiders than to outsiders. One of the most common ways in which insiders can mislead outsiders about the need for controls is termed a Ponzi scheme.

2.2 PONZI SCHEMES

In its original meaning, a *Ponzi scheme* is a criminal enterprise in which investors are tricked into believing that they will receive very high returns on their investments, but the early investors are paid out at high rates of return only with the payments coming from the cash invested by later investors. The illusion of high returns can be pretty convincing. After all, you can actually see the early investors receiving their high returns in cash, and the con men running these schemes can produce very plausible lies about the purported source of the returns. As a result, the pace of new investment can be intense, enabling the illusion of profit to be maintained over a fairly long time period. It's a vicious cycle—the eagerness of new investors to place money in the scheme leads to the heightened ability to make investments appear highly profitable, which leads to even greater eagerness of new investors. However, ultimately, any Ponzi scheme must collapse, as there is no ultimate source of investment return (in fact, investment return is quite negative, as the flow of new investment must also be partially diverted to the criminals profiting from it). Ponzi schemes are also sometimes called *pyramid schemes* and bear a close resemblance to chain letter frauds.

When I wrote the immediately preceding paragraph for the first edition of this book in 2003, I felt the need to thoroughly explain what a Ponzi scheme is. Today, it is probably not necessary, as Bernie Madoff has regrettably given us all an exhaustive lesson in how a Ponzi scheme is run.

The original meaning of Ponzi schemes has been broadened by risk managers to include situations in which firms are misled as to the profitability of a business line by the inadequate segregation of profits on newly acquired assets and returns on older assets.

Let's consider a typical example. Suppose a trading desk has entered into marketing a new type of path-dependent option. The desk expects substantially more customer demand for buying these options than for selling them. They intend to manage the resulting risk with dynamic hedging using forwards and more standard options. As we will see when discussing path-dependent options in Section 12.3, it is very difficult to try to estimate in advance how successful a dynamic hedging strategy for path-dependent options will be.

In such circumstances, the pricing of the option to the client must be based on an estimate of the future cost of the dynamic hedging, applying some conservatism to try to cover the uncertainty. Let's assume that a typical trade has a seven-year maturity, and that the customer pays \$8 million and the firm pays \$5 million to purchase the initial hedge. Of the remaining \$3 million, we'll assume that the desk is estimating dynamic hedging costs of \$1 million over the two years, but the uncertainty of these costs leads to

setting up a \$2 million initial allowance (or reserve) to cover the hedging costs, leaving \$1 million to be booked as up-front profit.

Suppose the trading desk has made a serious error in predicting the hedging costs, and the hedging costs actually end up around \$5 million, leading to a net loss of \$2 million on every transaction booked. You may not be able to do anything about deals already contracted, but you would at least hope to get feedback from the losses encountered on these deals in time to stop booking new deals or else raise your price to a more sustainable level. This should happen if P&L reporting is adequately detailed, so you can see the losses mounting up on the hedging of these trades (this is called *hedge slippage*).

However, it is often difficult to keep track of exactly how to allocate a day's trading gains and losses to the book of deals being hedged. You want to at least know that trading losses are occurring so you can investigate the causes. The most severe problem would be if you didn't realize that trades were losing money. How could this happen? If P&L reporting is not adequately differentiated between the existing business and new business, then the overall trading operation can continue to look profitable by just doing enough new business. Every time a new deal is booked, \$1 million goes immediately into P&L. Of course, the more deals that are booked, the larger the hedging losses that must be overcome, so even more new trades are needed to swamp the hedging losses. The resemblance to a Ponzi scheme should now be obvious.

One key difference is that in its original meaning, the Ponzi scheme is a deliberate scam. The financial situation described is far more likely to arise without any deliberate intent. However, those in the front office, based on their close knowledge of the trading book, will often suspect that this situation exists before any outsiders do, but may not want to upset the apple cart. They would be jeopardizing bonuses that can be collected up front on presumed earnings. They may also be willing to take the risk that they can find a way to turn the situation around based on their greater participation in future upside than future downside. They may choose to hide the situation from outsiders who they suspect would not give them the latitude to take such risks. So moral hazard can turn an accidentally originated Ponzi scheme into one that is very close to deliberate.

As a historical footnote, the Ponzi scheme derives its name from Charles Ponzi, a Boston-based swindler of the 1920s (though it was not the first Ponzi scheme—William “520 Percent” Miller ran one in Brooklyn around 1900; an excellent 1905 play by Harley Granville-Barker, *The Vosey Inheritance*, which has been revived frequently over the past decade, revolves around a lawyer specializing in trusts and estates trying to train his son to take over

the management of his Ponzi scheme). The following account of Charles Ponzi is drawn from Sifakis (1982):

[Ponzi] discovered he could buy up international postal-union reply coupons at depressed prices and sell them in the United States at a profit up to 50 percent. It was, in fact, a classic get-rich-slowly operation, and as such, it bored Ponzi. So he figured out a better gimmick.

Ponzi figured out that telling people he was making the money and how he could make it was just as good as actually making it. He advertised a rate of return of 50 percent in three months. It was an offer people couldn't refuse, and money started to come rolling in.

When Ponzi actually started paying out interest, a deluge followed. On one monumental day in 1920, Ponzi's offices took in an incredible \$2 million from America's newest gamblers, the little people who squeezed money out of bank accounts, mattresses, piggy banks, and cookie jars. There were days when Ponzi's office looked like a hurricane had hit it. Incoming cash had to be stuffed in closets, desk drawers and even wastebaskets. Of course, the more that came in, the more Ponzi paid out.

As long as new funds were coming in, Ponzi could continue to make payments. However, as with all pyramid schemes, the bubble had to burst. A newspaper published some damaging material about his past, including time spent in prison. New investors started to hesitate.

Ponzi's fragile scheme collapsed, since it required an unending flow of cash. His books, such as they were, showed a deficit of somewhere between \$5 and \$10 million, or perhaps even more. No one ever knew for sure.

2.3 ADVERSE SELECTION

Let's return to the situation described previously. Suppose our accounting is good enough to catch the hedge slippage before it does too much damage. We stop booking new deals of this type, but we may find we have booked a disturbingly large number of these deals before the cutoff. If our customers have figured out the degree to which we are underpricing the structure before we do, then they may try to complete as many deals as they can before we wise up. This pattern has frequently been seen in the financial markets. For

example, the last firms that figured out how to correctly price volatility skew into barrier options found that their customers had loaded up on trades that the less correct models were underpricing. A common convention is to label this situation as *adverse selection* as a parallel to a similar concern among insurance firms, which worry that those customers with failing health will be more eager to purchase insurance than those with better health, taking advantage of the fact that a person knows more about his own health than an insurance company can learn (Wilson 1989). So adverse selection is like moral hazard since it is based on information asymmetry; the difference is that moral hazard is concerned with the degree of risk that might be taken based on this asymmetry, whereas adverse selection is concerned with a difference in purchasing behavior. In 2001, George Akerlof, Michael Spence, and Joseph Stiglitz won the Nobel Prize in economics for their work on adverse selection and its application to a broad class of economic issues.

Concern about the risk from adverse selection motivates risk managers' concern about the composition of a trading desk's customer base. The key question is: What proportion of trades is with counterparties who are likely to possess an informational advantage relative to the firm's traders? As a general rule, you prefer to see a higher proportion of trades with individuals and nonfinancial corporations that are likely trading to meet hedging or investment needs rather than seeking to exploit informational advantage. Alarm is raised when an overwhelming proportion of trades is with other professional traders, particularly ones who are likely to see greater deal flow or have a greater proportion of trades with individuals and nonfinancial corporations than your firm's traders. Seeing greater deal flow can give a firm an informational advantage by having a more accurate sense of supply-and-demand pressures on the market. A greater proportion of customers who are not professional traders yields two further potential informational advantages:

1. At times you work with such customers over a long period of time to structure a large transaction. This gives the traders advance knowledge of supply and demand that has not been seen in the market yet.
2. Working on complex structures with customers gives traders a more intimate knowledge of the structure's risks. They can choose to retain those risks that this knowledge shows them are more easily manageable and attempt to pass less manageable risks on to other traders.

Traders may tend to underestimate the degree to which their profitability is due to customer deal flow and overestimate the degree to which it is due to anticipating market movements. This can be dangerous if it encourages them to aggressively take risks in markets in which they do not possess this customer flow advantage. A striking example I once observed was a foreign

exchange (FX) trader who had a phenomenally successful track record of producing profits at a large market-making firm. Convinced of his prowess in predicting market movements, he accepted a lucrative offer to move to a far smaller firm. He was back at his old job in less than year, confessing he simply had not realized how much of his success was due to the advantages of customer deal flow.

A pithy, if inelegant, statement of this principle was attributed to the head of mortgage-backed trading at Kidder Peabody: “We don’t want to make money trading against smart traders; we want to make money selling to stupid customers.” Of course, *stupid* needs to be understood here as macho Wall Street lingo for *informationally disadvantaged*. It’s the sort of talk that is meant to be heard only in locker rooms and on trading floors. An unfriendly leak resulted in his quote appearing on the front page of the *Wall Street Journal*. It is delightful to imagine the dialogue of some of his subsequent conversations with the firm’s customers.

2.4 THE WINNER’S CURSE

In response to the risks of adverse selection, traders may exhibit confidence that this is not something they need to worry about. After all, adverse selection impacts only those with less knowledge than the market. It is a rare trader who is not convinced that she possesses far more knowledge than the rest of the market—belief in one’s judgment is virtually a necessity for succeeding in this demanding profession. Whether the firm’s management shares the trader’s confidence may be another story. However, even if it does, the trader must still overcome another hurdle—the *winner’s curse*, the economic anomaly that says that in an auction, even those possessing (insider) knowledge tend to overpay.

The winner’s curse was first identified in conjunction with bidding for oil leases, but has since been applied to many other situations, such as corporate takeovers. My favorite explanation of the mechanism that leads to the winner’s curse comes from Thaler (1992):

Next time you find yourself a little short of cash for a night on the town, try the following experiment in your neighborhood tavern. Take a jar and fill it with coins, noting the total value of the coins. Now auction off the jar to the assembled masses at the bar (offering to pay the winning bidder in bills to control for penny aversion). Chances are very high that the following results will be obtained:

1. *The average bid will be significantly less than the value of the coins. (Bidders are risk averse.)*

2. *The winning bid will exceed the value of the jar.*

In conducting this demonstration, you will have simultaneously obtained the funding necessary for your evening's entertainment and enlightened the patrons of the tavern about the perils of the winner's curse.

When applied to trading, the winner's curse is most often seen in market making for less liquid products, where opinions on the true value of a transaction may vary more widely. Market makers are in competition with one another in pricing these products. The firm that evaluates a particular product as having a higher value than its competition is most likely to be winning the lion's share of these deals. Consider a market for options on stock baskets. As we will discuss in Section 12.4, a liquid market rarely exists for these instruments, so pricing depends on different estimates of correlation between stocks in a basket. The firm that has the lowest estimate for correlation between technology stocks will wind up with the most aggressive bids for baskets of technology stocks and will book a large share of these deals. Another firm that has the lowest estimate for correlation between financial industry stocks will book the largest share of those deals.

An anecdotal illustration comes from Neil Chriss. When Chriss was trading volatility swaps at Goldman Sachs, they would line up five or six dealers to give them quotes and would always hit the highest bid or lift the lowest offer. The dealers knew they were doing this and were very uneasy about it, limiting the size of trades they would accommodate. One dealer, on winning a bid, told Chriss, "I am always uncomfortable when I win a trade with you, as I know I was the best bid on top of five other smart guys. What did I do wrong?"

Adverse selection can be controlled by gaining expertise and increasing the proportion of business done with ultimate users rather than with other market makers. However, the winner's curse can be controlled only by either avoiding auction environments or adequately factoring in a further pricing conservatism beyond risk aversion. It provides a powerful motivation for conservatism in pricing and recognizing profits for those situations such as one-way markets (see Section 6.1.3) in which it is difficult to find prices at which risks can be exited.

We demonstrate the mechanism of the winner's curse with a simple numerical example involving a market with only three firms, two buyers, and one seller. The results are shown in Table 2.1.

We consider two different situations. In the first, direct negotiation occurs on the price between the seller and a single buyer. In the second, both buyers participate in an auction.

TABLE 2.1 The Winner's Curse

Deal	Actual Value	Seller's Offer	Buyer 1		Buyer 2		Auction P&L	
			Bid	P&L	Bid	P&L	Buyer 1	Buyer 2
1	1.56	2.10	2.00	0.00	2.40	-0.69	0.00	-0.84
2	2.66	1.40	1.50	1.21	1.80	1.06	0.00	0.86
3	3.16	3.20	2.10	0.00	3.10	0.00	0.00	0.06
4	1.96	3.40	1.60	0.00	2.20	0.00	0.00	-0.24
5	1.36	1.90	1.70	0.00	1.20	0.00	-0.34	0.00
6	4.46	4.30	4.80	-0.09	3.00	0.00	-0.34	0.00
7	3.16	2.60	3.50	0.11	2.90	0.41	-0.34	0.00
8	1.96	1.80	1.40	0.00	2.10	0.01	0.00	-0.14
9	1.56	2.70	1.40	0.00	1.10	0.00	0.16	0.00
10	2.16	2.10	2.50	-0.14	2.70	-0.24	0.00	-0.54
Average	2.40	2.55	2.25		2.25			
Total				1.09		0.55	-0.86	-0.84
Correlation with actual value	63.2%	83.3%		72.2%				

Notes

The column headed Actual Value shows what the deals are really worth.

The columns headed Buyer 1 Bid and Buyer 2 Bid are the bid price of buyers of deals.

The seller's asking prices, in the column headed Seller's Offer, are conservative (higher by .15 on average than true prices).

The buyers are posting conservative bids (lower by .15 on average than true prices).

The columns headed Buyer 1 P&L and Buyer 2 P&L are the buyers' profits if they negotiate to the average of the seller's asked and buyer's bid.

The columns headed Auction P&L Buyer 1 and Auction P&L Buyer 2 show the profits of the buyers in an auction.

There are 10 transactions that the seller might sell to the buyers. Neither the buyers nor the seller is certain of the true value of these transactions (for example, they might depend on future dynamic hedging costs, which depend on the evolution of future prices, which different firms estimate using different probability distributions). After the fact, we know the true realized value of each transaction, as shown in column 2 of the

table. Buyer 1's knowledge of this market is superior to buyer 2's, and both have superior knowledge compared to the seller. This can be seen by the correlations between realized value and each party's estimate of transaction value (83.3% for buyer 1, 72.2% for buyer 2, and 63.2% for the seller). The consequences of this informational advantage are that both buyer 1 and buyer 2 make a profit at the expense of the seller in direct negotiations, and that buyer 1's profit in this situation is higher than buyer 2's profit.

In the direct negotiation situation, we assume that the buyer, being risk averse, has successfully biased his bids down to be on average lower than the realized value, and the seller, being risk averse, has successfully biased his asked prices up to be on average higher than realized value. We assume no transaction takes place if the buyer's bid is lower than the seller's asked. If the buyer's bid exceeds the seller's asked, we assume the transaction takes place at the average price between these two prices. As a result, buyer 1 has a total P&L of +1.09, and buyer 2 has a total P&L of +0.55.

Now consider what happens in the auction when the buyers have to compete for the seller's business, a situation very typical for market making firms that must offer competitive price quotations to try to win customer business from other market makers. The seller no longer relies on his own estimate of value, but simply does business at the better bid price between the two firms. Even though both firms continue to successfully bias their bids down on average from realized values, both wind up losing money in total, with buyer 1 having a P&L of -0.86 and buyer 2 having a P&L of -0.84 . This is because they no longer have gains on trades that they seriously undervalued to balance out losses on trades that they seriously overvalued, since they tend to lose trades that they undervalue to the other bidder. This illustrates the winner's curse.

The spreadsheet **WinnersCurse** on the course website shows the consequences of changing some of the assumptions in this example.

2.5 MARKET MAKING VERSUS POSITION TAKING

An important institutional distinction between participants in the financial markets that we will refer to on several occasions throughout this book is between *market making* and *position taking*:

- Market making (also called *book running* or the *sell side*) consists of making two-way markets by engaging in (nearly) simultaneous buying and selling of the same instruments, attempting to keep position

holdings to a minimum and to profit primarily through the difference between (nearly) simultaneous buy and sell prices.

- Position taking (also called *market using*, *price taking*, *speculation*, or the *buy side*) consists of deliberately taking positions on one side or the other of a market, hoping to profit by the market moving in your favor between the time of purchase and the time of sale. Positions may be taken on behalf of a firm (in which case it is often labeled *proprietary trading*) or on behalf of an individual client or a group of clients, such as a mutual fund, hedge fund, or managed investment account.

Some time lag nearly always occurs between the purchase and sale involved in market making. Depending on the length of time and degree of deliberate choice of the resulting positions, these may be labeled position-taking aspects of market making. Market making almost always involves risk because you cannot often buy and sell exactly simultaneously. The market maker makes a guess on market direction by its posted price, but the bid-ask spread can outweigh even a persistent error in directional guess as long as the error is small. (In Exercise 9.1, you'll be asked to build a simulation to test out the degree to which this is true.) The experience and information gained from seeing so much flow means you most likely will develop the ability to be right on direction on average. However, the position taker has the advantage over the market maker of not needing to be in the market every day. Therefore, the position taker can stay away from the market except when possessed of a strong opinion. The market maker cannot do this; staying away from the market would jeopardize the franchise.

The different objectives of market makers and position takers tend to be reflected in different attitudes toward the use of models and valuation techniques. A position taker generally uses models as forecasting tools to arrive at a best estimate of what a position will be worth at the conclusion of a time period tied to an anticipated event. The position taker will pay attention to the market price of the position during that time period to determine the best time to exit the position and to check whether new information is coming into the market. However, a position taker will generally not be overly concerned by prices moving against the position. Since the position taker is usually waiting for an event to occur, price movements prior to the time the event is expected are not that relevant. A frequently heard statement among position takers is: "If I liked the position at the price I bought it, I like it even better at a lower price."

By contrast, a market maker generally uses models to perform risk decomposition in order to evaluate alternative current prices at which a position can be exited. The market maker will pay close attention to current market prices as the key indicator of how quickly inventory can be reduced.

The direction in which prices will move over the longer term is of little concern compared to determining what price will currently balance supply and demand.

An amusing analogy can be made to gambling on sports. Position takers correspond to the gamblers who place their bets based on an analysis of which team is going to win and by what margin. Market makers correspond to the bookmakers whose sole concern is to move the odds quoted to a point that will even out the amount bet on each side. The bookmaker's concern is not over which team wins or loses, but over the evenness of the amounts wagered. Close to even amounts let the bookmakers come out ahead based on the spread or *vigorish* in the odds, regardless of the outcome of the game. Uneven amounts turn the bookmaker into just another gambler who will win or lose depending on the outcome of the game.

As explained in Section 1.1, the focus of this book is on the active use of trading in liquid markets to manage risk. This view is more obviously aligned with market making than with position taking. In fact, the arbitrage-based models that are so prominent in mathematical finance have been developed largely to support market making. Position takers tend more toward the use of econometric forecasting models. In Section 6.1.7, we will further discuss the issue of the extent to which position takers should adopt the risk management discipline that has been developed for market makers.

Some authors distinguish a third type of financial market participant besides market makers and position takers—the *arbitrageurs*. I believe it is more useful to classify arbitrage trading as a subcategory of position taking. Pure arbitrage, in its original meaning of taking offsetting positions in closely related markets that generate a riskless profit, is rarely encountered in current financial markets, given the speed and efficiency with which liquid prices are disseminated. What is now labeled arbitrage is almost always a trade that offers a low but relatively certain return. The motivations and uses of models by those seeking to benefit from such positions are usually closely aligned with other position takers.

A good example is *merger arbitrage* (sometimes misleadingly called *risk arbitrage*). Suppose that Company A and Company B have announced a forthcoming merger in which two shares of A's stock will be traded for one share of B's stock. If the current forward prices of these stocks to the announced merger date are \$50 for A and \$102 for B, an arbitrage position would consist of a forward purchase of two shares of A for \$100 and a forward sale of one share of B for \$102. On the merger date, the two shares of A purchased will be traded for one share of B, which will be delivered into the forward sale. This nets a sure \$2, but only if the merger goes through as announced. If the merger fails, this trade could show a substantial loss.

Merger arbitrageurs are position takers who evaluate the probability of mergers breaking apart and study the size of loss that might result. They are prototypical forecasters of events with generally little concern for market price swings prior to the occurrence of the event.

For further reading on the economics and institutional structure of market making and position taking, a book I would recommend very highly is Harris (2003). Anyone involved in risk management should attempt to gain insight into how risk management is viewed by traders. While friendship and conversation are the best way to approach this, it is also helpful to read about risk management from a trader's perspective. The best book of this type I have encountered is Brown (2012).

Operational Risk

Operational risk is usually defined in the negative—it includes all of the risks that are not categorized as either market or credit risk. The industry does not yet have consensus on this terminology. Some firms use the term operational risk to cover a subset of the risks other than market and credit risk. For further discussion, see Jameson (1998a). Broadly speaking, these risks are the most difficult to quantify.

One attempt at a more positive definition that has been gaining some currency has been made by the Basel Committee on Banking Supervision: “the risk of direct or indirect loss resulting from inadequate or failed internal processes, people, or systems, or from external events.” Another attempt would be to break apart risk into three pieces. View a financial firm as the sum total of all the contracts it enters into. The firm can suffer losses on the contracts in one of three ways:

1. Obligations in contracts may be performed exactly as expected, but changes in economic conditions might make the sum of all contracted actions an undesired outcome. This is market risk.
2. The other parties to some of the contracts may fail to perform as specified. This is credit risk.
3. The firm may be misled about what the contracted actions are or the consequences of these actions. This is operational risk.

Operational risk is virtually all risk that cannot be managed through the use of liquid markets, so, as argued in Chapter 1, it does not fall within the scope of financial risk management. In this way, it is very much like the risks traditionally managed by insurance companies. Indeed, one of the primary tools for managing operational risk is to try to buy protection from insurance companies, as we’ll discuss in Section 3.8. But, even though the financial risk management approach does not apply, these are risks that arise as

a result of trading and so are intertwined with financial risk management, justifying a quick survey of these issues in this book.

Operational risk can be subdivided into the following categories:

- **Operations risk** is the risk that deficiencies in information systems or internal controls will result in unexpected loss. Operations risk can be further subdivided into the risk of fraud, risk of nondeliberate incorrect information, disaster risk, and personnel risk.
- **Legal risk** is the risk that the terms or conditions of a contract or agreement will prove unenforceable due to legal defects in the contract or in related documentation and procedures. Another type of legal risk is the risk that actions of the firm's employees will have been found to be illegal and subject the firm to substantial penalties. Legal risk includes regulatory risk.
- **Reputational risk** is the risk that the enforcement of contract provisions will prove too costly in terms of damage to the firm's reputation as a desirable firm for customers to do future business with.
- **Accounting risk** is the risk that an error in accounting practice will necessitate a restatement of earnings, which adversely affects the investors' or customers' perception of the firm.
- **Funding liquidity risk** is the risk that an institution will have to pay higher than prevailing market rates for its funding due to either the investors' perception that the credit quality of the institution is impaired (possibly due to earnings problems or capital structure problems) or the overly heavy use of particular funding sources within a given time period, with the large size of transactions impacting funding cost.
- **Enterprise risk** is the risk of loss due to change in the overall business climate, such as the needs of customers, actions of competitors, and pace of technological innovation.

This chapter briefly discusses each of these risks and possible controls, and then presents an overview of how these risks can be identified and the extent to which they can be quantified.

A valuable source of ideas on operational risk and control procedures is the *Trading and Capital-Markets Activities Manual* of the Federal Reserve System. I have used it as a foundation for several of the points in this chapter and recommend that readers interested in this topic look closely at the following sections: 2050.1 and 2060.1 (Operations and Systems Risk), 2070.1 (Legal Risk), 2150.1 (Ethics), 3005.1 (Funding Liquidity Risk), and 2040.1 (the subsection on New Products).

3.1 OPERATIONS RISK

Operations risk can be further subdivided into the risk of fraud, risk of non-deliberate incorrect information, disaster risk, and personnel risk.

3.1.1 The Risk of Fraud

The actual diversion of cash can take the form of creating unauthorized payments, conducting transactions at prices that are not the best available in return for bribes, or utilizing one's position to engage in profitable personal trading at the expense of the firm's profits.

Deception about earnings, in order to generate unearned bonuses or further one's career (or simply avoid being fired), can take the form of recording trades at incorrect prices or misreporting the current value of positions. We'll encounter examples of such deceptions that occurred at Kidder Peabody, Barings, Allied Irish Bank (AIB), and Société Générale in Section 4.1. Section 4.1, covering financial disasters that were due to misleading reporting, should be read in conjunction with this section.

Deception about positions, in order to appear to be operating within limits when an individual is actually outside them or to mislead management about the size of positions being taken, is done in order to preserve freedom of action—avoiding requirements to close down positions. This can be because a trader has a different belief about market movements than management or a different view toward risk than management (the moral hazard issue discussed in Section 2.1). Deception about positions can entail the outright misreporting of positions through the failure to enter transactions (*tickets in the drawer*) or manipulation of management reporting, or hiding positions by arranging for them to be temporarily held by another party with an unrecorded promise to take the position back (*parking*).

Going back 30 years or so, the oral tradition within control functions was to worry about position falsification primarily by traders simply not entering some of their trades onto the firm's books and records (tickets in the drawer). This approach seems to be on the decline, presumably because the possibility of the fraud being exposed through an inquiry from a counterparty to an unrecorded trade is too great. What seems to have replaced it is the entry of fictitious trades designed either to offset the risk position of actual trades, making the net risk look small, or to create bogus profit and loss (P&L) to disguise actual earnings. Since fictitious trades lack a real counterparty, they cannot be exposed through action of a counterparty but can be uncovered only by internal controls. The creator of fraud is in an ongoing battle with control personnel—the control personnel have the

advantage that uncovering only one clear-cut case of falsification is enough to uncover the fraud; the advantage of the creator of the fraud is the multiplicity of methods the creator can employ to discourage this discovery.

The most fundamental control for preventing fraud is by separating the responsibilities between the front office and the support staff (middle office, back office, and controllers), making sure that all entries of transactions and management reporting systems are under the complete control of the support staff. To make this separation of responsibilities work, the support staff must have a separate line of reporting from the front-office staff and compensation that is reasonably independent of the reported earnings of the business area being supported. As much as possible, the reporting lines and compensation structure should align support staff interests with those of management rather than with those of the front office. However, even the best-designed structures of this type are subject to pressures in the direction of alignment of support staff interests with front-office interests. Constant vigilance is required to fight against this. These pressures include:

- Support staff compensation cannot be completely independent from trading performance. At a minimum, unsuccessful results for trading may lead to the shrinking or elimination of a trading operation along with associated support staff positions. Since trading profits are the ultimate source from which expenses get paid, it is difficult to avoid some linkage between the trading performance and level of compensation. Section 4.1.4 presents a vivid example of how this pressure was felt in practice at AIB.
- Front-office personnel almost always command higher compensation and prestige than members of the support staff, usually considerably higher. Often, support staff members are hoping to eventually move into front-office positions. Front-office staff can afford to offer informal incentives to the support staff for cooperation such as helping them seek front-office jobs, giving access to perks such as lavish meals and free tickets to otherwise unavailable sports events, and even offering outright cash bribes. The higher prestige of front-office positions and the reality of the greater market experience of front-office personnel relative to support personnel can be utilized to place tremendous pressure on the support staff to adopt front-office views.
- Since the support staff has responsibilities for supporting the front office as well as for supporting management, their ratings for job performance are often heavily dependent on the views of front-office personnel, who are likely to be working far more closely with them than management personnel.

In addition to the separation of responsibilities, controls include:

- Support staff procedures should be thoroughly documented. Making these as unambiguous as possible lessens the scope for front-office influence.
- Trader lines should be recorded to create a potential source for spotting evidence of collusion with brokers or traders at other firms.
- Make sure that trades are entered into the firm's systems as close to execution time as possible. The further away from execution time you get, the greater the possibility that subsequent market movements will create a temptation to hide or otherwise misrepresent the transaction.
- Review all trades to look for prices that appear off-market, and perform a thorough investigation of any trades identified as such.
- Make sure that all market quotes used to value positions come into support staff, not front-office personnel, and are polled from as large a universe of sources as possible.
- Provide daily explanations of profit and loss (P&L) change and cash needs produced by the support staff. Incorrect reporting of positions can often be identified by the inability to explain P&L and cash movements based on the reported positions.
- Every customer confirmation of a new trade or a payment required by a previous trade should be reviewed by the support staff for consistency with transactions and positions being reported. All customer complaints should be reviewed by the support staff, not just front-office personnel. The confirmation process should be conducted only with support personnel at other firms, not with front-office personnel at other firms.
- Have clear policies about unacceptable practices that are consistently enforced. Deliberate actions to hide a position must entail strong penalties, without exceptions for star traders.
- Personal trading of both front-office and support personnel should be closely monitored.
- Tight controls should be placed on after-hours and off-premises trading to ensure that transactions cannot be omitted from the firm's records.
- Broker usage should be monitored for suspicious patterns—undue concentrations of business that might be compensation for supplying off-market quotes or direct bribery.
- Firms should insist on performing thorough background checks of a potential customer's creditworthiness and business reputation before entering into transactions. In other words, they should refuse to deal with customers they do not know, even on a fully collateralized basis. Unknown customers could be in collusion with the firm's personnel for off-market trading or parking.

- Systems security measures should be in place to ensure that no one other than authorized support personnel can make entries or changes to management information systems. In particular, no front-office personnel should have such access.
- The firm's auditors should perform a periodic review of all operating procedures.
- Control functions must budget some spare capacity for investigative work. The advantage of the control functions relative to an attempt at concealing positions is that only one clear-cut instance of concealment needs to be uncovered to expose a fraud and that any significant attempt at concealment will create a large number of warning indicators that can potentially trigger an investigation. But the disadvantage of control functions relative to an attempt at concealing positions is that a skillful perpetrator of fraud can be expected to be adroit at offering superficially plausible explanations of unusual patterns.

If an investigation is being done in a little spare time of control personnel with a full plate of daily responsibilities, it will be too easy for them to try to wrap it up quickly by accepting a plausible explanation. If control personnel have some budgeted time for conducting such investigations and know that the thoroughness of performance of this task will be part of their job evaluations, it is far more likely that they will perform the extra work needed to uncover the true situation. This will lead to other benefits as well, since thorough investigation of unusual patterns may turn up other gaps in the control system, such as the need for new risk measures or accidental errors in recording positions.

- Control functions should maintain some central registry of investigations they have conducted (along with outcomes). Even if a perpetrator of fraud has been successful in fooling control personnel in several investigations, the unusual number of investigations that the perpetrator's activity is engendering may itself be a clue that leads to a more thorough investigation.
- An overriding concern must be to protect control personnel against bullying. To have a good chance of uncovering frauds, investigations need to be launched based on warning indicators that will create many false positives. This means that the control personnel will go into the investigation knowing that the most likely outcome will be a finding that nothing is wrong. The very fact that they are conducting an investigation is likely to be resented by traders (as a waste of time and as an indicator that their honesty is being questioned). If control personnel feel they are going to be berated by the traders when their investigation finds nothing wrong, then they are likely to conduct fewer and more superficial investigations. Trading management needs to make sure that they give

control personnel the proper backing and try to explain to traders the motivation for such investigations. The careful documentation of major incidents of fraud and the difficulty in detecting them can provide trading managers with tools to use in making this case.

3.1.2 The Risk of Nondeliberate Incorrect Information

It is far more common to have incorrect P&L and position information due to human or systems error than incorrect P&L and position information due to fraud. Many of the controls for nondeliberate incorrect information are similar to the controls for fraud. The separation of responsibilities is effective in having several sets of eyes looking at the entry of a trade, reducing the chance that a single individual's error will impact positions. Checking confirmations and payment instructions against position entries, P&L and cash reconciliation, and the investigation of off-market trades are just as effective in spotting inadvertent errors as they are in spotting fraudulent entries. Equally close attention needs to be paid to making sure customers have posted collateral required by contracts to avoid inadvertently taking unauthorized credit risk. (For further discussions of the role of collateral in managing credit risk, see Sections 4.1.1, 10.1.4, 14.2, and 14.3.3.)

It is every bit as important to have front-office personnel involved in reconciliation (to take advantage of their superior market knowledge and intuitive feel for the size of their P&L and positions) as it is to have support personnel involved (to take advantage of their independence). Front-office personnel must be held responsible for the accuracy of the records of their P&L and positions, and cannot be allowed to place all the blame for incorrect reports on support personnel, in order to ensure that they will place sufficient importance on this reconciliation. Front offices should be required to produce daily projections of closing positions and P&L moves based on their own informal records, prior to seeing the official reports of positions and P&L, and should reconcile significant differences between the two.

To prevent incorrect P&L and position information, it is important to ensure that adequate support personnel and system resources are available, both in quantity and in quality, relative to the size and complexity of trading. Careful attention needs to be paid to planning staff and system upgrades to anticipate growth in trading volume. Management needs to be ready to resist premature approval of a new business if support resources cannot keep pace with front-office development.

Should model risk be regarded as an operations risk issue? The viewpoint of this book is that model risk is primarily a market risk issue, since the proper selection and calibration to market prices of models and the

provision for adequate reserves against model uncertainty are best dealt with by the market risk discipline. Chapter 6 will elucidate this view. However, the proper implementation of models and the assurance that system changes are undertaken with the proper controls are best dealt with by the operations risk discipline. An area independent of model and system developers and the front office should be established to perform quality assurance testing of system implementation and modifications, and to review the adequacy of system documentation.

3.1.3 Disaster Risk

The adequacy of support personnel and system resources for reporting P&L and positions must also be ensured in the event of a physical disaster. Examples of such disasters would include a power failure, fire, or explosion that closes down a trading facility and/or its supporting systems. Another example would be a computer system problem, such as a virus or error with consequences far-reaching enough to jeopardize the entire support structure (the most famous example is the Y2K crisis). Resource adequacy cannot be limited to just the ability to keep track of existing positions. It is also necessary to allow continued trading in a sufficiently controlled environment, at least at a level that will permit the ongoing management of existing positions.

The steps to deal with disaster risk begin with the development of a detailed contingency plan, which includes plans for backup computer systems, frequently updated backup data sets, backup power sources, and a backup trading floor. The adequacy of contingency plans must be judged against the likelihood that both the primary and backup facilities will be impacted by the same event. This concern was sharpened by the tragic events of September 11, 2001, when Bank of New York had both its primary and secondary trading systems, which were located in separate but nearby buildings, knocked out at the same time. This has caused many financial firms to rethink the degree of geographic separation that should be required between alternative sites.

Widespread computer errors that cut across all systems of the firm (backup as well as primary) are particularly worrisome. For example, the only way around the Y2K bug was to get a complete fix in place and thoroughly tested prior to the onset of the potential problem.

3.1.4 Personnel Risk

Investment banking firms have a history of raiding a competitor's personnel and hiring, en masse, an entire group of traders along with key support

staff. This can have the same impact on the raided firm as a physical disaster, but it has a longer recovery time, since replacement personnel must be identified, hired, and trained. Protective steps are to utilize cross-training and occasional backup duties as widely as possible to ensure that personnel are available to at least temporarily take over the duties of departed personnel. The requirements for thorough documentation of systems and procedures are also important.

3.2 LEGAL RISK

There are two types of legal risk: (1) the risk that contracts will prove unenforceable and (2) the risk that actions of the firm's employees will be found to be illegal, subjecting the firm to substantial penalties. We will examine both in turn.

3.2.1 The Risk of Unenforceable Contracts

The legal risk that the terms or conditions of a contract will prove unenforceable due to legal defects can prove a more serious problem than the credit risk that a counterparty does not have the financial capacity to perform on a contract. If a contract is found to be unenforceable, it may simultaneously impact a large number of contracts and have exactly the same impact on a trading firm as if a large number of counterparties defaulted simultaneously. A classic case of this was the finding by British courts that derivative contracts with British municipalities were *ultra vires*; that is, they were not contracts that the municipalities were legally authorized to enter into. This simultaneously canceled all outstanding derivatives contracts that financial firms had with British municipalities. For more detail, see Malcolm, Sharma, and Tanega (1999, 149–150). Another reason why legal risk can be more serious than credit risk is that it suffers more from adverse selection. Counterparty default is generally unrelated to whether the counterparty owes money or is owed money. However, lawsuits occur only when counterparties owe money.

The major mitigants to legal risk are:

- Thoroughly reviewing contract terms by experienced lawyers to ensure that language is properly drafted and that the contracted activities are authorized for the contracting parties.
- Thoroughly documenting what terms have been agreed to.
- Restricting dealings to reputable counterparties (know your customer).
- Placing limits on exposure to legal interpretations.

- Ensuring that contracts specify that legal jurisdiction resides with court systems that have experience in dealing with the particular issues involved and have previously demonstrated fairness in dealing with such cases.

A thorough review of contract terms may require lawyers with specialized legal knowledge of particular subject areas of law and legal jurisdiction (such as laws of particular countries, states, and districts), including knowledge of how courts and juries in a jurisdiction tend to interpret the law as well as applicable precedents. This often requires that legal work be contracted to outside counsel who specialize in certain areas and jurisdictions. However, care must be exercised to prevent front-office areas, which have a vested interest in seeing that a transaction gets done, from using this process to shop for a legal opinion, hiring a legal firm that can be counted on to provide a favorable opinion. The process of outside contracting of legal opinions must be controlled by an in-house legal department or a single trusted outside legal firm that can be counted on to offer independent judgments in the interest of the trading firm when this conflicts with the interest of individual front-office areas within the firm.

Adequate and clear legal language may prove useless if sufficient documentation has not been obtained showing customer agreement to the language. The most important measure in this regard is a strong commitment to following up verbal trade agreements with well-documented confirmations and signed legal agreements. This requires adequate documentation staff within trade support functions and the discipline to turn down potentially profitable business from counterparties that do not follow through on the required documentation. The enforcement of these rules is often placed within the credit risk function. Documentation should include written confirmation that a counterparty's board of directors and senior management have knowledge of the activities being contracted and have authorized the officers of the counterparty firm with which the trading firm is dealing to enter into such contracts on behalf of the counterparty firm. It is also useful to record all conversations between the counterparties and trading firm personnel so that disputes as to what terms were verbally agreed to can be settled equitably, without resorting to costly legal proceedings.

Firms have started to worry about what may be termed *legal-basis risk*. This arises when a firm treats transactions with two different customers as offsetting and hence without market risk (although not without credit risk). However, it may turn out that slightly different wording in the two contracts means that they are not truly offsetting in all circumstances. Although carefully vetting contractual language is a necessary countermeasure, an even better preventative is to use standardized contractual language as much as

possible to make it easier to spot differences. The International Swaps and Derivatives Association (ISDA) has been working to develop standardized language that can be used in derivatives contracts. See Section 13.1.1.2 for more details.

In addition to enforcing documentation rules, the credit risk function also needs to restrict the extension of credit to reputable counterparties. It is necessary to recognize that the willingness of a counterparty to meet contractual obligations is every bit as important as its financial ability to meet those obligations. A counterparty that does not have a good business reputation to protect may feel free to look for the slightest pretext to enter a legal challenge to meet its contractual obligations. Even if a firm has legal right strongly on its side, dealing with such a client may be very costly due to the expense of litigation and the threat of using a lawsuit as an excuse for a fishing expedition discovery process designed to uncover internal corporate information that can cause public embarrassment. The threat of such costs may incline a firm to settle for less than the full amount contractually owed, which serves as an incentive for unscrupulous firms to delay the settlement of legitimate claims. By contrast, a firm or an individual whose reputation for ethical business dealings is one of its assets will actually lean in the direction of making payments that meet its understanding of its obligations, even when the formal contract has been imperfectly drawn.

Because it is extremely difficult to quantify legal risk, firms may overlook the usefulness of quantitative limits to control exposure. Consider an example of a particular legal interpretation that has the potential to void all contracts of a specific type. The firm's legal consultants can issue opinions on the degree of likelihood that such an interpretation will be issued in the future by a court or regulatory body. Ultimately, business management must make a judgment on whether the economic benefits of the contract, relative to alternative ways of achieving the desired financial result, outweigh this risk. On a single deal, this is a binary decision—either you enter into the contract or you don't. There are few circumstances under which protection against an unfavorable contract interpretation can be purchased, making legal risk very different from market or credit risk. However, this is all the more reason to place a quantitative limit on the total size of contracts subject to all being voided by a single interpretation, where the size of the contract can be quantified by the potential loss from being voided. Quantitative limits place a control on risk, can be sized based on the degree of economic benefit relative to the perceived degree of legal uncertainty, and provide a framework for ensuring that individual deal approval is limited to those with the greatest potential benefit relative to potential loss.

One particular issue of legal risk that often causes concern is how bankruptcy courts will treat contractual obligations. When a counterparty goes

into default, the counterparty's reputation and desire to deal fairly no longer serve as a bulwark against litigation risk. In bankruptcy, all of the bankrupt firm's creditors become competitors in legal actions to gain as much of a share of the remaining assets as possible. Even when legal documents have been well drawn to provide specific collateral against an obligation or specific netting arrangements between derivative contracts on which the bankrupt firm owes and is owed money, other creditors may try to convince bankruptcy courts that it is only fair that they receive a share of the collateral or derivatives on which the bankrupt firm is owed money. Bankruptcy courts have been known to issue some very surprising rulings in these circumstances.

Contractual intention can be voided not only by courts, but also by regulatory authorities or legislatures, which may issue rules that make certain contractual provisions unenforceable. Financial institutions can and do mount lobbying campaigns against such changes, but other parties may be as effective or more effective in lobbying on the other side. Financial firms often need to analyze what they believe is the prospect for future regulatory actions in order to determine whether certain current business will prove to be worthwhile.

More detail on legal risk and how to control it can be found in Chapter 7 of Malcolm et al. (1999).

3.2.2 The Risk of Illegal Actions

The possibility of a firm's employees engaging in actions found to be illegal bears a very close relationship to reputational risk, which is examined in the next section. Any legal proceedings against a firm have the potential to damage the firm's reputation and the willingness of clients to engage its services. Even when legal proceedings don't result in a judgment against the firm, the publicity about the allegations and embarrassing disclosures in the legal discovery process can still impair reputation. And actions that can generate negative press, even if not rising to the level of illegality, can have a similar effect on reputation. One of the most effective screens for acceptable behavior remains the classic "Would you be comfortable seeing a description of this practice on the front page of the *Wall Street Journal*?"

The primary focus of legal and reputational risk has always been on the fiduciary responsibilities owed by a firm to its clients, particularly its less sophisticated clients. But recent cases have extended concern to damages that the client may inflict on others that the firm may be seen as having abetted. Section 4.3.2 on the losses in lawsuits of JPMorgan Chase and Citigroup for having been party to the Enron deception of investors is a good case study in this respect.

3.3 REPUTATIONAL RISK

Firms need to be sure not only that contract provisions are legally enforceable, but also that the process of enforcing their legal rights will not damage their business reputation. Even if a contract is strictly legal and enforceable, if its terms seem palpably unfair or can be portrayed as taking advantage of a client, the enforcement of the legal claims may be as damaging to the firm (or more so) as the inability to enforce the claims would have been. All transactions need to be reviewed by business managers from the viewpoint of whether the transaction is one that the client fully understands and it can reasonably be interpreted as a sensible action for the client to take. Ever since the Bankers Trust (BT) fiasco with Procter & Gamble (P&G) and Gibson Greetings, described in Section 4.3.1, all firms have placed increased emphasis on processes to ensure that transactions are appropriate or suitable for the client. The following processes are included:

- Conduct a careful review of all marketing materials to make sure that transactions have been fully explained, no misleading claims have been made, and no ambiguity exists as to whether the financial firm is simply acting as deal structurer or is also acting as an adviser to the client with fiduciary responsibility for the soundness of its advice. A full explanation of transactions may need to include simulations of possible outcomes, including stress situations.
- Make certain that any request from a client for a mark-to-market valuation of an existing transaction is supplied by support personnel using objective standards and not by marketing personnel who may have motivations to mislead the client as to the true performance of the transaction. Further, all valuations supplied need to be clearly labeled as to whether they are actual prices at which the trading firm is prepared to deal or simply indications of the general market level.
- Rank clients by their degree of financial sophistication and transactions by their degree of complexity, and ensure that a proper fit exists between the two. In cases where complex transactions are negotiated with less sophisticated clients, extra care needs to be taken to ensure that any advice given to the client by marketing personnel is consistent with their knowledge of the client's needs.
- Verify that clients fully understand the nature of the transactions they are undertaking, including written confirmation of such assurances from senior managers in some cases, based on the size and complexity of deals. These steps to ensure appropriateness and suitability are important not only to guard a trading firm's reputation, but, in extreme

cases, to also serve as protection against litigation. Note that the need to ensure the suitability of transactions to clients and the need to provide clients with evaluations that the trading firm can certify as reliable limit a trading firm's ability to simply serve as a credit intermediary between two counterparties using back-to-back derivatives.

3.4 ACCOUNTING RISK

Accounting risk can be viewed as a form of reputational risk. When a firm makes serious accounting errors, requiring the restatement of past earnings, it does not lead to any net loss of cash to the firm, as in cases of fraud, operations errors, or incorrectly drawn contracts. However, it can damage investor, creditor, and regulator confidence in the accuracy of information that the firm supplies about its financial health. This loss of confidence can be so severe that it threatens the firm's continued existence, as the Kidder Peabody financial disaster, discussed in Section 4.1.2, illustrates.

Measures to control accounting risk are similar in nature to those needed to control legal risk. Instead of needing knowledge of legal issues and precedents and how courts tend to interpret the law, knowledge of generally accepted accounting principles (GAAP) and how accounting boards of standards and regulatory authorities tend to interpret these principles is needed. The need for specialized knowledge by accounting jurisdiction is similar to the need for specialized knowledge by legal jurisdiction. The need to obtain independent accounting opinions and avoid opinion shopping parallels those considerations for legal risk. The need for thorough documentation showing that accounting rules are being followed parallels the need for thorough documentation of contractual understandings. The need for limits on exposure to accounting policies open to interpretation parallels the need for limits on exposure to legal interpretation.

3.5 FUNDING LIQUIDITY RISK

Funding liquidity risk should be clearly differentiated from the liquidity risk we discussed as part of market risk in Section 1.2, which is sometimes called *asset liquidity risk*.

Funding liquidity risk has two fundamental components:

1. The risk that investors' perception of the firm's credit quality will become impaired, thereby raising the firm's funding costs relative to the costs of competitors across all funding sources utilized.

2. The overly heavy use of a particular funding source in a given time period, raising the firm's funding cost relative to that of competitors for that particular funding source only.

Controlling the cost of the firm's liabilities by managing investors' perceptions of the firm's credit quality is the flip side of the coin of credit risk's management of the credit quality of the firm's assets. Crises in investor confidence are usually triggered by problems with earnings or the inadequacy of capital. As a result, they are functions of the overall management of the firm's business. The chief financial officer of the firm has particular responsibility for controlling funding liquidity risk by explaining the earnings situation to financial analysts and rating agencies and ensuring that capital levels are maintained to meet both regulatory guidelines and the expectations of financial analysts and rating agencies. Specific funding liquidity responsibilities of the treasury function of the firm include ensuring that any such crisis is not exacerbated by having to raise too much funding from the market at a time of crisis. Preferably, the firm should be able to reduce to a bare minimum its funding during a crisis period to gain time for the firm to improve its fundamental financial condition and tell its side of the story effectively to financial analysts, rating agencies, and individual investors.

The ability to avoid too much market funding in these circumstances requires:

- Long-term plans to get more funding from stable sources less sensitive to a firm's credit rating (such as retail deposits and transaction balances), to lengthen the maturity of market funding, to create cushions of market funding to tap in emergencies by raising less than the full amount of potential funds available, and to arrange backup lines of credit.
- Information systems to project periods of large funding needs in order to spread out the period of time over which such funding is raised. Of particular importance is the use of funding needs projections to avoid having funding requirements over a short period being so heavy that they trigger a crisis of investor confidence.
- Well-developed contingency plans for handling a funding crisis, which could include steps such as selling liquid assets, unwinding liquid derivatives positions that tie up collateral, and utilizing untapped cushions of funding and backup lines of credit.

The treasury function's management of particular funding sources to avoid overuse is also tied to information systems that can project future funding needs. It may be necessary to restrict particular types of investment

or derivative transactions that depend on access to particular funding sources to be profitable. For example, some transactions are profitable only if off-balance-sheet commercial paper funding can be obtained, bypassing the need for capital to be held against on-balance-sheet assets. However, the treasury function may need to limit the total amount of commercial paper being rolled over in any particular period to reduce the risk of having to pay a premium for such funding.

3.6 ENTERPRISE RISK

Enterprise risk can be tied to the fixed nature of many of the costs of engaging in a particular line of business. Even heavily personnel-intensive businesses, such as trading, still have fixed cost components such as buildings, computer and communications equipment, and some base level of employee compensation below which a firm loses its ability to remain in the business line through downturns in activity. However, these fixed costs entail the risk of losses to the extent that the amount of business that can be attracted in a downturn cannot cover the fixed costs.

By its nature, the management of enterprise risk belongs more naturally to individual business managers than to a corporate-wide risk function. Usually, the corporate-wide operational risk function will restrict itself to attempting to include some measurement of enterprise risk in the risk-adjusted return on capital (RAROC) or shareholder value added (SVA) measures.

3.7 IDENTIFICATION OF RISKS

In Damon Runyon's short story on which the musical *Guys and Dolls* is based, a gambler named Sky Masterson relates the following advice he received from his father:

Son, no matter how far you travel or how smart you get, always remember this: Someday, somewhere, a guy is going to come to you and show you a nice brand-new deck of cards on which the seal is never broken, and this guy is going to offer to bet you that the jack of spades will jump out of the deck and squirt cider in your ear. But, son, do not bet him, for as sure as you do you are going to get an ear full of cider.

The equivalent of this story for a risk manager is the trader or marketer who informs you that "There is absolutely no risk of loss on this product."

As my experience with markets has grown, I have come to recognize this assertion as a sure harbinger of painful losses to come, either sooner or later. However, my first encounter with the statement came well before I was involved with the financial side of banking, when I was working in Chase Manhattan's operations research department on projects like the simulation of the truck routes that delivered checks from branches to the head office and the sorting machines that then processed the checks.

One day on the subway, I ran into someone I had worked with on these simulations, but had not seen in a few years. He told me about the wonderful new job he had heading up a unit of the bank that matched firms that wanted to borrow securities with those that wanted to lend them. The bank received a nice fee for the service and he was aggressively growing the business. The key to profitability was operational efficiency, at which he was an expert. He told me that since the bank was not a principal to any of these transactions, there was absolutely no risk of loss on the product.

The losses came a few years later. When Drysdale Securities, a large borrower of government securities, could not repay its borrowings, it turned out that considerable ambiguity existed about whether the lenders of the securities understood they were being borrowed by Chase or by Drysdale with Chase merely arranging the borrowing. The legal contracts under which the transactions had been executed were open to the interpretation that Chase was the principal. Chase lost \$285 million in settling these claims (see Section 4.1.1 for a more detailed discussion). My acquaintance, needless to say, lost his job.

Before a risk can be controlled, it must first be recognized. Often, the management team that is involved with the introduction of a new product may lack the experience to perceive a possibility of risk and as a result may fail to call in the expertise needed to control the risk. For example, if a new legal risk is not recognized, the firm's legal experts may never thoroughly review the existing contracts. This is why it has become the accepted best practice in the financial industry to establish a new-product review process in which more experienced business managers and experts in risk disciplines (such as market risk, credit risk, reputational risk, legal, finance, and audit) vet proposals for products to make sure risks are identified and controls are instituted.

3.8 OPERATIONAL RISK CAPITAL

We started this chapter by stating that operational risks do not fall under the scope of financial risk management, since they could not be managed using liquid markets. This is not to say that quantitative measures cannot

be developed for operational risk, just that the tools to do so will come from the traditional insurance industry and will be close to the tools used to manage exposure to physical disasters, such as hurricanes and nuclear plant breakdowns. There are certain common items in the tool kit of financial risk and insurance risk, given that they are both trying to measure exposures in the extreme tail of events, as discussed in Section 1.3. For example, you will see use of simulation, extreme value theory, and stress scenarios in both. But the specific techniques discussed in this book, very closely tied to relating loss estimation and control to liquid market price movements, cannot be applied.

The primary impetus for developing quantitative measures of operational risk has been a desire to develop a methodology for operational risk capital to complement the measures of capital allocated for market risk and credit risk. In particular, the push by the Basel Committee on Banking Supervision to promulgate international standards requiring all banks to allocate capital against operational risk has spurred much work on how to quantify this capital requirement.

Operational risk capital can be approached in two ways—from the bottom up and from the top down. The bottom-up approach emphasizes quantitative measures of factors that contribute to operational risk. Some possibilities are:

- Audit scores as a measure of operations risk.
- Counts of unreconciled items or error rates as a measure of operations risk.
- Measures of delay in obtaining signed confirmations as a measure of legal risk.

Although these measures provide good incentives, tying reduction in capital to desirable improvements in controls, it is very difficult to establish links between these measures and the possible sizes of losses. Some firms are pursuing research on this, but supporting data is scarce.

The top-down approach emphasizes the historical volatility of earnings. This measure provides a direct link to the size of losses and can include all operational risks, even enterprise risk. But what incentive does this measure provide to reducing operational risk? No credit is given to a program that clears up back-office problems or places new controls on suitability.

Neither of these approaches bears much resemblance to the use of actual market prices for reduction of risk, which we discussed in Chapter 1 as the hallmark of financial risk management. To the extent market prices are available for some operational risks, it would come from the insurance

market, since it is possible to purchase insurance against some types of risk of fraud, operations errors, disasters, loss of personnel, legal liability, and accounting errors.

An up-to-date and thorough introduction to methodological approaches to quantifying operational risk capital and the regulatory background of the Basel Committee initiatives can be found in the closely related books Moosa (2007), particularly Chapters 5, 6, and 7, and Moosa (2008), particularly Chapters 4, 5, and 6.

Financial Disasters

One of the fundamental goals of financial risk management is to avoid the types of disasters that can threaten the viability of a firm. So we should expect that a study of such events that have occurred in the past will prove instructive. A complete catalog of all such incidents is beyond the scope of this book, but I have tried to include the most enlightening examples that relate to the operation of financial markets, as this is the book's primary focus.

A broad categorization of financial disasters involves a three-part division:

1. Cases in which the firm or its investors and lenders were seriously misled about the size and nature of the positions it had.
2. Cases in which the firm and its investors and lenders had reasonable knowledge of its positions, but had losses resulting from unexpectedly large market moves.
3. Cases in which losses did not result from positions held by the firm, but instead resulted from fiduciary or reputational exposure to positions held by the firm's customers.

4.1 DISASTERS DUE TO MISLEADING REPORTING

A striking feature of all the financial disasters we will study involving cases in which a firm or its investors and lenders have been misled about the size and nature of its positions is that they all involve a significant degree of deliberation on the part of some individuals to create or exploit incorrect information. This is not to say situations do not exist in which firms are misled without any deliberation on the part of any individual. Everyone who has been in the financial industry for some time knows of many instances when everyone at the firm was misled about the nature of positions because a ticket was entered into a system incorrectly. Most typically,

this will represent a purchase entered as a sale, or vice versa. However, although the size of such errors and the time it takes to detect them can sometimes lead to substantial losses, I am not aware of any such incident that has resulted in losses that were large enough to threaten the viability of a firm.

An error in legal interpretation can also seriously mislead a firm about its positions without any deliberate exploitation of the situation. However, such cases, although they can result in large losses, tend to be spread across many firms rather than concentrated at a single firm, perhaps because lawyers tend to check potentially controversial legal opinions with one another. The best-known case of this type was when derivatives contracted by British municipalities were voided. See Section 3.2.

If we accept that all cases of financial disaster due to firms being misled about their positions involve some degree of complicity on the part of some individuals, we cannot regard them completely as cases of incorrectly reported positions. Some of the individuals involved know the correct positions, at least approximately, whereas others are thoroughly misinformed. Understanding such cases therefore requires examining two different questions:

1. Why does the first group persist in taking large positions they know can lead to large losses for the firm despite their knowledge of the positions?
2. How do they succeed in keeping this knowledge from the second group, who we can presume would put a stop to the position taking if they were fully informed?

I will suggest that the answer to the first question tends to be fairly uniform across disasters, while the answer to the second question varies.

The willingness to take large risky positions is driven by moral hazard. As we saw in our discussion of moral hazard in Section 2.1, it represents an asymmetry in reward structure and an asymmetry in information; in other words, the group with the best information on the nature of the risk of a position has a greater participation in potential upside than potential downside. This often leads insiders to desire large risky positions that offer them commensurately large potential gains. The idea is that traders own an option on their profits; therefore, they will gain from increasing volatility, as we discussed in Section 2.1. The normal counterweights against this are the attempts by representatives of senior management, stockholders, creditors, and government regulators, who all own a larger share of the potential downside than the traders, to place controls on the amount of risk taken. However, when those who could exercise this control substantially

lack knowledge of the positions, the temptation exists for traders to exploit the control weakness to run inflated positions. This action often leads to another motivation spurring the growth of risky positions—the Ponzi scheme, as discussed in Section 2.2.

Some traders who take risky positions that are unauthorized but disguised by a control weakness will make profits on these positions. These positions are then possibly closed down without anyone being the wiser. However, some unauthorized positions will lead to losses, and traders will be strongly tempted to take on even larger, riskier positions in an attempt to cover up unauthorized losses. This is where the Ponzi scheme comes in. I think it helps to explain how losses from unauthorized positions can grow to be so overwhelmingly large. Stigum (1989) quotes an “astute trader” with regard to the losses in the Chase/Drysdale financial disaster: “I find it puzzling that Drysdale could lose so much so fast. If you charged me to lose one-fourth of a billion, I think it would be hard to do; I would probably end up making money some of the time because I would buy something going down and it would go up. They must have been extraordinarily good at losing money.” I would suggest that the reason traders whose positions are unauthorized can be so “extraordinarily good at losing money” is that normal constraints that force them to justify positions to outsiders are lacking and small unauthorized losses already put them at risk of their jobs and reputations. With no significant downside left, truly reckless positions are undertaken in an attempt to make enough money to cover the previous losses. This is closely related to double-or-nothing betting strategies, which can start with very small stakes and quickly mushroom to extraordinary levels in an effort to get back to even.

This snowballing pattern can be seen in many financial disasters. Nick Leeson’s losses on behalf of Barings were just \$21 million in 1993, \$185 million in 1994, and \$619 million in just the first two months of 1995 (Chew 1996, Table 10.2). John Rusnak’s unauthorized trading at Allied Irish Bank (AIB) accumulated losses of \$90 million in its first five years through 1999, \$210 million in 2000, and \$374 million in 2001 (Ludwig 2002, Section H). Joseph Jett’s phantom trades at Kidder Peabody started off small and ended with booked trades in excess of the quantity of all bonds the U.S. Treasury had issued.

The key to preventing financial disasters based on misrepresented positions is therefore the ability to spot unauthorized position taking in a timely enough fashion to prevent this explosive growth in position size. The lessons we can learn from these cases primarily center on why it took so long for knowledge of the misreported positions to spread from an insider group to the firm’s management. We will examine each case by providing a brief summary of how the unauthorized position arose, how it failed to come

to management's attention, and what lessons can be learned. In each instance, I provide references for those seeking more detailed knowledge of the case. General conclusions based on the cases in this section can be found in Section 3.1.1.

4.1.1 Chase Manhattan Bank/Drysdale Securities

4.1.1.1 Incident In three months of 1976, Drysdale Government Securities, a newly founded subsidiary of an established firm, succeeded in obtaining unsecured borrowing of about \$300 million by exploiting a flaw in the market practices for computing the value of U.S. government bond collateral. This unsecured borrowing exceeded any amount Drysdale would have been approved for, given that the firm had only \$20 million in capital. Drysdale used the borrowed money to take outright positions in bond markets. When the traders lost money on the positions they put on, they lacked cash with which to pay back their borrowings. Drysdale went bankrupt, losing virtually all of the \$300 million in unsecured borrowings. Chase Manhattan absorbed almost all of these losses because it had brokered most of Drysdale's securities borrowings. Although Chase employees believed they were only acting as agents on these transactions and were not taking any direct risk on behalf of Chase, the legal documentation of the securities borrowings did not support their claim.

4.1.1.2 Result Chase's financial viability was not threatened by losses of this size, but the losses were large enough to severely damage its reputation and stock valuation for several years.

4.1.1.3 How the Unauthorized Positions Arose Misrepresentation in obtaining loans is unfortunately not that uncommon in bank lending. A classic example would be Anthony De Angelis, the "Salad Oil King," who, in 1963, obtained \$175 million in loans supposedly secured by large salad oil holdings, which turned out to be vast drums filled with water with a thin layer of salad oil floating on top. Lending officers who came to check on their collateral were bamboozled into only looking at a sample from the top of each tank.

The following are some reasons for featuring the Drysdale shenanigans in this section rather than discussing any number of other cases of misrepresentation:

- Drysdale utilized a weakness in trading markets to obtain its funds.
- Drysdale lost the borrowed money in the financial markets.
- It is highly unusual for a single firm to bear this large a proportion of this large a borrowing sting.

There is not much question as to how Drysdale managed to obtain the unsecured funds. The firm took systematic advantage of a computational shortcut in determining the value of borrowed securities. To save time and effort, borrowed securities were routinely valued as collateral without accounting for accrued coupon interest. By seeking to borrow large amounts of securities with high coupons and a short time left until the next coupon date, Drysdale could take maximum advantage of the difference in the amount of cash the borrowed security could be sold for (which included accrued interest) and the amount of cash collateral that needed to be posted against the borrowed security (which did not include accrued interest).

4.1.1.4 How the Unauthorized Positions Failed to Be Detected Chase Manhattan allowed such a sizable position to be built up largely because it believed that the firm's capital was not at risk. The relatively inexperienced managers running the securities borrowing and lending operation were convinced they were simply acting as intermediaries between Drysdale and a large group of bond lenders. Through their inexperience, they failed both to realize that the wording in the borrowing agreements would most likely be found by a court to indicate that Chase was taking full responsibility for payments due against the securities borrowings and to realize the need for experienced legal counsel to review the contracts.

4.1.1.5 How the Unauthorized Positions Were Eventually Detected There was some limit to the size of bond positions Drysdale could borrow, even given the assumption that the borrowings were fully collateralized. At some point, the size of the losses exceeded the amount of unauthorized borrowings Drysdale could raise and the firm had to declare bankruptcy.

4.1.1.6 Lessons Learned The securities industry as a whole learned that it needed to make its methods for computing collateral value on bond borrowings more precise. Chase, and other firms that may have had similar control deficiencies, learned the need for a process that forced areas contemplating new product offerings to receive prior approval from representatives of the principal risk control functions within the firm (see Section 3.7).

4.1.1.7 Further Reading Chapter 14 of Stigum (1989) gives a detailed description of the Chase/Drysdale incident, some prior misadventures in bond borrowing collateralization, and the subsequent market reforms.

4.1.2 Kidder Peabody

4.1.2.1 Incident Between 1992 and 1994, Joseph Jett, head of the government bond trading desk at Kidder Peabody, entered into a series of trades

that were incorrectly reported in the firm's accounting system, artificially inflating reported profits. When this was ultimately corrected in April 1994, \$350 million in previously reported gains had to be reversed.

4.1.2.2 Result Although Jett's trades had not resulted in any actual loss of cash for Kidder, the announcement of such a massive misreporting of earnings triggered a substantial loss of confidence in the competence of the firm's management by customers and General Electric, which owned Kidder. In October 1994, General Electric sold Kidder to PaineWebber, which dismantled the firm.

4.1.2.3 How the Unauthorized Positions Arose A flaw in accounting for forward transactions in the computer system for government bond trading failed to take into account the present valuing of the forward. This enabled a trader purchasing a cash bond and delivering it at a forward price to book an instant profit. Over the period between booking and delivery, the profit would inevitably dissipate, since the cash position had a financing cost that was unmatched by any financing gain on the forward position.

Had the computer system been used as it was originally intended (for a handful of forward trades with only a few days to forward delivery), the size of error would have been small. However, the system permitted entry not only of contracted forward trades, but also of intended forward delivery of bonds to the U.S. Treasury, which did not actually need to be acted on, but could be rolled forward into further intentions to deliver in the future. Both the size of the forward positions and the length of the forward delivery period were constantly increased to magnify the accounting error. This permitted a classic Ponzi scheme of ever-mounting hypothetical profits covering the fact that previously promised profits never materialized.

Although it has never been completely clear how thoroughly Jett understood the full mechanics of the illusion, he had certainly worked out the link between his entry of forward trades and the recording of profit, and increasingly exploited the opportunity.

4.1.2.4 How the Unauthorized Positions Failed to Be Detected Suspicions regarding the source of Jett's extraordinary profit performance were widespread throughout the episode. It was broadly perceived that no plausible account was being offered of a successful trading strategy that would explain the size of reported earnings. On several occasions, accusations were made that spelled out exactly the mechanism behind the inflated reporting. Jett seemed to have had a talent for developing explanations that succeeded in totally confusing everyone (including, perhaps, himself) as to what was going on. However, he was clearly aided and abetted by a management satisfied

enough not to take too close a look at what seemed like a magical source of profits.

4.1.2.5 How the Unauthorized Positions Were Eventually Detected Large increases in the size of his reported positions and earnings eventually triggered a more thorough investigation of Jett's operation.

4.1.2.6 Lessons to Be Learned Two lessons can be drawn from this: Always investigate a stream of large unexpected profits thoroughly and make sure you completely understand the source. Periodically review models and systems to see if changes in the way they are being used require changes in simplifying assumptions (see Section 8.2.8).

4.1.2.7 Further Reading Jett has written a detailed account of the whole affair (see Jett 1999). However, his talent for obscurity remains and it is not possible to tell from his account just what he believes generated either his large profits or the subsequent losses. For an account of the mechanics of the deception, one must rely on the investigation conducted by Gary Lynch on behalf of Kidder. Summaries of this investigation can be found in Hansell (1997), Mayer (1995), and Weiss (1994).

4.1.3 Barings Bank

4.1.3.1 Incident The incident involved the loss of roughly \$1.25 billion due to the unauthorized trading activities during 1993 to 1995 of a single, relatively junior trader named Nick Leeson.

4.1.3.2 Result The size of the losses relative to Barings Bank's capital along with potential additional losses on outstanding trades forced Barings into bankruptcy in February 1995.

4.1.3.3 How the Unauthorized Positions Arose Leeson, who was supposed to be running a low-risk, limited return arbitrage business for Barings in Singapore, was actually taking increasingly large speculative positions in Japanese stocks and interest rate futures and options. He disguised his speculative position taking by reporting that he was taking the positions on behalf of fictitious customers. By booking the losses to these nonexistent customer accounts, he was able to manufacture fairly substantial reported profits for his own accounts, enabling him to earn a \$720,000 bonus in 1994.

4.1.3.4 How the Unauthorized Positions Failed to Be Detected A certain amount of credit must be given to Leeson's industriousness in perpetrating a

deliberate fraud. He worked hard at creating false accounts and was able to exploit his knowledge of weaknesses in the firm's controls. However, anyone reading an account of the incident will have to give primary credit to the stupendous incompetence on the part of Barings' management, which ignored every known control rule and failed to act on myriad obvious indications of something being wrong. What is particularly amazing is that all those trades were carried out in exchange-traded markets that require immediate cash settlement of all positions, thereby severely limiting the ability to hide positions (although Leeson did even manage to get some false reporting past the futures exchange to reduce the amount of cash required).

The most blatant of management failures was an attempt to save money by allowing Leeson to function as head of trading and the back office at an isolated branch. Even when auditors' reports warned about the danger of allowing Leeson to settle his own trades, thereby depriving the firm of an independent check on his activities, Barings' management persisted in their folly. Equally damning was management's failure to inquire how a low-risk trading strategy was supposedly generating such a large profit. Even when covering these supposed customer losses on the exchanges required Barings to send massive amounts of cash to the Singapore branch, no inquiries were launched as to the cause. A large part of this failure can be attributed to the very poor structuring of management information so that different risk control areas could be looking at reports that did not tie together. The funding area would see a report indicating that cash was required to cover losses of a customer, not the firm, thereby avoiding alarm bells about the trading losses. A logical consequence is that credit exposure to customers must be large since the supposed covering of customer losses would entail a loan from Barings to the customer. However, information provided to the credit risk area was not integrated with information provided to funding and showed no such credit extension.

4.1.3.5 How the Unauthorized Positions Were Eventually Detected The size of losses Leeson was trying to cover up eventually got too overwhelming and he took flight, leaving behind an admission of irregularities.

4.1.3.6 Lessons to Be Learned One might be tempted to say that the primary lesson is that there are limits to how incompetent you can be and still hope to manage a major financial firm. However, to try to take away something positive, the major lessons would be the absolute necessity of an independent trading back office, the need to make thorough inquiries about unexpected sources of profit (or loss), and the need to make thorough inquiries about any large unanticipated movement of cash.

4.1.3.7 Further Reading A concise and excellent summary of the Barings case constitutes Chapter 10 of Chew (1996). Chapter 11 of Mayer (1997) contains less insight on the causes, but is strong on the financial and political maneuvers required to avoid serious damage to the financial system from the Barings failure. Leeson has written a full-length book that appears to be reasonably honest as to how he evaded detection (Leeson 1996). Fay (1996) and Rawnsley (1995) are also full-length accounts.

4.1.4 Allied Irish Bank (AIB)

4.1.4.1 Incident John Rusnak, a currency option trader in charge of a very small trading book in AIB's Allfirst First Maryland Bancorp subsidiary, entered into massive unauthorized trades during the period 1997 through 2002, ultimately resulting in \$691 million in losses.

4.1.4.2 Result This resulted in a major blow to AIB's reputation and stock price.

4.1.4.3 How the Unauthorized Positions Arose Rusnak was supposed to be running a small arbitrage between foreign exchange (FX) options and FX spot and forward markets. He was actually running large outright positions and disguising them from management.

4.1.4.4 How the Unauthorized Positions Failed to Be Detected To quote the investigating report, "Mr. Rusnak was unusually clever and devious." He invented imaginary trades that offset his real trades, making his trading positions appear small. He persuaded back-office personnel not to check these bogus trades. He obtained cash to cover his losses by selling deep-in-the-money options, which provided cash up front in exchange for a high probability of needing to pay out even more cash at a later date, and covered up his position by offsetting these real trades with further imaginary trades. He entered false positions into the firm's system for calculating value at risk (VaR) to mislead managers about the size of his positions.

In many ways, Rusnak's pattern of behavior was a close copy of Nick Leeson's at Barings, using similar imaginary transactions to cover up real ones. Rusnak operated without Leeson's advantage of running his own back office, but had the offsetting advantage that he was operating in an over-the-counter market in which there was not an immediate need to put up cash against losses. He also was extremely modest in the amount of false profit he claimed so he did not set off the warning flags of large unexplained profits from small operations that Leeson and Jett at Kidder Peabody had triggered in their desire to collect large bonuses.

Like Barings, AIB's management and risk control units demonstrated a fairly startling level of incompetence in failing to figure out that something was amiss. AIB at least has the excuse that Rusnak's business continued to look small and insignificant, so it never drew much management attention. However, the scope and length of time over which Rusnak's deception continued provided ample opportunity for even the most minimal level of controls to catch up with him.

The most egregious was the back office's failure to confirm all trades. Rusnak succeeded in convincing back-office personnel that not all of these trades needed to be confirmed. He relied partly on an argument that trades whose initial payments offset one another didn't really need to be checked since they did not give rise to net immediate cash flow, ignoring the fact that the purported trades had different terms and hence significant impact on future cash flows. He relied partly on booking imaginary trades with counterparties in the Asian time zone, making confirmation for U.S.-based back-office staff a potentially unpleasant task involving middle-of-the-night phone calls, perhaps making it easier to persuade them that this work was not really necessary. He also relied on arguments that costs should be cut by weakening or eliminating key controls.

Once this outside control was missing, the way was opened for the ongoing manipulation of trading records. Auditors could have caught this, but the spot audits performed used far too small a sample. Suspicious movements in cash balances, daily trading profit and loss (P&L), sizes of gross positions, and levels of daily turnover were all ignored by Rusnak's managers through a combination of inexperience in FX options and overreliance on trust in Rusnak's supposedly excellent character as a substitute for vigilant supervision. His management was too willing to withhold information from control functions and too compliant with Rusnak's bullying of operations personnel as part of a general culture of hostility toward control staff. This is precisely the sort of front-office pressure that reduces support staff independence, which was referred to in Section 3.1.1.

4.1.4.5 How the Unauthorized Positions Were Eventually Detected In December 2001, a back-office supervisor noticed trade tickets that did not have confirmations attached. When informed that the back-office personnel did not believe all trades required confirmations, he insisted that confirmation be sought for existing unconfirmed trades. Although it took some time for the instructions to be carried out, when they finally were carried out in early February 2002, despite some efforts by Rusnak to forge written confirmations and bully the back office into not seeking verbal confirmations, his fraud was brought to light within a few days.

4.1.4.6 Lessons to Be Learned This incident does not provide many new lessons beyond the lessons that should already have been learned from Barings. This case does emphasize the need to avoid engaging in small ventures in which the firm lacks any depth of expertise—there is simply too much reliance on the knowledge and probity of a single individual.

On the positive side, the investigative report on this fraud has provided risk control units throughout the financial industry with a set of delicious quotes that are sure to be trotted out anytime they feel threatened by cost-cutting measures or front-office bullying and lack of cooperation. The following are a few choice samples from Ludwig (2002):

- When one risk control analyst questioned why a risk measurement system was taking market inputs from a front-office-controlled system rather than from an independent source, she was told that AIB “would not pay for a \$10,000 data feed from Reuters to the back office.”
- When questioned about confirmations, “Mr. Rusnak became angry. He said he was making money for the bank, and that if the back office continued to question everything he did, they would drive him to quit. . . . Mr. Rusnak’s supervisor warned that if Mr. Rusnak left the bank, the loss of his profitable trading would force job cuts in the back office.”
- “When required, Mr. Rusnak was able to use a strong personality to bully those who questioned him, particularly in Operations.” His supervisors “tolerated numerous instances of severe friction between Mr. Rusnak and the back-office staff.”
- Rusnak’s supervisor “discouraged outside control groups from gaining access to information in his area and reflexively supported Mr. Rusnak whenever questions about his trading arose.”
- “[I]n response to general efforts to reduce expense and increase revenues, the Allfirst treasurer permitted the weakening or elimination of key controls for which he was responsible. . . . Mr. Rusnak was able to manipulate this concern for additional cost cutting into his fraud.”

4.1.4.7 Further Reading I have relied heavily on the very thorough report issued by Ludwig (2002).

4.1.5 Union Bank of Switzerland (UBS)

4.1.5.1 Incident This incident involves losses of between \$400 million and \$700 million in equity derivatives during 1997, which appear to have been exacerbated by lack of internal controls. A loss of \$700 million during 1998 was due to a large position in Long-Term Capital Management (LTCM).

4.1.5.2 Result The 1997 losses forced UBS into a merger on unfavorable terms with Swiss Bank Corporation (SBC) at the end of 1997. The 1998 losses came after that merger.

4.1.5.3 Were the Positions Unauthorized? Less is known about the UBS disaster than the other incidents discussed in this chapter. Even the size of the losses has never been fully disclosed. Considerable controversy exists about whether the 1997 losses just reflected poor decision making or unlucky outcomes or whether an improper control structure led to positions that management would not have authorized. The 1998 losses were the result of a position that certainly had been approved by the UBS management, but evidence suggests that it failed to receive adequate scrutiny from the firm's risk controllers and that it was not adequately disclosed to the SBC management that took over the firm.

What seems uncontroversial is that the equity derivatives business was being run without the degree of management oversight that would be normally expected in a firm of the size and sophistication of UBS, but there is disagreement about how much this situation contributed to the losses. The equity derivatives department was given an unusual degree of independence within the firm with little oversight by, or sharing of information with, the corporate risk managers. The person with senior risk management authority for the department doubled as head of quantitative analytics. As head of analytics, he was both a contributor to the business decisions he was responsible for reviewing and had his compensation tied to trading results, which are both violations of the fundamental principles of independent oversight.

The equity derivative losses appear to have been primarily due to four factors:

1. A change in British tax laws, which impacted the value of some long-dated stock options.
2. A large position in Japanese bank warrants, which was inadequately hedged against a significant drop in the underlying stocks (see the fuller description in Section 11.4).
3. An overly aggressive valuation of long-dated options on equity baskets, utilizing correlation assumptions that were out of line with those used by competitors.
4. Losses on other long-dated basket options, which may have been due to modeling deficiencies.

The first two transactions were ones where UBS had similar positions to many of its competitors so it would be difficult to accuse the firm of excessive risk taking, although its Japanese warrant positions appear to have

been unreasonably large relative to competitors. The last two problems appear to have been more unique to UBS. Many competitors made accusations that its prices for trades were off the market.

The losses related to LTCM came as the result of a position personally approved by Mathis Cabiallavetta, the UBS CEO, so they were certainly authorized in one sense. However, accusations have been made that the trades were approved without adequate review by risk control areas and were never properly represented in the firm's risk management systems. Although about 40 percent of the exposure represented a direct investment in LTCM that had large potential profits to weigh against the risk, about 60 percent of the exposure was an option written on the value of LTCM shares. However, there was no effective way in which such an option could be risk managed given the illiquidity of LTCM shares and restrictions that LTCM placed on UBS delta hedging the position (see the next-to-last paragraph in Section 11.1).

The imbalance in risk/reward trade-off for an option that was that difficult to risk manage had caused other investment banks to reject the proposed trade. UBS appears to have entered into the option because of its desire for a direct investment in LTCM, which LTCM tied to agreement to the option. Agreeing to this type of bundled transaction can certainly be a legitimate business decision, but it is unclear whether the full risk of the option had been analyzed by UBS or whether stress tests of the two positions taken together had been performed.

4.1.5.4 Lessons Learned This incident emphasizes the need for independent risk oversight.

4.1.5.5 Further Reading The fullest account of the equity derivative losses is contained in a book by Schutz (2000), which contains many lurid accusations about improper dealings between the equity derivatives department and senior management of the firm. Schutz has been accused of inaccuracy in some of these charges—see *Derivative Strategies* (1998) for details. There is also a good summary in the January 31, 1998, issue of the *Economist*.

A good account of the LTCM transaction is Shirreff (1998). Lowenstein (2000), an account of the LTCM collapse, also covers the UBS story in some detail.

4.1.6 Société Générale

4.1.6.1 Incident In January 2008, Société Générale reported trading losses of \$7.1 billion that the firm attributed to unauthorized activity by a junior trader, Jérôme Kerviel.

4.1.6.2 Result The large loss severely damaged Société Générale's reputation and required it to raise a large amount of new capital.

4.1.6.3 How the Unauthorized Positions Arose In this section and the next, I am drawing primarily on the Société Générale Special Committee of the Board of Directors Report to Shareholders of May 22, 2008 (I'll abbreviate references to it as SpecComm) and its accompanying Mission Green Report of the Société Générale General Inspection Department (I'll abbreviate it as MG).

Kerviel took very large unauthorized positions in equities and exchange-traded futures, beginning in July 2005 and ending when his concealment of positions was uncovered in January 2008. His primary method for concealing these unauthorized positions was to enter fictitious transactions that offset the risk and P&L of his true trades. The fictitious nature of these transactions was hidden mostly by creating transactions with forward start dates and then, relying on his knowledge of when control personnel would seek confirmation of a forward-dated trade, canceling the trade prior to the date that confirmation would be sought (Kerviel had previously worked in the middle office of the firm, which may have provided him with particular insight into the actions of control personnel). Not surprisingly, given his need to constantly replace canceled fictitious transactions with new ones, there were a large number of these trades, 947 transactions according to MG Focus 4. How was Kerviel able to enter this many fictitious trades before discovery of his fraud?

4.1.6.4 How the Unauthorized Positions Failed to Be Detected

Trade Cancellation There was no procedure in place that required control functions to confirm information entered for a trade that was then canceled and Kerviel knew this, nor was there a system in place for red-flagging an unusual level of trade cancellations. SpecComm, point 10, notes that the back and middle office gave "priority to the correct execution of trades" and showed "an absence of an adequate degree of sensitivity to fraud risks." The head of equity derivatives at a European bank is quoted as saying, "If he was able to cancel a trade and book a new one before the confirm was sent out, the clock [for obtaining confirmation] would start again. But at our bank, we actively monitor cancel-and-correct activity for each trader, which is standard practice at most institutions. It would stick out like a sore thumb if you had one trader who was perpetually cancelling and correcting trades" (Davies 2008). Hugo Banziger, chief risk officer of Deutsche Bank, is quoted as saying, "Routine IT controls can monitor unusual trades put on and cancelled—this is a particularly effective control mechanism" (Davies 2008). It

certainly appears from the account in MG that no such procedures were in place at Société Générale, and even the inquiry to confirm the counterparty on a canceled trade that eventually led to Kerviel's downfall in January 2008 appears to have been a matter of chance (MG Focus 6).

Supervision Kerviel's immediate manager resigned in January 2007. For two and a half months, until the manager was replaced, Kerviel's positions were validated by his desk's senior trader. Day-to-day supervision of Kerviel by the new manager, who started in April 2007, was weak (SpecComm, point 9; MG, page 6). While Kerviel had begun his fraudulent activities prior to January 2007, the size of his unauthorized positions increased explosively at this time (MG Focus 10).

Trading Assistant The trading assistant who worked with Kerviel in entering trades, who would have the most immediate potential access to seeing how he was manipulating the trading system, may have been operating in collusion with Kerviel. This has not been confirmed (MG, page 3, notes that this is an allegation under investigation by the courts), but, in any case, the trading assistant appears to have accepted Kerviel's directions without questioning. This would have helped Kerviel's credibility with control functions, since the trading assistant reported to a control function and was the primary point of contact of other control functions regarding Kerviel's positions (MG, page 4).

Vacation Policy The normal precaution of forcing a trader to take two consecutive weeks of vacation in a year, during which time his positions would be managed by another trader, was not followed (MG, page 7). This control could easily have caused the collapse of a scheme based on constant rolling forward of fraudulent trading entries.

Gross Positions There were no limits or other monitoring of Kerviel's gross positions, only his net positions (SpecComm, point 10, notes the "lack of certain controls liable to identify the fraudulent mechanisms, such as the control of the positions' nominal value"). Had gross positions been monitored, this would have revealed the abnormally large size of his activities and might have raised suspicions as to what the purpose was of such large positions. Henning Giescke, chief risk officer of the UniCredit Group, is quoted as saying, "In high-volume businesses, banks have to look at gross as well as net position. This allows an institution to look at each trader's book to see whether they are taking too much risk, regardless of whether the net position is neutral" (Davies 2008). The chief risk officer of a UK bank is quoted as saying, "To effectively manage basis risk, you have to be able to

see how the outright position—the notional—performs against the hedge. It is inconceivable such a sophisticated institution could have missed this. Modern systems are able to stress-test positions, and to do this you automatically need the notional amount” (Davies 2008). Kerviel’s unusually high amount of brokerage commissions (MG, page 6), related to his high level of gross positions, could also have provided a warning sign.

Cash and Collateral The use of fictitious transactions to conceal positions will often create positions of unusual size in cash and required collateral—since the fictitious trades do not generate any cash or collateral movements, there is nothing to balance out the cash and collateral needs of the real trades. This provides good opportunities for fraud detection. The reason that Société Générale’s control functions did not respond to these clues is that cash and collateral reports and inquiry procedures lacked sufficient granularity to detect unusual movements at the level of a single trader (MG Focus 13).

P&L Concealment of trading positions will not always lead to unusual earnings patterns. A trader who is trying to conceal losses may be satisfied simply to show a small positive P&L. But some fraudulent traders will show unusual profits, either because their unauthorized positions have resulted in large gains for which they want to be compensated or because their success in hiding losing positions encourages them to also claim some phantom gains to fund bonuses. Kerviel was reporting trading gains in excess of levels his authorized position taking could have accounted for, and this should have given his management and the control functions a warning sign to investigate closely the source of his earnings (MG Focus 12). These warning signs were apparently not pursued.

4.1.6.5 How the Unauthorized Positions Were Eventually Detected One of Kerviel’s fictitious trades was identified as fabricated by control personnel as part of routine monitoring of positions, leading to a thorough investigation. Kerviel’s attempts to deflect the inquiry by forging confirmations proved fruitless. It appears that it was just a matter of chance that this particular inquiry led to identification of the fraud.

4.1.6.6 Lessons to Be Learned What new lessons can we draw from this control failure? From one point of view, the answer is not much—Kerviel’s methods for eluding scrutiny of his positions were very close to those used in previous incidents such as those of Kidder Peabody, Barings, and Allied Irish Bank. But, from another viewpoint, we can learn quite a bit, since clear patterns are emerging when we look across episodes.

The obvious lessons for control personnel are to tighten procedures that may lead to detection of fictitious trade entries. Corresponding to the points raised in Section 4.1.6.4, the specific lessons follow.

Trade Cancellation Institute systems for monitoring patterns of trade cancellation. Flag any trader who appears to be using an unusually high number of such cancellations. Any trader flagged should have a reasonably large sample of the cancellations checked to make sure that they represent real trades by checking details of the transaction with the counterparty.

Supervision Control personnel should be aware of situations in which traders are being supervised by temporary or new managers. Tightened control procedures should be employed.

Trading Assistant Control personnel must remember that even in situations where there is no suspicion of dishonesty, trading assistants are often under intense pressure from the traders with whom they work closely. Their job performance ratings and future career paths often depend on the trader, regardless of official reporting lines. The greater prestige, experience, and possible bullying tactics of a trader can often convince a trading assistant to see things from the trader's viewpoint. Other control personnel must be cognizant of these realities and not place too much reliance on the presumed independence of the trading assistant.

Vacation Policy Rules for mandatory time away from work should be enforced.

Gross Positions Gross positions must be monitored and highlighted in control reports. This is particularly important since unusually high ratios of gross to net positions are a warning sign of potentially inadequately measured basis risk as well as a possible flag for unauthorized activities. The Kidder Peabody and Allied Irish Bank frauds could also have been uncovered by investigating unusually high ratios of gross to net trading.

Cash and Collateral Cash and collateral requirements should be monitored down to the individual trader level. Better monitoring of cash and credit flows would have also been instrumental in uncovering the Barings and Allied Irish Bank frauds.

P&L Any patterns of P&L that are unusual relative to expectations need to be identified and investigated by both management and the control functions. Identification of unusual patterns can be comparisons to historical

experience, to budgeted targets, and to the performance of traders with similar levels of authority. Investigation of suspicious earnings patterns could also have led to earlier discovery of the Kidder Peabody and Barings frauds.

4.1.6.7 Further Reading I have relied primarily on the Société Générale Special Committee of the Board of Directors Report to Shareholders (2008) and its accompanying Société Générale Mission Green Report (2008).

4.1.7 Other Cases

Other disasters involving unauthorized positions are covered more briefly, because they had less of an impact on the firm involved, because it is harder to uncover details on what occurred, or because they do not have any lessons to teach beyond those of the cases already discussed:

- Toshihida Iguchi of Daiwa Bank's New York office lost \$1.1 billion trading Treasury bonds between 1984 and 1995. He hid his losses and made his operation appear to be quite profitable by forging trading slips, which enabled him to sell without authorization bonds held in customer accounts to produce funds he could claim were part of his trading profit. His fraud was aided by a situation similar to Nick Leeson's at Barings—Iguchi was head of both trading and the back-office support function. In addition to the losses, Daiwa lost its license to trade in the United States, but this was primarily due to its failure to promptly disclose the fraud once senior executives of the firm learned of it. A more detailed account of this by Rob Jameson of ERisk can be found on their website, www.erisk.com.
- The Sumitomo Corporation of Japan lost \$2.6 billion in a failed attempt by Yasuo Hamanaka, a senior trader, to corner the world's copper market—that is, to drive up prices by controlling a large portion of the available supply. Sumitomo management claimed that Hamanaka had employed fraudulent means in hiding the size of his positions from them. Hamanaka claimed that he had disclosed the positions to senior management. Hamanaka was sent to jail for his actions. The available details are sketchy, but some can be found in Dwyer (1996), *Asiaweek* (1996), Kooi (1996), and McKay (1999).
- Askin Capital Management and Granite Capital, hedge funds that invested in mortgage securities, went bankrupt in 1994 with losses of \$600 million. It was revealed that David Askin, the manager of the funds, was valuing positions with his own marks substituted for dealer quotes and using these position values in reports to investors in the

funds and in marketing materials to attract new clients. For a brief discussion, see Mayer (1997).

- Merrill Lynch reportedly lost \$350 million in trading mortgage securities in 1987, due to risk reporting that used a 13-year duration for all securities created from a pool of 30-year mortgages. Although this duration is roughly correct for an undivided pool of 30-year mortgages, the correct duration is 30 years when the interest-only (IO) part is sold and the principal-only (PO) part is kept, as Merrill was doing. See Crouhy, Galai, and Mark (2001).
- National Westminster Bank in 1997 reported a loss on interest rate caps and swaptions of about \$140 million. The losses were attributed to trades dating back to 1994 and had been masked by deliberate use by traders of incorrect volatility inputs for less liquid maturities. The loss of confidence in management caused by this incident may have contributed to NatWest's sale to the Royal Bank of Scotland. I have heard from market sources that the traders were taking advantage of the middle-office saving costs by checking only a sample of volatility marks against market sources, although it is unclear how the traders were able to determine in advance which quotes would be checked. A more detailed account is Wolfe (2001).
- The large Swiss bank UBS in 2011 reported a loss of \$2.3 billion due to unauthorized trading by Kweku Adoboli, a relatively junior equity trader. This incident cost the CEO of UBS his job. Adoboli's ability to enter into unauthorized trades appears to have been engineered by means very similar to those of Kerviel in the Société Générale incident discussed in Section 4.1.6. He took advantage of intimate knowledge of back-office control procedures to identify a loophole. Trades with forward settlement greater than 15 days were not being immediately confirmed with counterparties; confirmation was delayed until closer to the settlement date. If the trade was canceled prior to the date on which the confirmation would have been confirmed, no confirmation ever took place. Adoboli appears to have been able to utilize this loophole to disguise his real positions by entering bogus offsetting forward positions and then canceling the fictitious positions prior to the date they would have been confirmed, replacing them with new fictitious forwards. For this to have gone on for any period of time, there must have also been flaws in UBS's monitoring of excessive cancellations. Due to an ongoing criminal prosecution against Adoboli at the time of my writing, not many public details are available. Wilson (2011) is a good summary of what is known about the mechanics of the unauthorized trades, and Broom (2011) summarizes the devastating impact the revelation of this faulty control environment had on UBS.

4.2 DISASTERS DUE TO LARGE MARKET MOVES

We will now look at financial disasters that were not caused by incorrect position information, but were caused by unanticipated market moves. The first question that should be asked is: How is a disaster possible if positions are known? No matter what strategy is chosen, as losses start to mount beyond acceptable bounds, why aren't the positions closed out? The answer is lack of liquidity. We will focus on this aspect of these disasters.

4.2.1 Long-Term Capital Management (LTCM)

The case we will consider at greatest length is that of the large hedge fund managed by Long-Term Capital Management (LTCM), which came close to bankruptcy in 1998. In many ways, it represents an ideal example for this type of case since all of its positions were marked to a market value daily, the market values were supplied by the dealers on the other end of each trade, no accusations have been made of anyone at LTCM providing misleading information about positions taken, and the near failure came in the midst of some of the largest market moves in recent memory.

To review the facts, LTCM had been formed in 1994 by about a dozen partners. Many of these partners had previously worked together at Salomon Brothers in a highly successful proprietary trading group. Over the period from 1994 until early 1998, the LTCM fund produced quite spectacular returns for its investors. From the beginning, the partners made clear that they would be highly secretive about the particulars of their investment portfolio, even by the standard of other hedge funds. (Since hedge funds are open only to wealthy investors and cannot be publicly offered the way mutual funds are, they are not subject to legal requirements to disclose their holdings.)

Within the firm, however, the management style favored sharing information openly, and essentially every investment decision was made by all the partners acting together, an approach that virtually eliminates the possibility of a rogue trader making decisions based on information concealed from other members of the firm. Although it is true that outside investors in the fund did not have access to much information beyond the month-end valuation of its assets and the track record of its performance, it is equally true that the investors knew these rules prior to their decision to invest. Since the partners who managed the fund were such strong believers in the fund that they had invested most of their net worth in it (several even borrowed to invest more than their net worth), their incentives were closely aligned with investors (in other words, there was little room for moral hazard). If anything, the concentration of partner assets in the fund should

have led to more risk-averse decision making than might have been optimal for outside investors, who invested only a small portion of their wealth in the fund, with the exception of UBS, discussed in Section 4.1.5.

In fact, even if investors had been given access to more information, there is little they could have done with it, since they were locked into their investments for extended time periods (generally, three years). This reflected the basic investment philosophy of LTCM, which was to locate trading opportunities that represented what the partners believed were temporary disruptions in price relationships due to short-term market pressures, which were almost certain to be reversed over longer time periods. To take advantage of such opportunities, they needed to know they had access to patient capital that would not be withdrawn if markets seemed to be temporarily going against them. This also helped to explain why LTCM was so secretive about its holdings. These were not quick in-and-out trades, but long-term holdings, and they needed to prevent other firms from learning the positions and trading against them.

The following are two examples of the types of positions the LTCM fund was taking:

1. LTCM was long U.S. interest rate swaps and short U.S. government bonds at a time when these spreads were at historically high levels. Over the life of the trade, this position will make money as long as the average spread between the London Interbank Offered Rate (LIBOR) at which swaps are reset (see Section 10.1.6) and the repurchase agreement (RP) rates at which government bonds are funded (see Section 10.1.2) is not higher than the spread at which the trade was entered into. Over longer time periods, the range for the average of LIBOR-RP spreads is not that wide, but in the short run, swap spreads can show large swings based on relative investor demand for the safety of governments versus the higher yield of corporate bonds (with corporate bond issuers then demanding interest rate swaps to convert fixed debt to floating debt).
2. LTCM sold equity options at historically high implied volatilities. Over the life of the trade, this position will make money if the actual volatility is lower than the implied volatility, but in the short run, investor demand for protection against stock market crashes can raise implied volatilities to very high levels. Perold (1999a) presents further analysis of why LTCM viewed these trades as excellent long-term investments and presents several other examples of positions it entered into.

One additional element was needed to obtain the potential returns LTCM was looking for. LTCM needed to be able to finance positions for longer terms in order to be able to ensure there was no pressure on them to

sell positions before they reached the price relationships LTCM was waiting for. However, the banks and investment banks who financed hedge fund positions were the very competitors that they least wanted to share information on holdings with. How were they to persuade firms to take credit risk without knowing much about the trading positions of the hedge fund?

To understand why the lenders were comfortable in doing this, we need to digress a moment into how credit works in a futures exchange. A futures exchange (see Section 14.2) represents the extreme of being willing to lend without knowledge of the borrower. Someone who purchases, for example, a bond for future delivery needs to deposit only a small percentage of the agreed purchase price as margin and does not need to disclose anything about one's financial condition. The futures exchange is counting on the nature of the transaction itself to provide confidence that money will not be lost in the transaction. This is because anytime the value of the bond falls, the purchaser is required to immediately provide added margin to fully cover the decline in value. If the purchaser does not do so, the position is closed out without delay. Loss is possible only if the price has declined so much since the last time the price fell and margin was added that the incremental price drop exceeds the amount of initial margin or if closing out the option results in a large price move. The probability of this occurring is kept low by setting initial margins high enough, restricting the size of position that can be taken by any one investor, and designing futures contracts to cover sufficiently standardized products to ensure enough liquidity that the closing out of a trade will not cause a big price jump.

LTCM wanted to deal in over-the-counter markets as well as on futures exchanges partly because it wanted to deal in some contracts more individually tailored than those available on futures exchanges and partly because of the position size restrictions of exchanges. However, the mechanism used to assure lenders in over-the-counter markets is similar—there is a requirement to cover declines in market value by immediately putting up cash. If a firm fails to put up the cash, then positions are closed out. LTCM almost always negotiated terms that avoided posting the initial margin. Lenders were satisfied with the lack of initial margin based on the size of the LTCM fund's equity, the track record of its excellent returns, and the firm's recognized investment management skills. Lenders retained the option of demanding initial margin if fund equity fell too much.

This dependence on short-term swings in valuation represented a potential Achilles' heel for LTCM's long-term focused investment strategy. Because the firm was seeking opportunities where market pressures were causing deviation from long-run relationships, a strong possibility always existed that these same market pressures would push the deviation even further. LTCM would then immediately need to come up with cash to fund the change in

market valuation. This would not be a problem if some of the trades were moving in its favor at the same time as others were moving against it, since LTCM would receive cash on upswings in value to balance having to put up cash on downswings (again, the same structure as exchange-traded futures). However, if many of its trades were to move against it in tandem, LTCM would need to raise cash quickly, either from investors or by cutting positions.

In the actual events of August and September 1998, this is exactly what led to LTCM's rapid downfall. The initial trigger was a combination of the Russian debt default of August, which unsettled the markets, and the June 1998 decision by Salomon Brothers to liquidate proprietary positions it was holding, which were similar to many of those held by LTCM. The LTCM fund's equity began to decline precipitously from \$4.1 billion as of the end of July 1998, and it was very reluctant to cut positions in a turbulent market in which any large position sale could easily move the valuations even further against it. This left the option of seeking new equity from investors. LTCM pursued this path vigorously, but the very act of doing so created two perverse effects. First, rumors of LTCM's predicament caused competitors to drive market prices even further against what they guessed were LTCM's positions, in anticipation of LTCM being forced to unload the positions at distressed prices. Second, to persuade potential investors to provide new money in the midst of volatile markets, LTCM was forced to disclose information about the actual positions it held. As competitors learned more about the actual positions, their pressure on market prices in the direction unfavorable to LTCM intensified.

As market valuations continued to move against LTCM and the lack of liquidity made it even more unlikely that reducing positions would be a viable plan, it became increasingly probable that in the absence of a truly large infusion of new equity, the LTCM fund would be bankrupt. Its creditors started to prepare to close out LTCM's positions, but quickly came to fear that they were so large and the markets were so illiquid that the creditors would suffer serious losses in the course of doing so. The lenders were also concerned that the impact of closing out these positions would depress values in the already fragile markets and thereby cause considerable damage to other positions held by the creditors and other investment firms they were financing.

Ultimately, 14 of the largest creditors, all major investment banks or commercial banks with large investment banking operations, contributed a fresh \$3.65 billion in equity investment into the LTCM fund to permit the firm to keep operating and allow for a substantial time period in which to close out positions. In return, the creditors received substantial control over fund management. The existing investors had their investments valued at

the then current market value of \$400 million, so they had only a 10 percent share in the positions of the fund. Although some of the partners remained employed to help wind down investments, it was the consortium of 14 creditors who now exercised control and insisted on winding down all positions.

As a result, the markets calmed down. By 2000, the fund had been wound down with the 14 creditors having recovered all of the equity they had invested and having avoided any losses on the LTCM positions they had held at the time of the bailout. This outcome lends support to two propositions: LTCM was largely right about the long-term values underlying its positions, and the creditors were right to see the primary problem as one of liquidity, which required patience to ride out.

Please note that the bailout was not primarily a rescue of LTCM's investors or management, but a rescue of LTCM's creditors by a concerted action of these creditors. Even recently, I continue to encounter the view that the bailout involved the use of U.S. government funds, helped the LTCM investors and management avoid the consequences of their mistakes, and therefore contributed to an attitude that some firms are "too big to fail" and so can afford to take extra risks because they can count on the government absorbing some of their losses.

I do not think evidence is available to support any of these claims. Interested readers can form their own conclusions by looking at the detailed account of the negotiations on the rescue package in Lowenstein (2000). An opposing viewpoint can be found in Shirreff (2000). The only government involvement was some coordination by the Federal Reserve, acting out of legitimate concern for the potential impact on the financial markets. The LTCM creditors took a risk by investing money in the fund, but did so in their own self-interest, believing (correctly, as it turns out) that they were thereby lowering their total risk of loss. LTCM's investors and managers had little left to lose at the point of the bailout since they could not lose more than their initial investment. It is true that, without a rescue, the fund would have been liquidated, which would have almost certainly wiped out the remaining \$400 million market value of the investors. However, in exchange for this rescue, they were able to retain only a 10 percent interest in the fund's positions, since the \$3.65 billion in new investment was explicitly not being used to enable new trades, but only to wind down the existing positions.

LTCM management was certainly aware of the potential for short-term market movements to disrupt the fund's fundamental trading strategy of focusing on longer-term relationships. The firm tried to limit this risk by insisting that its positions pass value at risk (VaR) tests based on whether potential losses over one month due to adverse market moves would reduce equity to unacceptable levels. Where LTCM seems to have fallen short of

best practices was a failure to supplement VaR measures with a full set of stress test scenarios (see Section 11.2). It did run stress versions of VaR based on a higher than historical level of correlations, but it is doubtful that this offers the same degree of conservatism as a set of fully worked-through scenarios.

A lesson that all market participants have learned from the LTCM incident is that a stress scenario is needed to look at the impact of a competitor holding similar positions exiting the market, as when Salomon decided to cut back on proprietary trading. However, even by best practice standards of the time, LTCM should have constructed a stress test based on common economic factors that could cause impacts across its positions, such as a flight to quality by investors, which would widen all credit spreads, including swap spreads, and increase premiums on buying protection against stock market crashes, hence increasing option volatility.

Another point on which LTCM's risk management could be criticized is a failure to account for the illiquidity of its largest positions in its VaR or stress runs. LTCM knew that the position valuations it was receiving from dealers did not reflect the concentration of some of LTCM's positions, either because dealers were not taking liquidity into account in their marks or because each dealer knew only a small part of LTCM's total size in its largest positions.

Two other criticisms have been made of LTCM's management of risk with which I disagree. One is that a simple computation of leverage would have shown that LTCM's positions were too risky. However, as will be seen in Section 13.2.4, leverage by itself is not an adequate measure of risk of default. It must be multiplied by volatility of the firm's assets. But this just gets us back to testing through VaR or stress scenarios. The second criticism is that LTCM showed unreasonable faith in the outcome of models. I see no evidence to support this claim. Major positions LTCM entered into—U.S. swap spreads to narrow, equity volatilities to decline—were ones that many proprietary position takers had entered into. For example, the bias in equity implied volatilities due to demand for downside protection by shareholders had long been widely recognized as a fairly certain profit opportunity for investors with long-enough time horizons. That some firms made more use of models to inform their trading judgments while others relied more on trader experience tells me nothing about the relative quality of their decision making.

Most of the focus of LTCM studies has been on the decision making of LTCM management and the losses of the investors. I believe this emphasis is misplaced. It is a fairly common occurrence, and to be expected, that investment funds will have severe drops in valuation. The bankruptcy of an investment fund does not ordinarily threaten the stability of the

financial system the way the bankruptcy of a firm that makes markets or is a critical part of the payments system would. It just represents the losses of a small number of investors. Nor is there a major difference in consequences between bankruptcy and a large loss short of bankruptcy for an investment fund. It shouldn't matter to investors whether a fund in which they have invested \$10 million goes bankrupt or a fund in which they have invested \$30 million loses a third of its value. By contrast, losses short of bankruptcy hurt only the stockholders of a bank, whereas bankruptcy of a bank could hurt depositors and lead to loss of confidence in the banking system.

The reason that an LTCM failure came close to disrupting the financial markets and required a major rescue operation was its potential impact on the creditors to LTCM, so we need to take a closer look at their role in the story. In retrospect, the creditors to LTCM believed they had been too lax in their credit standards, and the incident triggered a major industry study of credit practices relating to trading counterparties (Counterparty Risk Management Policy Group 1999).

Some suggestions for improved practices, many of which are extensively addressed in this study, have been:

- **A greater reluctance to allow trading without initial margin for counterparties whose principal business is investing and trading.** A counterparty that has other substantial business lines—for example, auto manufacturing or retail banking—is unlikely to have all of its economic resources threatened by a large move in financial markets. However, a firm that is primarily engaged in these markets is vulnerable to illiquidity spreading from one market to another as firms close out positions in one market to meet margin calls in another market. For such firms, initial margin is needed as a cushion against market volatility.
- **Factoring the potential costs of liquidating positions in an adverse market environment into estimates of the price at which trades can be unwound.** These estimates should be based on the size of positions as well as the general liquidity of the market (see Section 6.1.2). These potential liquidation costs should impact estimates of the amount of credit being extended and requirements for initial margin.
- **A push for greater disclosure by counterparties of their trading strategies and positions.** Reliance on historical records of return as an indicator of the volatility of a portfolio can be very misleading because it cannot capture the impact of changes in trading style (see Section 7.1). Increased allowance for liquidation costs of positions will be very inexact if the creditor only knows the positions that a counterparty holds with the creditor without knowing the impact of other positions held.

To try to deal with counterparties' legitimate fears that disclosure of their positions will lead to taking advantage of this knowledge, creditors are implementing more stringent internal policies against the sharing of information between the firm's credit officers and the firm's traders.

- **Better use of stress tests in assessing credit risk.** To some extent, this involves using more extreme stresses than were previously used in measuring risk to reflect the increased market volatility that has been experienced in recent years. However, a major emphasis is also on more integration of market risk and credit risk stress testing to take into account overlap in risks. In the LTCM case, this would have required recognition by a creditor to LTCM that many of the largest positions being held by LTCM were also being held by other investment funds to which the firm had counterparty credit exposure, as well as by the firm's own proprietary traders. A full stress test would then look at the losses that would be incurred by a large market move and subsequent decrease in liquidity across all of these similar positions.

A complete account of the LTCM case covering all aspects of the history of the fund and its managers, the involvement of creditors, and the negotiations over its rescue can be found in Lowenstein (2000). The Harvard Business School case studies of Perold (1999a, 1999b) provide a detailed but concise analysis of the fund's investment strategy and the dilemma that it faced in August 1998.

4.2.2 Metallgesellschaft (MG)

The disaster at Metallgesellschaft (MG) reveals another aspect of liquidity management. In 1992, an American subsidiary of MG, Metallgesellschaft Refining and Marketing (MGRM), began a program of entering into long-term contracts to supply customers with gas and oil products at fixed costs and to hedge these contracts with short-term gas and oil futures. Although some controversy exists about how effective this hedging strategy was from a P&L standpoint, as we'll discuss in just a moment, the fundamental consequence of this strategy for liquidity management is certain. The futures being used to hedge were exchange-traded instruments requiring daily cash settlement, as explained in Section 10.1.4. The long-term contracts with customers involved no such cash settlement. So no matter how effective the hedging strategy was, the consequence of a large downward move in gas and oil prices would be to require MGRM to pay cash against its futures positions that would be offset by money owed to MGRM by customers who would be paid in the future.

A properly designed hedge will reflect both the cash paid and the financing cost of that cash during the period until the customer payment is due and hence will be effective from a P&L standpoint. However, the funding must still be obtained, which can lead to funding liquidity risk (see Section 3.5). As we will discuss in Section 6.1.6, such cash needs must be planned in advance. Limits need to be set on positions based on the amount of cash shortfall that can be funded.

It appears that MGRM did not communicate to its parent company the possible need for such funding. In 1993, when a large decrease in gas and oil prices had resulted in funding needs of around \$900 million, the MG parent responded by closing down the futures positions, leaving unhedged exposure to gas and oil price increases through the customer contracts. Faced with this open exposure, MG negotiated unwinds of these contracts at unfavorable terms. It may be that MG, with lack of advance warning as to possible cash needs, responded to the demand for cash as a sign that the trading strategy was deeply flawed; if only Barings' management had reacted similarly.

As mentioned earlier, the MG incident spurred considerable debate as to whether MGRM's trading strategy was reasonable or fundamentally flawed. Most notably, Culp and Miller (1995a) wrote an article defending the reasonableness of the strategy, and Mello and Parsons (1995) wrote an article attacking the Culp and Miller conclusions, which were then defended by Culp and Miller (1995b). Although it is difficult to settle the factual arguments about the particular events in the MG case, I believe the following lessons can be drawn:

- It is often a key component of a market maker's business strategy to extend available liquidity in a market (see Section 10.2.2). This requires the use of shorter-term hedges against longer-term contracts. Experience shows that this can be successfully carried out when proper risk controls are applied.
- The uncertainty of roll cost is a key risk for strategies involving shorter-term hedges against longer-term risk. As explained in Section 10.2.2, this requires the use of valuation reserves based on conservative assumptions of future roll cost. MGRM does not appear to have utilized valuation reserves; it just based its valuation on the historical averages of roll costs.
- A firm running short-term hedges against longer-term risk requires the flexibility to choose the shorter-term hedge that offers the best trade-off between risk and reward. It may legitimately choose to follow a hedging strategy other than a theoretical minimum variance hedge, or choose not to hedge with the longest future available, based on liquidity

considerations, or take into account the expectation of positive roll cost as part of potential return. It is not reasonable to conclude, as Mello and Parsons (1995) do, that these choices indicate that the firm is engaged in pure speculation rather than hedging. At the same time, regardless of a firm's conclusions about probable return, its assessment of risk should include valuation reserves, as in the previous point, and volume limits based on reasonable stress testing of assumptions.

4.3 DISASTERS DUE TO THE CONDUCT OF CUSTOMER BUSINESS

In this section, we focus on disasters that did not involve any direct financial loss to the firm, but were completely a matter of reputational risk due to the conduct of customer business.

4.3.1 Bankers Trust (BT)

The classic case of this type is the Bankers Trust (BT) incident of 1994, when BT was sued by Procter & Gamble (P&G) and Gibson Greetings. Both P&G and Gibson claimed that they had suffered large losses in derivatives trades they had entered into with BT due to being misled by BT as to the nature of the positions. These were trades on which BT had little market or credit risk, since it had hedged the market risk on them with other derivatives and there was no credit issue of P&G or Gibson being unable to pay the amount they owed. However, the evidence uncovered in the course of legal discovery for these lawsuits was severely damaging to BT's reputation for fair business dealing, led to the resignation of the firm's CEO, and ultimately had such negative consequences for the bank's ability to do business that it was forced into an acquisition by Deutsche Bank, which essentially amounted to a dismemberment of the firm.

The exact terms of these derivative trades were quite complex and are not essential to understanding the incident. Interested readers are referred to Chew (1996, Chapter 2) for details. The key point is that the trades offered P&G and Gibson a reasonably probable but small reduction in funding expenses in exchange for a potentially large loss under some less probable circumstances. P&G and Gibson had been entering into such trades for several years prior to 1994 with good results. The derivatives were not tailored to any particular needs of P&G or Gibson in the sense that the circumstances under which the derivatives would lose them money were not designed to coincide with cases in which other P&G or Gibson positions would be making money. Their objective was just to reduce expected funding costs. Since

the only way to reduce costs in some cases is to raise them in others, P&G and Gibson can be presumed to have understood that they could lose money under some economic circumstances. On what basis could they claim that BT had misled them?

One element that established some prima facie suspicion of BT was the sheer complexity of the structures. It was hard to believe that BT's clients started out with any particular belief about whether there was a small enough probability of loss in such a structure to be comfortable entering into it. BT would have had to carefully explain all the intricacies of the payoffs to the clients for them to be fully informed.

Since it was quite clear that the exact nature of the structures hadn't been tailored to meet client needs, why had BT utilized so complex a design? The most probable reason was that the structures were designed to be complex enough to make it difficult for clients to comparison shop the pricing to competitor firms. However, this also made the clients highly dependent on BT on an ongoing basis. If they wanted to unwind the position, they couldn't count on getting a competitive quote from another firm.

BT claimed that it had adequately explained all the payoffs and risks to P&G and Gibson. But then came the discovery phase of the lawsuit. BT, like all trading firms, recorded all phone lines of traders and marketers as a means of resolving disputes about verbal contracts (see Sections 3.1.1 and 3.2). However, this recording also picked up internal conversations among BT personnel. When subpoenaed, they produced evidence of BT staff boasting of how thoroughly they had fooled the clients as to the true value of the trades and how little the clients understood the true risks. Further, the internal BT recordings showed that price quotes to P&G and Gibson were being manipulated to mislead them. At first, they were given valuations of the trades that were much too high to mask the degree of profit BT was able to book up front. Later, they were given valuations that were too low because this was BT's bid at which to buy back the trade or swap it into a new trade offering even more profit to BT. For more details on what was revealed in the recordings, see Holland and Himmelstein (1995).

The BT scandal caused all financial firms to tighten up their procedures for dealing with customers, both in better controls on matching the degree of complexity of trades to the degree of financial sophistication of customers and in providing for customers to obtain price quotes from an area independent of the front office. These measures were detailed in Section 3.3.

Another lesson that you would think would be learned is to be cautious about how you use any form of communication that can later be made public. BT's reputation was certainly hurt by the objective facts about its conduct, but it was even further damaged by the arrogant and insulting tone some of its employees used in referring to clients, which could be

documented through recorded conversations. However, even with such an instructive example, we have seen Merrill Lynch's reputation being damaged in 2002 by remarks its stock analysts made in e-mails and tape-recorded conversations (see the article "Value of Trust" in the June 6, 2002, *Economist*) along with a number of similar incidents surrounding Wall Street's relations with Enron (see the article "Banks on Trial" in the July 25, 2002, *Economist*).

4.3.2 JPMorgan, Citigroup, and Enron

Following the Bankers Trust incident, investment banks put in controls to guard against exploitation of customers. But it was not seen as part of a bank's responsibility to safeguard others from actions by the customer. This has changed as part of the fallout from Enron's 2001 bankruptcy. As part of the process leading up to the bankruptcy, it was revealed that Enron had for years been engaging in dubious accounting practices to hide the size of its borrowings from investors and lenders (it was their part in these shenanigans that brought an end to the major accounting firm Arthur Andersen). One of the ploys that Enron had used was to disguise a borrowing as an oil futures contract.

As a major player in the energy markets, it was to be expected that Enron would be heavily engaged in futures contracts on oil. But these particular futures contracts did not involve taking any position on oil price movements. Enron sold oil for future delivery, getting cash, and then agreed to buy back the oil that it delivered for a fixed price. So, in effect, no oil was ever delivered. When you canceled out the oil part of the trades, what was left was just an agreement for Enron to pay cash later for cash it had received up front—in practice, if not in legal terms, a loan. The advantage to Enron was that it did not have to report this in its public statements as a loan, making the firm appear more desirable as an investment and as a borrower.

When this was finally disclosed, JPMorgan Chase and Citigroup, Enron's principal counterparties on these trades, justified their activities by saying that they had not harmed Enron, their client, in any way, and that they had no part in determining how Enron had accounted for the transactions on its books; that was an issue between Enron and Arthur Andersen. JPMorgan and Citigroup had treated these transaction as loans in their own accounting and reporting to regulators, so they had not deceived their own investors or lenders.

But both JPMorgan and Citigroup clearly knew what Enron's intent was in entering into the transaction. In the end, they agreed to pay a combined \$286 million for "helping to commit a fraud" on Enron's shareholders. They

also agreed to put new controls in place to ascertain that their clients were accounting for derivative transactions with them in ways that were transparent to investors.

The precedent of this successful legal action caused other investment banks to commit to similar new controls. And yet we have recently witnessed charges against Goldman Sachs for helping Greece hide its level of indebtedness from its European Union partners by disguising debt as an interest rate swap, a mechanism very similar to that in the Enron case. The details here are that the swap was deliberately done at an off-market rate, creating an up-front payment to Greece that would of course need to be paid back by Greece, with suitable interest, over the course of the swap's life. The only reason for creating the swap at an off-market rate would appear to be letting Greece take out a loan that didn't need to show up on its books.

Details on the Enron case can be found in McLean and Elkind (2003, 159–160, 407–408). Details on the Greek case can be found in Dunbar and Martinuzzi (2012).

4.3.3 Other Cases

The following are some examples of other cases in which firms damaged their reputations by the manner in which they dealt with customers:

- Prudential-Bache Securities was found to have seriously misled thousands of customers concerning the risk of proposed investments in limited partnerships. In addition to damaging its reputation, Prudential-Bache had to pay more than \$1 billion in fines and settlements. An account of this incident can be found in Eichenwald (1995).
- In 1995, a fund manager at Morgan Grenfell Asset Management directed mutual fund investments into highly speculative stocks, utilizing shell companies to evade legal restrictions on the percentage of a firm's stock that could be owned by a single fund. In addition to damage to its reputation, Morgan Grenfell had to pay roughly \$600 million to compensate investors for resulting losses. A brief case account can be found in Garfield (1998).
- JPMorgan's reputation was damaged by allegations that it misled a group of South Korean corporate investors as to the risk in derivative trades that lost hundreds of millions of dollars based on the precipitous decline in the Thai baht exchange rate against the dollar in 1997. An account of these trades and the ensuing lawsuits can be found in Gillen, Lee, and Austin (1999).
- Many investment banks had their reputations damaged in the events leading up to the large fall in value of technology stocks in 2001 and

2002. Evidence showed that some widely followed stock market analysts working at investment banks had issued favorable recommendations for companies as a quid pro quo for underwriting business, with analyst bonuses tied to underwriting business generated. Regulators responded with fines for firms, bans from the industry for some analysts, and requirements for separation of the stock analysis function from the underwriting business. A summary account with references can be found in Lowenstein (2004, 212–213).

Reputational risk incidents that arose in connection with the 2007–2008 crisis are covered in Sections 5.2.1, 5.2.2, and 5.2.3.

The Systemic Disaster of 2007–2008

5.1 OVERVIEW

There can be little question that the global financial disaster of 2007–2008 stemmed fundamentally from events in the market for collateralized debt obligations (CDOs) backed by subprime mortgages. Firms that failed or needed government rescue either had large losses in these CDOs or else got caught up in events triggered by the difficulties of firms that did have large losses on these CDOs. In examining the crisis, this chapter therefore begins with a section (5.2) focusing on CDOs backed by subprime mortgages. Section 5.3 looks at how this crisis then spread from the institutions with heavy losses in the CDO market to other institutions—by contagion through credit exposure and by contagion through impact on markets. Then, Sections 5.4 and 5.5 examine lessons from the crisis for, respectively, risk managers and government regulators. Section 5.6 takes a brief look at lessons from the crisis that go beyond the scope of risk managers and government regulators.

Just to attempt to clear up one confusing bit of nomenclature at the beginning—CDOs on subprime mortgages were termed asset-backed securities (ABSs), so what were called ABS CDOs were actually CDO-squared products (see Section 13.4.2 for explanation of a CDO-squared). In fact, as documented in Cordell, Huang, and Williams (2012), a very substantial portion of subprime mortgage securities were CDO-squared products. But since the economic and analytic characteristics of CDO-squared products do not differ materially from primary CDO products, as discussed in Section 13.4.2, I will ignore this distinction in the remainder of this chapter.

CDOs were the genesis of this crisis, and they were also at the root of what made it so damaging to the world economy. Large losses at banks due to lending in boom times that later goes sour under more challenging economic circumstances are a part of a fairly predictable cycle. Despite these periodic large losses, lending tends to be a profitable business over time, a conclusion that the studies on the excess of credit spreads over long-term

loss rates would tend to support (see Hull 2012, Section 23.5, and Amato and Remolona 2003). Nor were losses on mortgage lending over this crisis period confined to holders of CDOs; large banks like Countrywide and Washington Mutual and government-sponsored agencies like Fannie Mae and Freddie Mac managed to be big losers without much participation in CDOs. See the Financial Crisis Inquiry Report (2011, 106–109, 248–250, 305–307) on Countrywide and Washington Mutual; see Acharya et al. (2011) on the government-sponsored agencies.

What was different about credit losses that resulted from CDOs rather than from loans? The illusion that CDOs were bringing more liquidity to the mortgage lending market resulted in an exacerbation of what might otherwise have been a far more manageable downturn. As we'll see in Section 5.2, treating the CDOs as if they were liquid securities rather than illiquid loans helped to fuel an expansion in lending far beyond what probably would have occurred without it. Then, when it became clear that the alleged liquidity wasn't really there, the commitment of the investment banks to accounting for CDOs as if they were liquid assets turned what would have been longer-term losses to be dealt with over the length of a credit cycle into immediate requirements for raising new capital. This quickly led to contagion in which markets, securities, and firms not originally involved in the CDO market got heavily impacted as well. We'll follow this aspect of the story in Section 5.3.

The focus in this chapter is on those aspects of the crisis that are most directly relevant to risk management. For those looking for a broader view of the crisis, the Financial Crisis Inquiry Report (2011), referenced as FCIR (2011), and the guide to the literature in Lo (2012) are good starting points.

In my narrative and analysis and my lessons for risk managers section, I acknowledge that I cannot draw upon the firsthand detailed familiarity that would come from working in risk management at one of the affected institutions during the crisis period—I retired from JPMorgan Chase in 2004 and was working during the crisis period primarily as an educator. I have based my account on a combination of what is in the public record, what I have gleaned from conversations with people who were on the inside during the crisis, and what I have seen as an independent consultant to some of the impacted firms in the aftermath of the crisis. Balancing this, my lack of participation in the crisis leaves me relatively free of any axes to grind, positions to defend, or constraints due to confidentiality (though I can't identify, either explicitly or by implication, specific clients I worked for after the crisis).

In my analysis of lessons for regulators, I summarize the major proposals that have been offered, but restrict my own suggestions to those where my experience and judgment as a risk manager offer some direct benefit.

This means that I need to leave to others analysis of critical issues where my knowledge is less germane. To take one typical example, a good deal of policy discussion following the crisis has to do with issues surrounding how narrow you want to make the role of commercial banks—proposals like the Volcker rule banning proprietary trading by any firm with implicit government guarantees or suggestions about reimposition of Glass-Steagall-like restrictions on the mixing of commercial banking and investment banking. These proposals involve trade-offs between reducing the probability of future disasters versus the possible negative impacts on a country's economic growth by reducing financial innovation or by hurting the lending capacity of commercial banks by limiting their sources of revenue. As a risk manager, I have been trained in analyzing risks that arise within a given institutional structure and not in evaluating the economic impact of different institutional structures. On that which I cannot speak with insight, I will remain silent.

5.2 THE CRISIS IN CDOs OF SUBPRIME MORTGAGES

It is not surprising that a disaster of the magnitude of the 2007–2008 CDO crisis had many causes and has led to a sizable literature of exposition. While I draw on many books and articles in what follows, I would like to particularly draw the reader's attention to four relatively short articles that I find especially incisive in their analysis: Davidson (2007), Ashcroft and Schuermann (2008), Hull (2009), and Brunnermeier (2009). For those interested in further reading about the causes of the crisis and implications for the future, Lo (2012) is a concise guide to the best of the academic and journalistic literature on the topic, while Oyama (2010, Chapter 3) provides an excellent summary of reports that have been issued by various regulatory agencies and industry groups. Also recommended is the FCIR (2011) compiled by the Financial Crisis Inquiry Commission that was authorized by the U.S. Congress.

In trying to look at all the causes, I have divided up the narrative into separate sections on the institutions with different roles in the CDO process. Section 5.2.1 covers the originators of subprime mortgages, Section 5.2.2 the issuers of the CDOs backed by these mortgages, Section 5.2.3 the rating agencies whose input was critical in the decision making of investors buying the mortgages, and Section 5.2.4 investors who suffered the actual losses when the CDOs lost value. Section 5.2.5 looks at those investment banks that had substantial direct exposure to the CDO losses, a subset of the CDO issuers studied in Section 5.2.2.

In many ways the investment banks that, as we shall see, had by far the most catastrophic losses in the crisis are the most puzzling of the institutional

groupings. First, the “originate to distribute” paradigm of investment banking would call for only temporary use of a firm’s balance sheet, yet the CDO positions on which they were exposed were long-standing and seemingly permanent. Second, the sophistication that should have resulted from origination of the CDOs should have made the investment banks far less vulnerable to being misled as to their value and riskiness than ordinary investors would be. And third, their well-established risk management processes should have served as a check on such large and reckless exposures. We will spend some time in understanding how all these barriers were breached.

Finally, in Section 5.2.6, we will consider the AAA-rated insurance companies, American International Group (AIG) and the monoline insurers, who became entangled in the crisis and who wrecked valuable franchises in pursuit of business completely tangential to their core competencies.

5.2.1 Subprime Mortgage Originators

One point on which everyone examining the crisis can agree is that a significant contributor was the lax standards and misaligned incentives of the originators of subprime mortgages. Since the originators knew the mortgages were going to be bundled for purchase by an investor, the originators had no direct financial stake in the ultimate value of the mortgages. But the originators had a strong incentive to originate as many loans as possible, given that they were being paid a fee for originations and given the heavy demand by CDO creators for new product.

To take just a few excerpts from the postmortems:

- Brunnermeier (2009): Mortgage brokers offered teaser rates, no-documentation mortgages, piggyback mortgages (a combination of two mortgages that eliminates the need for a down payment), and “no income, no job or assets” (NINJA) loans.
- Hull (2009): “Mortgage brokers started to increase their lending standards in about 2000. . . . How could mortgage brokers and mortgage lenders keep increasing their profits? Their problem was that as house prices rose it was more difficult for first-time buyers to afford a house. In order to attract new entrants into the housing market, they had to find ways to relax their lending standards even more—and that is exactly what they did. The amount lent as a percentage of the house price increased. Adjustable rate mortgages (ARMs) were developed where there was a low ‘teaser’ rate of interest that would last for two or three years and be followed by a rate that was much higher. . . . Lenders also became much more cavalier in the way they reviewed mortgage applications. Indeed, the applicant’s income and other information reported

on the mortgage form were frequently not checked.” Even loan-to-value ratios and FICO scores (the credit score of the home buyer) reported to investors became suspect as “the property assessors who determined the value of a house at the time of mortgage application sometimes succumbed to pressure from lenders to come up with high values” and “potential borrowers were sometimes counseled to take certain actions that would improve their FICO scores.”

- Michael Youngblood, head of asset-backed securities research at Friedman, Billings, Ramsey, is quoted by Peter Coy in the March 2, 2007, issue of *BusinessWeek* as stating that there was “a sudden but little-noticed shift in lenders’ strategy that occurred at the end of 2005: Lenders went from competing for customers on price (by lowering rates) to competing for customers on easy terms (by lowering lending standards).”

The incentives and the results seem clear. What is less clear is why other parties didn’t perceive this incentive structure and begin to exercise caution as evidence of lax standards started to mount. To cite just one example of concerns expressed at the time, a July 15, 2005, *New York Times* article by Edmund Andrews states that the areas that bank regulators find most worrisome “include granting loans equal to 100 percent of the value of homes; granting large loans without due attention to the likelihood of higher monthly payments in the future; and granting ‘no-doc’ (no documentation) or ‘low-doc’ loans that require little or no proof of income or assets.” This article quotes Barbara Grunkemeyer, deputy controller for credit risk at the Office of the Comptroller of the Currency: “You have to ask yourself, why would [a borrower] be willing to pay a quarter-percent more when he could have gotten a lower rate by giving a copy of his pay stub and a W-2 form. There’s a reason they’ve been called ‘liar’s loans.’”

According to Davidson (2007), “mortgage market participants have long recognized that there is substantial risk in acquiring loans originated by someone else” and so require representations and warrants from the originator. If loans sold are later found not to meet the guidelines of the purchaser, the originator must repurchase the loans. But as the push for more new product to feed CDO issuance intensified, more marginal originators became part of the pipeline. The thin capitalization of these newer originators decreased the value of any promises to repurchase mortgages that were not as represented. But the CDO creators, the rating agencies, and the more sophisticated investors should all have been aware of this thinner capital cushion. Why didn’t this lead to more caution? We look at some specific reasons in the sections that follow. One general possibility, suggested by Brunnermeier (2009), is that the assumption that “background checks are unnecessary because house prices could only rise, and a borrower could thus

always refinance a loan using the increased value of the house” may have caused both originators and potential watchdogs to relax their vigilance.

5.2.2 CDO Creators

Much of the blame for the problems with subprime mortgage CDOs must be allocated to the investment banks that created the CDOs. They certainly possessed the greatest amount of expertise, with highly compensated and very experienced structurers, marketers, traders, researchers, and risk managers specializing in mortgage markets and securitization. If any party was well positioned to be aware of the shortcuts that were being taken by the mortgage originators and to spot the potential dangers, the CDO creators were it.

Certainly, the CDO creators cannot claim that they were misled by the rating agencies. As we will see in the next section, the investment banks had full access to the models the rating agencies used in determining ratings. In the process of playing with those models to determine how to optimally structure new issues, the CDO creators probably gained more intimate knowledge of these models than the people within the rating agencies who built them. And the investment banks could bring far more resources than the rating agencies into play, in terms of ability to pay high compensation to attract the best modeling talent (see Tett 2009, 100).

The easy answer is just to focus on incentives. Since the CDO creators were operating on an originate to distribute business model in which all the CDO risk would eventually end up elsewhere, their incentives, like those of the mortgage originators, were to create as much product as possible, since fees earned were tied to volume sold, and to do their best to minimize anyone’s perception of possible loss. One could argue that this is failing to give the investment banks sufficient credit for concern for their longer-term reputations with investors and future losses through possible lawsuits, but after their collectively dismal record in hyping technology initial public offerings (IPOs) in the late 1990s, it would be hard to take that argument very seriously.

But incentives can’t be the whole story, for two reasons. One is that the investors should have been aware of these incentives and the track record the investment banks had shown when faced with these temptations in the past, and so should have exercised their own due diligence. The second is that many of these investment banks failed to execute their desired originate to distribute strategy so egregiously that they wound up being the largest losers when CDO values started to decline. We’ll look at how this occurred in Section 5.2.5; first, let’s see why the investors were willing to trust the CDO creators to the degree they did.

Part of the reason for this trust was undoubtedly comfort that came from the supposedly independent review role of the rating agencies. Why that trust was misplaced we'll examine in the next section. Part came from the skillful marketing of the investment banks, which did their best to convince investors that gains and losses on CDOs would all be about esoteric issues like correlation assumptions, on which the investment banks would be happy to give tutorials to investors, ignoring issues like quality of underlying loans on which the CDO creators possessed insider knowledge that investors could not hope to obtain. And part of the reason for this trust was a structural feature that was supposed to align the interests of the CDO creator with the interests of investors: the retention of the first-loss piece of the CDO by the creator, the so-called equity tranche.

The retention of this first-loss piece meant that this part of the CDO would absorb all of the losses up to some given point and that investors could suffer losses only if the equity tranche was wiped out. The theory was that the investment banks had to closely monitor the quality of assets going into the CDO to avoid large losses on this first-loss piece. The problem was that profits from the tranches that were sold to investors became so lucrative that the CDO creators stopped caring about how much they lost on the equity tranche. According to Hull (2009), "the equity tranche was often regarded as a 'free good.' The originators had obtained adequate compensation for the mortgages from the sales of the other tranches to investors." So much profit had been generated that they could afford to take a full loss on the equity tranche and still come out ahead, or they could afford to purchase protection on the equity tranche from a hedge fund.

5.2.3 Rating Agencies

The rating agencies—Standard & Poor's (S&P), Moody's Investors Service, and Fitch—all badly damaged their reputations by the role they played in providing ratings on CDOs backed by subprime mortgages. They have been the subjects of major investigations, and their role in CDO ratings has led to questions being raised about the role they play in all ratings, including their long-established core business of rating corporate debt. In some ways, their story resembles that of the insurers we will look at later in Section 5.2.6 who jeopardized their core franchise in pursuit of new business.

And yet, the rating agencies had a more plausible case than the insurers that this new business line was related to existing competency. Unlike the insurers, who entered a market that could have survived without them, the rating agencies had a role that was critical to the existence of the CDO business. Most debt investors, from long habit, would have been extremely uncomfortable investing without an agency rating; many were legally

prohibited from investing in debt that did not have a particular minimum rating—it was considered too risky. Ratings tied to probability of repayment were the rating agencies' bread and butter. And they did have several decades' worth of successful experience in rating structured debt that related to mortgages, credit cards, auto loans, and CDOs based on corporate debt. But in 2007 and 2008, the ratings on existing CDOs were downgraded far more violently than any other class of rated securities ever had been, sowing widespread distrust in the agencies. Where did this business model break down?

Many critics in the wake of the CDO crisis point to conflict of interest as the main flaw in the rating agency structure: the rating agencies were being paid by the firms whose bonds they were rating. But that flaw has always existed for all agency ratings, including the core business of rating corporate debt. What is probably more germane is the very close relationship developed between the rating agencies and the investment bank structurers creating the bonds. Structurers had full access to the agency ratings models and a great deal of freedom in deciding what mortgages would go into a CDO. They could play with the structure until they optimized the disconnect between the risk represented by the rating and the true risk, maximizing their profits (see Brunnermeier 2009, 82). There is no comparable freedom to easily change corporate structure. Furthermore, a corporation that does not get the rating it wants will still continue in business and so may choose to pay for the rating anyway. A CDO not getting the rating it wants will not come to market, so the only way rating agencies could get paid is if CDOs did come to market; for further discussion of this point, see Davidson (2007, 4). There is considerable evidence that has come to light since the crisis that rating agencies did succumb to the pressure to find ways to give CDO structures the ratings they needed (see, for example, McLean and Nocera 2010, Chapter 8, and Lowenstein 2010, Chapter 4).

Another significant flaw in the analogy between traditional agency ratings of corporate debt and agency ratings of CDOs was that the ratings methodology for CDOs required the agencies to make forecasts about the state of the economy whereas corporate debt ratings did not. This point is made well by Ashcroft and Schuermann (2008, Section 5.5): CDO ratings “rely heavily on a forecast of economic conditions. Note that a corporate credit rating is based on the agency’s assessment that a firm will default during neutral economic conditions (i.e., full employment at the national and industry level).” (This corresponds to the point made in Section 13.2.1.1, about agency ratings being through-the-cycle and not point-in-the-cycle.) In CDO modeling, by contrast, “uncertainty about the level of loss in the mortgage pool is driven completely by changes in economic conditions” (such as the expected default rates of mortgages, which are closely tied to forecasts of real estate prices). Furthermore, CDO ratings “depend heavily

on quantitative models while corporate debt ratings rely heavily on analyst judgment.” This meant that rating agency senior management, experienced in corporate debt ratings, had little intuition for what was going on in the CDO ratings. And neither rating agency management nor investors had been warned about the precipitous decline in ratings a change in economic outlook could entail, in contrast to the far more steady corporate debt ratings. (CDO ratings are more volatile than corporate debt ratings both because they depend on economic forecasts and because the CDO tranching process concentrates sensitivity to the economy in the higher-rated tranches—see Section 13.4.4.) While these are probably the two most important factors in the rating agency failure, other issues of some weight were:

- The failure of rating agencies to monitor the deteriorating credit standards of the subprime mortgage originators. There was certainly enough publicity about this issue that rating agencies should have been aware of a need to perform some due diligence. Former Moody’s managing director Jerome Fons in testimony to the Financial Crisis Inquiry stated that “never once was it raised to this group [Moody’s high-level Structured Credit committee] or put on our agenda that the decline in quality that was going into pools, the impact possibly on ratings . . .” (FCIR 2011, 121).
- Even if the rating agencies didn’t believe it was their responsibility to check on the mortgage originators, they should at least have been questioning the relevance of historical default data to a rapidly changing situation. It wasn’t just issues being raised about slipping credit standards that should have triggered such questioning, but the sheer explosive growth of the market, which should have been enough to make the relevance of data from prior eras doubtful (compare with Section 8.2.8.2).

The rating agencies’ response to these criticisms was to claim the transparency of their CDO ratings models as a virtue (see, for example, Tett 2009, 100). Anyone could see exactly what the model was doing, the agencies implied, so why blame us if you were later disappointed in the results? This was disingenuous in two directions. First, as emphasized by Tett, it was the very openness and transparency of the models that made them so easy for sophisticated structurers to manipulate. And second, the vast majority of investors certainly lacked the sophistication to understand the workings of the models and had far less capability than the rating agencies for checking loan quality and relevance of historical data.

Why hadn’t these issues surfaced in the reasonably long history of agency ratings of other structured securities? I haven’t seen an analysis of this, but I suspect that while some of these issues were present for other

structured securities, they did not have as strong an impact as they did on the subprime mortgage CDOs. For example, subprime mortgage CDOs are particularly dependent on the state of the national real estate market.

5.2.4 Investors

In many ways, the investors in CDOs can be regarded as the key players in the whole structure. It was the large appetite of investors to own CDO tranches that drove the growth of the market and set the incentives for all the other players. There was a large and diverse universe of these investors, including mutual funds, pension funds, insurance companies, hedge funds, high net worth individuals, and smaller banks (those not involved in the creation of CDOs). It was the CDO investors who were the claimed victims of fraud and misrepresentation by the other players. And it was the CDO investors who, in theory, should have been the ones to suffer the bulk of losses that occurred in the market meltdown.

But somehow, it did not work out that way. The major institutions that suffered the greatest reverses and either went bankrupt or required government bailouts were not primarily the investors but rather the investment banks that created the CDOs. The best overall summary that is available of losses due to the crisis is the International Monetary Fund analysis of April 2009 (International Monetary Fund 2009, Table 1.3) that concluded that out of roughly \$1 trillion in losses on U.S.-originated mortgage CDOs, 60 percent was lost by banks, 25 percent by U.S. government-sponsored enterprises (GSEs), 10 percent by insurers, and only 5 percent by hedge funds, pension funds, and other nonbank financial institutions.

Still, the investors did suffer substantial losses, as can be seen by just looking at the damage claims in lawsuits filed against mortgage originators and CDO creators. The FCIR (2011, 225) asserts that “as of mid-2010, court actions embroiled almost all major loan originators and underwriters—there were more than 400 lawsuits related to breaches of representations and warranties, by one estimate”; for an updated account of the many lawsuits that have been filed, see the Structured Finance Litigation blog: www.structuredfinancelitigation.com). The theory of these lawsuits and of many articles that have been written on the crisis is that deliberately misleading action was taken to entice investors to buy these securities. The previous three sections contain much evidence to support such claims, so there is at least a significant extent to which investors were misled. The question I want to ask here is: To what degree was that the entire story and to what extent were investors knowingly taking on significant risk?

This question is one that has much relevance for risk managers in trying to learn lessons from the crisis. If there were clear signs of riskiness

that investors failed to understand or chose not to focus on, then we have material that can be used in designing better risk management procedures for the future. The principal arguments that investors were to at least some significant degree aware of the risk they were taking on are first that CDO tranches were yielding considerably higher returns than corporate bonds with comparable credit ratings and second that the very illiquidity of the tranches should have been a warning sign against placing too much faith in what they were being told. The historical data I have been able to access shows a steady yield advantage of about 80 basis points for AAA-rated subprime CDO tranches over AAA-rated corporate bonds throughout the period from 2000 to 2006.

5.2.5 Investment Banks

As already noted in Section 5.2.4, investment banks that were among the major creators of CDOs were also the group that suffered the heaviest losses in the 2007–2008 meltdown. This can be seen from the previously cited International Monetary Fund (IMF) analysis that found that 60 percent of the \$1 trillion in losses on U.S. originated mortgage CDOs came from banks while only 5 percent came from the mutual funds, pension funds, hedge funds, and other nonbank financial institutions that were the primary clients to which the investment banks marketed the CDO tranches. It is true that 10 percent of the losses came from insurance companies and many insurance companies were among the primary clients to whom CDO tranches were marketed. But a good portion of the insurance company loss is attributable to AIG and the monoline insurers, and, as we detail in Section 5.2.6, AIG and the monoline insurers can more reasonably be viewed as partners of the investment banks in CDO creation than they can be viewed as clients.

It is also true that the IMF analysis does not distinguish how much of the \$600 billion in losses came from investment banks that were CDO creators and how much was due to smaller banks that may have been clients. But an analysis by the Federal Reserve Bank of Philadelphia (Cordell et al. 2012, Table 11) shows losses of 72% on the \$223 billion of mortgage-backed CDOs originated in 2006 and 84% on the \$163 billion of mortgage-backed CDOs originated in 2007. With loss levels this high, a substantial portion of the losses had to be going to the super-senior tranches that were primarily held by the investment banks that were CDO creators, and Table 12(b) from the same report shows 67% losses on senior AAA tranches originated in 2006 and 76% losses on senior AAA tranches originated in 2007. At the level of an individual investment bank, UBS, which made a public and thorough report to shareholders in April 2008 of the fallout of the crisis, reported 2007 losses related to the U.S. residential mortgage

market of \$18.7 billion, with about \$12 billion due to CDO positions. By early 2009, estimates of total write-downs and credit losses on U.S. financial assets were \$48.6 billion for UBS, \$67.2 billion for Citigroup, and \$55.9 billion for Merrill Lynch (see Zandi 2009, Table 11.2).

This is both unfortunate and surprising: unfortunate, because concentrated losses by large banks are far more damaging to the economy than the same amount of losses spread out over smaller banks and investors; surprising, because the sophistication, intimate familiarity with the product, and originate to distribute business model should all have worked to protect the investment banks.

How then did investment banks wind up with so much mortgage CDO exposure? The initial mechanics of the situation are fairly straightforward. Clients were eager to purchase CDO tranches, thereby selling protection against mortgage defaults, but they were interested only in the *mezzanine* tranches that carried intermediate expected loss. The highest expected loss tranches, the so-called equity tranches, attracting the first losses, could not have achieved investment-grade ratings and were not considered suitable investment vehicles for most clients (though some hedge funds did take on this risk, mostly through derivatives). Also, it was considered appropriate that the CDO creator hold the equity tranche, as explained in Section 5.2.2. The tranches with the lowest expected loss, termed *super-senior* because they supposedly had a statistical probability of loss even lower than AAA-rated corporate bonds, did not have a strong client demand. Because of their very low loss expectation, they carried very low credit spreads, just a few basis points, and it was virtually impossible to find a client that wanted to use valuable balance sheet room to earn such a meager return. (It might be thought that super-senior tranches would be a possibly attractive investment as an alternative to Treasury securities that had similarly low returns, but Treasury securities had many advantages in terms of liquidity and attractive repurchase rate funding opportunities that super-seniors lacked.)

Here was a dilemma for the investment banks. To create more mezzanine tranches for which there was high demand, they also needed to create super-senior tranches for which there was virtually no demand. Of course, one alternative would have been to substantially raise the yield on the super-seniors to the point that demand was created, but this would have so severely cut into the profitability of the overall transaction that it wasn't seriously considered. Their only alternatives were to stop the flow of lucrative new business or to pile up super-senior tranches on their own balance sheets. They almost all chose the latter option. As Chuck Prince, the soon-to-be-ex-CEO of Citigroup, infamously said in July 2007, "As long as the music is playing, you've got to get up and dance. We're still dancing" (FCIR 2011, 175). It was this continuous buildup of super-seniors, totally lacking

a liquid market, that was the source of almost all of the large CDO losses suffered by the investment banks. For example, the UBS report to shareholders showed that about \$9 billion of its \$12 billion 2007 losses on CDOs were due to super-senior tranches. Other large investment banks that followed this pattern included Citigroup, Merrill Lynch, Morgan Stanley, and Bear Stearns (see Tett 2009, Chapter 9). Writing generally about investment banks that experienced large losses in 2007, the Senior Supervisors Group report of March 2008 on page 8 states that “some firms continued to underwrite or increase their exposures until the summer of 2007 despite an array of data indicating rising stress in the subprime mortgage market and worsening credit market conditions.”

If management of these banks had placed sensible limits on the size of super-senior holdings or had insisted on mark-to-market valuations of the holdings that reflected their total lack of liquidity (thereby lowering the profit that could be recognized on new CDO issuance and shrinking bonus pools), the entire mortgage CDO creation process would have come to a halt at a fairly early stage and the damage to the financial industry and the world economy would not have been nearly as severe. As Richard Bookstaber, an experienced senior risk manager, put it in his testimony before the Financial Crisis Inquiry Commission, “As everybody in any business knows, if inventory is growing, that means you’re not pricing it correctly. . . . It was a hidden subsidy to the CDO business by mispricing” (FCIR 2011, 196). What stopped reasonable action from being taken? The banks seemed to be operating as if they possessed a split personality. In one part of the firm, the CDO creation teams were behaving as if all risk was being taken on by clients, as if the originate to distribute mechanism was operating smoothly. This left them free to ignore warning signs about the increasingly poor quality of the mortgages being originated and about the potential impact on losses if the housing price bubble burst. In another part of the firm, super-senior tranche holdings were growing by leaps and bounds.

One possible answer is that the traders and structurers who had the greatest degree of knowledge of the situation just didn’t care about the health of the firm and so did their best to mislead senior managers and risk managers about what was really going on. All they cared about was generating one more round of spectacular bonuses. They treated risk managers and senior management as just another set of clients to whom product needed to be sold—in this case, super-senior tranches. While this no doubt contains an element of truth, it can’t be the entire story. Risk management of investment banks has always been built upon a healthy skepticism about the motivations of front-office personnel, as we saw in Section 2.1 and as we will consider at greater length in Chapter 6. So let’s try looking at some other possible explanations. We’ll look at supporting evidence for them in

this section, and then draw on this material to examine risk management lessons in Section 5.4.5, using the same headings in both sections.

Before going any further, let me clear up one possible source of confusion. In the midst of the crisis, there were many news reports concerning disputes over the marking-to-market of distressed securities—should firms holding securities experiencing what was hoped to be a temporary bout of illiquidity show losses based on the fire-sale prices at which these securities were trading in the market? Since these disputes occurred during the same period that large losses were being recorded by the investment banks on their super-senior tranches, it might have seemed that the super-senior losses were at least partially an accounting fiction. But as we’ve just recounted, the super-senior tranches never had a liquid market at any time, so their marks were always based just on the best judgment as to ultimate losses. Whatever the merits of the accounting debate over other securities that were caught up in the crisis (we’ll have more to say about this in Section 5.3.2), the losses reported on super-seniors always represented best estimates of true ultimate cost.

5.2.5.1 Reliance on Inadequate Derivatives Protection One fairly common response to the inability to find clients to buy super-senior tranches was to hold on to the super-senior tranches but hedge the risk with derivatives. This should clearly have been viewed with suspicion by risk managers—if you couldn’t find clients willing to buy super-seniors, why were you able to find clients willing to take on the risk through derivatives? Wasn’t there some substantial difference in the amount of risk being shed in the two different transactions?

By saying that suspicion should have been aroused, I do not mean that it was obvious that the risk was not being fully hedged, just that thorough analysis should have been initiated. I have seen cases in which firms were willing to fully absorb risk but had limitations on balance sheet usage, perhaps because of lack of access to good funding sources or perhaps due to statutory restrictions. Analysis in these cases showed that the sellers of derivative protection were providing sufficient collateral and margining to keep risk very close to what would have been achieved with an outright sale, though with different funding requirements. (Discussion of collateral and margining will be found in Section 14.3.3.)

Had thorough analysis been performed in the case of the derivatives hedging super-seniors, a very different picture would have emerged. Many collateralization and margining agreements were either nonexistent or of very limited value. For example, Lowenstein (2010, Chapter 9) reports that Vikrim Pandit, on taking over as CEO of Citigroup, was “stunned to hear” from New York State’s top insurance regulator that “Citigroup’s insurance

did not entitle it to payments as the prices of CDOs declined.” Citi had “insurance on defaults, not on market value.” Given the long-dated nature of the CDOs, “Citi (and every other bank with insurance) would have to wait years to file claims, at which point the insurers could be out of business.” This was very typical of insurance purchased (whether through insurance contracts or through derivatives) from the major suppliers of super-senior insurance, AIG and the monoline insurers (whose role we will look at more closely in the next section). AIG did offer some collateralization, partly to gain a competitive advantage on the monoline insurers, which offered none (McLean and Nocera 2010, 190–191). But some of this was weak collateralization that would be triggered only under extreme circumstances, by which time AIG might already be facing difficulties (as proved to be the case).

When we look at risk management lessons in Section 5.4, we’ll do a detailed analysis of all the alarm bells this arrangement should have sounded. But, as I will detail there, the risk management methodology for identifying the large gap in risk reduction between outright sale and insurance protection was well known and thoroughly disseminated well before these deals were booked. If this was not highlighted to senior management and regulators, it constituted a major breach of risk managers’ responsibilities.

5.2.5.2 Reliance on Off-Balance-Sheet Vehicles If you couldn’t find clients interested in holding super-seniors because of their very thin spreads over funding costs, there was one more trick that could be used: If you set up an entity that could hold the super-seniors and issue short-dated AAA-rated debt against them, the normal upward slope of the yield curve would provide enough cushion to generate some extra spread to entice investors in short-dated AAA debt (Tett 2009, Chapter 6).

The primary practioner of this bit of financial legerdemain was Citigroup, which began placing its super-seniors into *structured investment vehicles* (SIVs). SIVs were officially independent enterprises whose commitments Citi had no legal responsibility for and so did not have to be consolidated onto Citi’s balance sheet (leading to their classification as *off-balance-sheet* vehicles). But SIVs were funded by commercial paper (CP), and commercial paper investors would invest only in AAA-rated entities. Even if the rating agencies regarded the super-senior tranches as AAA, the short-dated funding and long-dated assets of the SIVs raised the issue of what would happen to the CP holders if new CP investors could not be found when the old CP matured. To obtain an AAA rating for the SIV, Citi needed to offer liquidity puts that would allow the SIV to sell the super-seniors back to Citi at par if the SIV ran into problems funding them (McLean and Nocera 2010, 240–241). Citi wrote about \$25 billion of these liquidity puts.

The key risk management question would now be what probability of loss to assign to these liquidity puts. The attentive reader will not be surprised that Citi's internal risk models estimated so remote a possibility of the liquidity puts being triggered that they only needed to hold 0.16% in capital against the put (FCIR 2011, 138). Hence only \$40 million in capital would be required against Citi's \$25 billion in liquidity puts. And there seems no evidence that Citi continued to view the super-seniors placed into the SIVs as still being part of its risk book. But clearly, placing the super-seniors into an SIV made practically no difference to Citi's risk position. In the event that there would be losses on the super-seniors, it would be virtually certain that the liquidity put would be exercised. This is another clear case of violation of one of the well-established rules of risk management, the need to account for wrong-way risk (see Section 14.3.4 for more explanation of wrong-way risk).

5.2.5.3 Use of Faulty CDO Models Felix Salmon's February 2009 story for *Wired* magazine, "Recipe for Disaster: The Formula That Killed Wall Street" (Salmon 2009) brought David Li's version of the Vasicek model to the attention of a wider audience than financial industry quants (see Section 13.3.3 for a description of the model). The article led off with statements such as: "One result of the [2008 financial system] collapse has been the end of financial economics as something to be celebrated rather than feared. And Li's Gaussian copula formula will go down in history as instrumental in causing the unfathomable losses that brought the world financial system to its knees." With this as background, I was somewhat surprised in my survey of the principal book-length writings and journal articles on the crisis to see scant mention of either the Li model or the Vasicek model. Did faulty CDO modeling play a significant role in the crisis?

The case for faulty CDO models playing, at best, a minor role in the crisis would go as follows:

- The Li model was primarily being used as an interpolation tool from more common tranches for which price quotes could be obtained to less common tranches. As such, its use was very similar to that of the Black-Scholes model in interpolation of options prices and the use of fitting to a correlation skew implied by the market (see Section 13.4.2) as part of the interpolation shows that the Gaussian copula assumptions of the Li model were not being taken very seriously by the traders using it.
- The Li model was also being used as an aid to intuition (see Section 13.3.3) and as such it did its job admirably. In fact, it was particularly valuable in letting users see the degree of systematic risk embedded

in different tranches, which should have directed attention to the riskiness of super-senior tranches (see Section 13.4.4).

- The emphasis on the correct estimation of correlation levels and the shape of the correlation copula was very important for traders making decisions on the value of tranches. Had the tranches been liquid, this would also have been important for risk managers, in estimating where liquid positions could be exited. But given the illiquidity of super-senior tranches, stress testing large changes in the common factor, closely linked to real estate prices, was overwhelmingly more important for risk managers than stress testing of either correlation level or copula shape.
- When investment banks wanted to perform more fundamental analysis of tranche pricing and risk, they were hardly lacking for more sophisticated versions of CDO models, as the discussion in Sections 13.3.3 and 13.4.2 clearly show. Many of the models cited in these sections date from the first half of the 2000s decade and were widely available—often referenced and explained in papers published by investment bank research teams, in the well-known book by Schonbucher (2003), and in many issues of *Risk* magazine from that period.

And yet there is one key way in which CDO models utilized by investment banks in this period were misleading. Too much emphasis was placed on fitting model parameters to observed market prices without an adequate consideration of the degree of illiquidity that pervaded many sectors of this market, including the entire super-senior sector. This may have helped encourage the definitively faulty analysis we discuss in Sections 5.2.5.6 and 5.2.5.8.

5.2.5.4 Reliance on External Ratings It is uncontroversial that the rating agencies played a significant role in fueling the demand for CDO tranches by investors. But could they have also played a role in the willingness of investment banks to tolerate so large an exposure to super-seniors? At first glance, this seems preposterous. As we noted in Section 5.2.2, the investment banks in their role as CDO creators had intimate knowledge of the rating agency models and knew the extent to which they had manipulated those models. How could they then rely on those models to take comfort with their exposure?

And yet one finds in the March 2008 UBS report to shareholders (UBS 2008, Section 5.3.2) that the UBS market risk control group's "VaR methodologies relied on the AAA ratings of the Super Senior positions. The AAA ratings determined the relevant product-type time series to be used in calculating VaR. . . . As a consequence, even unhedged Super Senior positions contributed little to VaR utilization." Tett (2009, 139) quotes Peter Kurer, a member of UBS's board, as saying, "Frankly most of us had not even heard

the word ‘super-senior’ until the summer of 2007. We were just told by our risk people that these instruments are Triple-A, like Treasury bonds.” Anecdotal accounts I have heard from other investment bank risk managers indicate that UBS was not alone in utilizing AAA ratings of tranches as an invitation to calculate risk statistics for them based on time series of price moves of AAA-rated corporate bonds. The March 2008 Senior Supervisors Group report on the risk management practices of investment banks leading up to the crisis states that at some firms “internal risk capital measures that relied too much on agency ratings underestimated the true price of the risk of such positions” and that some firms “tended to assume that they could apply the low historical return volatility of corporate credits rated Aaa to super-senior tranches of CDOs” (p. 5). It further states, “Given that the firms surveyed for this review are major participants in credit markets, some firms’ dependence on external assessments such as ratings agencies’ views of the risk inherent in these securities contrasts with more sophisticated internal processes they already maintained to assess credit risk in other business lines” (p. 3).

The impression left is consistent with the picture of front-office personnel not sharing their knowledge of rating agency model limitations with risk managers. We address the lessons for risk managers in Section 5.4.1.

5.2.5.5 Overreliance on VaR Measures As we have just seen, UBS (and, anecdotally, some other investment banks) used the AAA ratings of super-seniors as a shortcut in VaR calculations, essentially treating any AAA-rated security as if its price movements could be represented by a time series drawn from AAA corporate bond prices. This was clearly an error—as discussed in Section 13.4.4, the volatility of tranche prices is expected to be quite different from the volatility of corporate bonds of the same rating. But an even more important question is: Why were firms even bothering to calculate VaR, a measure of vulnerability to short-term price fluctuations, for an instrument as illiquid as super-seniors?

Now, perhaps this was just a calculation of VaR for a liquid proxy hedge of the super-seniors, and the bulk of the risk was going to be evaluated elsewhere (a measure I will strongly advocate, in Sections 6.1.2 and 8.4, for highly illiquid instruments). If this was the case, then even the use of the computational shortcut might be justified—you would be choosing a portfolio of AAA corporate bonds as your liquid proxy hedge. It may not be the best choice, but as long as you are calculating the long-term risk of the hedge separately no great harm will be done. But this does not appear to be the way UBS (or, anecdotally, some other investment banks) were operating. VaR was intended to be the primary risk measure for the super-seniors. Quoting UBS (2008, Section 6.3.2), “Investment bank business planning

relied upon VaR, which appears to be the key risk parameter in the planning process. When the market dislocation unfolded, it became apparent that the risk measure methodology had not appropriately captured the risk inherent in the businesses having Subprime exposure.” Dash and Creswell (2008) relate that “when examiners from the Securities and Exchange Commission began scrutinizing Citigroup’s subprime mortgage holdings after Bear Stearns’s problems surfaced, the bank told them that the probability of these mortgages defaulting was so tiny that they excluded them from their risk analysis.”

This brings us to the broader question of the extent to which the illiquidity of the super-seniors was being factored into risk measurement.

5.2.5.6 Failing to Account for the Illiquidity of Super-Senior Tranches The illiquidity of super-senior tranches should have been evident to anyone involved in investment banking, even those most remote from direct trading and marketing of CDOs, just by the fact that it was such a problem to find willing buyers. But the Senior Supervisors Group report of March 2008 finds that “firms that faced more significant challenges in late 2007 . . . continued to price the super-senior tranches of CDOs at or close to par despite observable deterioration in the performance of the underlying . . . collateral and declining market liquidity” (p. 3). The UBS report to shareholders Section 6.3.6.4 states that “The Super Senior notes were always treated as trading book (i.e., the book for assets intended for resale in the short term), notwithstanding the fact that there does not appear to have been a liquid secondary market and that the business tended to retain the Super Senior tranche.”

Why were firms treating such clearly illiquid instruments as liquid? One clear motivation is alluded to in the same section of the UBS report: “Treatment under the ‘banking book’ would have significantly changed the economics of the CDO desk business as this would have increased the required regulatory capital charges.” Classifying assets in the trading book, available for resale in the short term, attracted more favorable capital treatment than the same assets placed in the banking book, intended to be held. Note that this is just a statement of intention—nothing stops you from selling assets in the banking book; loan sales occur all the time. But this statement of intention was allowed to impact required regulatory capital, a major driver of the economics of a product. This loophole was closed after the CDO-fueled crisis revealed its shortsightedness; the Bank for International Settlements (BIS) Incremental Risk Capital Guidelines of July 2009 made capital requirements for credit products held in the trading book and banking book essentially equivalent—see PricewaterhouseCoopers (2011, Section 4.6.3.5). It also impacted the balance sheet reporting that might impact public perception of the degree of liquidity of the firm’s assets.

My guess, and it's only a guess, is that the mechanism that operated at some firms was that the potential liquidity of CDOs, including super-seniors, had been emphasized in order to obtain the favorable capital treatment—securities are, in general, more liquid and likely to be sold than loans are. While this accounting decision should not have forced a similar classification by risk managers, it is not uncommon for this kind of distinction between accounting principles and risk management principles to get blurred.

5.2.5.7 Inadequate Stress Tests Another possibility is that there was widespread conviction that risks that threatened mezzanine tranches could not spread to super-senior tranches. I find this difficult to accept, since the simplest possible CDO model could easily show the vulnerability of even super-senior tranches to a large downturn in housing prices, the sort of economic stress scenario that risk management groups are supposed to run routinely (see Section 13.4.4 on the usefulness of the Vasicek model in analyzing vulnerability of senior tranches to systematic risk).

One viewpoint I have frequently encountered in conversations with risk managers who were caught up in the crisis goes something like this: “Place yourself back in 2006 and suppose you were to stress test your CDO portfolio. Suppose that you chose to shock U.S. house prices down 30 percent to evaluate the impact on the prices of super-senior tranches that you held. You would have been laughed out of the room—no one would have found this a plausible stress test scenario.” A published account of a closely related incident can be found in Lewis (2011, 211–212). With all due deference to the fact that I was not actively involved in risk management of any of the impacted firms during this critical period, I must respectfully but strongly dissent from this view.

First, looking at the history of super-senior tranches at many of the impacted firms, you find an active interest in purchasing protection on these tranches from AIG (see Tett 2009, 134–136; FCIR 2011, 139–142, 202–204). It is only when AIG's appetite for selling protection dried up that firms turned to either absorbing the risk completely or utilizing clearly inadequate substitutes, such as buying uncollateralized protection from inadequately capitalized monoline insurers. If losses on super-seniors weren't going to occur under any *plausible* shock, why spend money and effort on buying protection? The rejoinder might be that this was “just to keep the risk managers (or the accountants or the regulators) happy.” But keeping risk managers or accountants or regulators happy means addressing a shock that *they* would find plausible; what made them stop finding it plausible at just the moment the protection could no longer be purchased?

Second, it is not difficult to find mainstream economic analysis that viewed a large drop in housing prices as not just plausible but reasonably

probable. Just using the *Economist* magazine as a representative voice, one finds articles in the issues of December 9, 2004 (“Flimsy Foundations”); December 8, 2005 (“Hear That Hissing Sound?”); and September 7, 2006 (“Checking the Thermostat”), all talking about U.S. house prices being overvalued by amounts ranging from 20 percent to 50 percent and all talking about the serious possibility of the “bubble bursting.” This was not some then-unknown junior economist crying in the wilderness; this was in a prominent mainstream publication that is required weekly reading for virtually everyone in the financial industry. And the opinions were backed by detailed statistical analysis of historical relationships of housing prices to rental prices and to incomes. At the same time, the Yale economist Robert Shiller, already prominent for the timely concerns he had expressed about the Internet bubble and noted for his expertise in the field of housing prices, was quoted by David Leonhardt in the *New York Times* on August 21, 2005, as “arguing that the housing craze is another bubble destined to end badly, just as every other real-estate boom on record has. . . . He predicts that prices could fall 40 percent in inflation-adjusted terms over the next generation.”

Now certainly there is room for disagreement among economists and financial analysts. Someone making a strong and detailed argument for a given viewpoint is no reason it can’t be rejected as a most likely or even reasonably probable view. But to reject it as a plausible view I find disingenuous. My guess would be that it is far more likely that risk managers were buying into a wholly unsupported view of the liquidity of the super-seniors, as documented in the preceding subsection. And if you are treating super-seniors as liquid, then of course a drop of 40 percent in housing prices over the next generation is none of your concern since you only need to be worried about what might be reflected in the market over a period of a few weeks.

5.2.5.8 Inadequate Analysis of Statistical Hedging Faced with the inability to fully eliminate super-senior exposure, some investment banks very sensibly began seeking more liquid hedges that would eliminate at least some of the exposure. The question is not whether this was a prudent strategy (it was), but whether risk managers adequately analyzed the resulting risk. One case in which they conspicuously did not do so is at UBS. Section 4.2.3 of UBS (2008) states that the Amplified Mortgage Portfolio (AMPS) consisted of super-senior positions “where the risk of loss was initially hedged through the purchase of protection on a portion of the nominal position. . . . This level of hedging was based on statistical analyses of historical price movements that indicated that such protection was sufficient to protect UBS from any losses on the position.” In Section 6.2.3, the report states that once

hedged through AMPS trades, the super-senior positions were considered fully hedged and therefore did not appear in either VaR or stress test reports. The report further notes, in Section 6.3.6.1, that even though an internal audit had “identified certain risks in the Subprime trading books, senior risk control did not appear to take those issues into account when concluding that positions were hedged.”

To get a better understanding of statistical hedging, we need to add just a bit of complexity to the basic picture we have painted of trading in mortgage CDOs. In addition to the tranches that were based on dividing up actual pools of mortgages, some synthetic tranches based on reference portfolios began to trade (see Section 13.4.1 for details on synthetic tranches). To some extent, these synthetic tranches were just side bets between investors who wanted to sell protection on mezzanine tranches (the vast majority) and a few investors looking to buy protection, either as an offset to previous sales or because of a belief that mortgage defaults were going to exceed market expectations. The investment banks’ involvement with these side bets would have been that of market maker in a reasonably liquid market. But to some extent, these synthetic tranches offered an opportunity to investment banks looking to reduce their exposures to mortgage tranches. An entertaining and informative book focused on the market for synthetic tranches of subprime mortgages is Lewis (2011). Lewis provides a detailed narrative of the role these synthetic tranches played in generating large profits by hedge fund managers, such as John Paulson and Steve Eisman, as well as traders for investment banks, such as Deutsche Bank’s Greg Lippmann, on bets that mortgage defaults would exceed expectations.

The same market fundamentals drove this market as drove the market for pool tranches, namely the strong investor demand for selling protection on mezzanine tranches and little interest in either equity or super-senior tranches. So the synthetic tranches did not offer a direct offset to warehoused super-senior exposure. But synthetic tranches did make it possible for investment banks to buy more protection on mezzanine tranches than they had created through the pool tranching process. So they could consider offsetting some of their super-senior position in a particular portfolio by buying mezzanine protection on a reference portfolio either identical to or closely related to the portfolio the super-seniors were exposed to. Here’s where the statistical analysis came in: What was the best dollar volume of a mezzanine tranche to buy protection on to hedge a given volume of super-senior tranche, and just how large was the risk offset?

As you would expect from the large difference between hedging against changes in credit spread and hedging against changes in default exposure, illustrated in Section 13.1.2.2, there was going to be a large residual risk in some direction. And given that it would have been prohibitively expensive to purchase

true default protection for super-seniors using mezzanine tranches, you can be certain that the hedges actually employed were primarily hedges against credit spread movement, not against default. This highlights just how misleading it was, and how easy it should have been to spot the error of UBS treating statistical hedges as fully eliminating risk. Even in the far more liquid vanilla options market, no one treats positions that are “neutral in the Greeks” as having no residual risk (see Section 11.4, particularly the discussion of Table 11.6).

5.2.5.9 Too Big to Fail Finally, there is the question of why, leaving aside any probabilistic analysis of risk, the sheer size of the positions didn’t trigger alarms. Let me offer an analogy directly from my own experience. During the late 1990s, I was in charge of risk management for Chase’s derivatives business. A very conspicuous part of that business was the new, rapidly growing, and very profitable CDO business, based on commercial loans, not residential mortgages. But like the residential mortgage CDOs of the mid-2000s decade, the commercial loan CDOs of the late 1990s were starting to run into an accumulation of super-senior risk that the bank was finding difficult to buy protection on. While I was, whether correctly or incorrectly, quite convinced that the probability of loss on this super-senior risk was extremely low, making presentations to the firm’s risk committee supporting this view, I was just as strong in my opposition to the continued buildup of super-senior risk on the firm’s books. Even though limitations on the growth of super-seniors ultimately meant limitations on the growth of the very profitable CDO business as a whole (for reasons similar to those discussed earlier for mortgage CDOs), the skeptical views of me and my similarly minded risk management colleagues prevailed. Super-seniors were piling on exposure to what was already the firm’s largest vulnerability as a major commercial lender, exposure to a drastic economic downturn. No matter how remote a possibility we might have regarded such a downturn, it was not a scenario we could completely dismiss. Tett (2009, 65–66) reports a similar decision-making process around the same time at JPMorgan, prior to the merger with Chase. While all anecdotal recollections of past risk management triumphs, perhaps including my own, should be taken with a grain of salt, what I saw of JPMorgan’s exposures going into the merger were consistent with Tett’s account.

Arguments were offered by a few front-office people at the time of this decision that “in the case of that drastic an economic downturn, the firm will need to be rescued by the Federal Reserve anyway, so what difference does the size of the rescue make?” These arguments were considered wholly without merit by both risk managers and senior management. But one wonders if perhaps this kind of view was behind some of the decision making in 2005–2007.

In the wake of the 2007–2008 collapse, suspicions have certainly been expressed that this confidence that regulators and the government owned the downside on big bets was explicitly or implicitly part of the calculation that drove the CDO-creation machine past reasonable limits. “Moral hazard,” “Greenspan put,” and “too big to fail” have all become part of the common vocabulary used in the postmortem analyses of these decisions (see, for example, FCIR 2011, 57, 61, 341, 356). It is certainly in line with the moral hazard story we told in Section 2.1. And the greater the belief that your firm’s outrageous positions are not out of line with the outrageous positions of your competitors, the greater the tendency for arguments based on ultimate regulator rescue to gain traction.

The usual counter of those who find these arguments specious, leaving aside considerations of morality that might not be shared by all discussants, can be summed up in a well-circulated, but presumably apocryphal, story that goes back at least to the 1970s. In this story, the CEO of a large commercial bank attending an industry conference finds himself at a men’s room urinal next to the crotchety and brusque chairman of the Federal Reserve (in those days, all Fed chairmen were expected to be crotchety and brusque—and male). Looking around and seeing no one else in the room, he whispers to the chairman, “Just between us, would the Fed come to our rescue in a crisis?” The chairman, without looking up, responds, “That is a question I would need to discuss with your successor.”

The moral of the story is supposed to be that the penalties for putting your firm in the position of being rescued are personally severe. And certainly one sees evidence of the regulators attempting to enforce this, going out of their way to demand that the price JPMorgan paid for Bear Stearns was punitive to the Bear Stearns stockholders, which included most of the firm’s longtime employees (see McLean and Nocera 2010, 347). And those of us who fought against a “too big to fail” mentality can point to the benefits to firms like Goldman Sachs, JPMorgan, and Deutsche Bank, whose need for government assistance was much less pronounced than Citigroup or Merrill Lynch or UBS. But in the modern era of outsized compensation for senior executives and star traders, which may include so-called golden parachutes protecting them against the personal consequences of failure, is the government ownership of the downside becoming too great a temptation for risk takers?

5.2.6 Insurers

Compared to the voluminous literature about the investment banks in the CDO meltdown, far less has been written about the insurance companies whose sale of protection for super-senior tranches led to the destruction of

valuable business franchises. And what has been written about the insurance companies is mostly from the standpoint of the errors investment banks made in their reliance on this insurance. My primary source for what follows is FCIR (2011), which addresses AIG on pages 139–142, 200–202, 243–244, 265–274, 344–352, and 376–379 and the monoline insurance companies on pages 204–206 and 276–278.

For the most part, these insurance companies appear to have regarded the super-senior tranches of subprime mortgage CDOs as being virtually without risk of loss. Their analysis can therefore be subject to the same critical examination we have just been through in the previous section for the investment banks. But there is one major difference: The investment banks possessed some genuine expertise in evaluation and modeling of subprime mortgages and of CDO structures. The insurance companies possessed none of this expertise and just relied on analysis by the investment banks and rating agencies for their assurance that risk of loss was practically nonexistent.

A telling quote comes from Alan Roseman, CEO of ACA Insurance, one of the monoline insurers: “We were providing hedges on market volatility to institutional counterparties. . . . We were positioned, we believed, to take the volatility because we didn’t have to post collateral against the changes in market value to our counterparty . . . [and] we were told by the ratings agencies that rated us that mark-to-market variations [were] not important to our rating, from a financial strength point of view at the insurance company” (FCIR 2011, 276). If this attitude was typical, then the insurers were operating on the premise that there was no genuine risk of loss on the super-seniors, just annoying fluctuations in mark-to-market accounting, presumably due to technical liquidity factors. This view would see the insurers collecting a fee for providing an accounting arbitrage as opposed to being paid for absorbing risk (the accounting arbitrage would arise from an uninsured super-senior holding at an investment bank being subject to mark-to-market earnings fluctuations; once insured, it would no longer need to be marked to market and the insurers did not have mark-to-market accounting). If you are just being paid for an accounting arbitrage, then you don’t require any expertise in assessing risk, just a knowledge of accounting rules.

While the evidence for how typical this view was is not clear, it certainly does appear that little concern was shown by any of the insurers involved for making their own assessment of credit risk. The only one of these insurers that did begin to show some concern about the volume of exposure they were taking on was AIG (FCIR 2011, 200–201), but its slowdown in taking on CDO risk still left it holding \$79 billion in CDO exposure. MBIA, Inc., another of the monoline insurers, stated, according to Norris (2009), that “‘the due diligence standard for a monoline insurer, which MBIA followed,’ did not involve looking into the quality of the securities underlying

the securities being insured . . . it primarily relied on the assurances by Merrill Lynch and the credit ratings of Moody's and Standard & Poor's." While this was part of an MBIA suit brought against Merrill Lynch and so might be expected to exaggerate MBIA's lack of sophistication, it is still revealing that such a claim would be even plausible relative to a business line in which the insurers had bet their entire franchises.

5.3 THE SPREAD OF THE CRISIS

The crisis that began in the subprime mortgage CDO market spread to markets, instruments, and institutions that had no direct involvement with either mortgages or CDOs. There were two primary paths through which this spread: contagion through credit exposure to impacted firms, which we examine in Section 5.3.1, and contagion through market impact, which we examine in Section 5.3.2.

5.3.1 Credit Contagion

The most direct way for a crisis to spread is through credit exposure to impacted firms. This was certainly a prime ingredient in the 2007–2008 crisis.

One of the major paths for credit contagion was the great extent to which financial firms had large counterparty credit exposure to one another through the derivatives markets. I am not including in this the CDO-related counterparty exposure of many firms to AIG and the monoline insurers, since this was part of the fundamental process creating the crisis. But many firms that may have had no dealings in CDOs had heavy exposure to firms that did have large CDO losses on other derivative contracts such as interest rate and foreign exchange swaps. And this was a decided worry for regulators, as they had to decide on how to handle firms approaching bankruptcy. One can see in the reporting on regulators' decisions during this period just how big a worry this was (see, for example, FCIR 2011, 291, 329). Some contracts would not be backed by collateral and would result in outright loss; even where there was collateral, there would still be losses resulting from the market impact of so many counterparties simultaneously rushing to sell the collateral and to replace the defaulted derivatives positions. Not only did regulators need to worry about the direct impact on derivatives counterparties of a default, but they also had to be concerned about the potential freezing of derivatives markets as worry about defaults would cause reluctance to enter into new contracts. This in turn could worsen the situation for

counterparties of a defaulting firm, since they might have difficulty finding a replacement for a defaulted derivative contract, exacerbating the original loss. The bankruptcy of one firm might then drive other firms into bankruptcy in an ever widening circle.

It is easy to understand the frustration of regulators at being placed in this position. Derivatives trading had originally been almost exclusively conducted on exchanges that had well-developed procedures for minimizing credit exposure. A major argument of large investment banks in setting up over-the-counter derivatives markets as alternatives to exchange-traded derivatives was that they had the credit systems and expertise that were capable of managing the extra credit risk that would arise. But now they had apparently done so poor a job of managing this that they needed to be bailed out by regulators, and this only a decade after the Long-Term Capital Management crisis had supposedly led to reforms in counterparty credit management (see Section 4.2.1). Another major path for credit contagion was through financial firms that had made direct loans to firms whose CDO positions threatened them with bankruptcy. This direct lending was primarily in very short maturity instruments, such as commercial paper. Because of the short maturity and the previous sound financial status of major financial firms, this paper was very highly rated by the rating agencies and was supposed to be a very safe investment. Money market mutual funds that bought a diverse portfolio of this paper were considered nearly as sound as government-guaranteed bank deposits, and regulators worried about the impact on small investors if defaults on commercial paper drove big money market funds to the point of “breaking the buck,” that is, not having sufficient funds to pay back investors’ principal. This, too, was a major concern for regulators as they considered how to deal with firms close to bankruptcy (see Sorkin 2010, Chapter 17).

5.3.2 Market Contagion

If you look at Table 13.6, you will see the normal cyclic pattern of defaults of corporate borrowers. Lending institutions have survived this cyclic pattern for decades, building up reserves and capital during times of low defaults that can be used as a buffer against times of higher defaults. But providing credit is a business that requires patience—on the part of bank management, of bank regulators, and of those who invest in banks. When defaults start occurring at an accelerated pace, banks will start cutting back on the volume of new loans, but they won’t start panicking and trying to sell off large blocks of their remaining loans.

Credit derivatives, such as credit default swaps and CDOs, brought the promise of increased liquidity to the business of bank lending. When used

reasonably, these instruments can be part of a blended strategy, in which some portions of the loan portfolio are judged liquid and managed accordingly while other portions continue to be viewed as illiquid, with a management approach that matches their lack of liquidity. Chapter 13 of this book, and particularly Section 13.3, outlines what I consider an appropriate blend of tools for managing a portfolio that contains both liquid and illiquid credit exposure.

By falsely labeling all of the subprime mortgage CDOs as liquid in their desire to obtain more favorable regulatory capital treatment, the investment banks created a dilemma. When falling housing prices started to threaten widening defaults on these CDOs, there was no cushion of reserves or capital to allow for patience, as would have been the case if a large portion had properly been labeled illiquid. And the accounting for liquid instruments meant that banks had to recognize earnings losses immediately through mark-to-market accounting and therefore needed to immediately take action to get capital ratios back to allowable levels, since earnings gains and losses immediately impact capital.

When a truly liquid position suffers a mark-to-market loss that requires an increase in capital, there is a readily available remedy: sell some of the liquid position to reduce the need for capital and also reduce a possible source of further losses requiring capital. This is why mark-to-market accounting, stop-loss limits, and capital allocations designed to allow liquidation of positions over a temporary period of illiquidity fit so well with liquid positions (discussed at greater length in Section 6.1.1). But if you have been only pretending that a position is liquid, you don't have this option. Since large losses usually occur in periods of economic stress, when raising new capital from investors is difficult and costly, your only remaining choice is selling other positions that truly are liquid to reduce the need for capital. But that doesn't get the illiquid positions off your books, and if they continue to lose money, you will need to go through this cycle all over again. This is a sketch of how losses on illiquid positions treated as liquid can spread a crisis to other markets by continued forced selling of positions that were not related to the illiquid positions.

This is roughly what occurred during the 2007–2008 period, but it was exacerbated by the realization that positions that had been labeled as liquid and as virtually immune to losses were in fact very vulnerable. This raised the level of suspicion in the markets about any asset or derivative that might conceivably have some type of hidden risk. This was another factor in driving down prices and drying up liquidity in other markets (see, for example, Greenlaw et al. 2008, Section 2.1).

The full mechanics through which depressed asset values lead to market contagion, with many illustrations from the 2007–2008 crisis, are

covered in more detail in Duffie (2011, Chapter 3). The discussion of this in Brunnermeier (2009, 92–94) is also useful.

5.4 LESSONS FROM THE CRISIS FOR RISK MANAGERS

My numbering of subsections is designed to allow easy reference to the discussion of the mechanics of the crisis in Sections 5.2 and 5.3. Sections 5.4.1 through 5.4.6 correspond to sections 5.2.1 through 5.2.6, respectively, while Section 5.4.7 corresponds to Section 5.3.1, and Section 5.4.8 corresponds to Section 5.3.2.

5.4.1 Subprime Mortgage Originators

From the viewpoint of internal risk management, the lessons of the crisis regarding both subprime mortgage originators and CDO creators center on issues of legal and reputational risk. These nonquantitative areas of risk management do not align with the focus of this book, which is on risks that can be managed through liquid markets. The comments that I do have on legal and reputational risk can be found in Chapters 3 and 4, particularly Sections 3.2.2 and 3.3.

5.4.2 CDO Creators

My comments for CDO creators are identical to those for subprime mortgage originators in Section 5.4.1.

5.4.3 Rating Agencies

Key risk management lessons that are reinforced by the rating agency experience leading up to the crisis are the need for a strong separation between models used for risk management and input from traders and structurers (see Section 8.4.3) and the need to have data analysis be responsive to large changes in the market environment (see Section 8.2.8.2).

5.4.4 Investors

The key risk management lesson we can draw from the experience of investors is the need for extreme skepticism in looking at marketing claims that you are getting superior returns without taking on additional risk, particularly when your ability to exit trades is limited by illiquidity. This point is discussed at greater length when looking at the experience of insurers in Section 5.4.6—everything said there can be applied here.

5.4.5 Investment Banks

The numbering of these subsections has been designed to correspond to the related discussion in Sections 5.2.5.1 through 5.2.5.9.

5.4.5.1 Reliance on Inadequate Derivatives Protection The tools needed to analyze the risk of uncollateralized and weakly collateralized derivatives protection were well known both in the academic literature and in common practice well before these transactions were booked. First, even well-collateralized protection of illiquid transactions leaves a great deal of remaining risk, since in the event of counterparty default you may have great difficulty in finding a substitute insurance provider (see the bullet point regarding derivatives with actuarial risk in Section 14.3.3). Second, lack of collateralization or weak collateralization needs to be part of the calculation of counterparty credit risk, as emphasized throughout Section 14.3.3. Third, these trades were classical examples of wrong-way risk, since the circumstances in which the counterparties would need to make insurance payments would be major economic downturns likely to impact the creditworthiness of the counterparties themselves. These trades were also wrong-way because it was well known that the insurance firms entering into them were entering into many billions of dollars of similar trades with other investment banks; the circumstances that would cause them to have to pay on one trade were highly likely to make them pay on similar trades, and they clearly did not have the financial resources to make payments under all of these contracts (this point is further elaborated in Section 5.2.6). See Section 14.3.4 for a discussion of wrong-way risk and how to account for it in calculations of counterparty credit risk. Indeed, some of these trades border on being the types of transactions Section 14.3.4 discusses as being so wrong-way that they should be counted as offering no protection at all—look at the discussion of “end of the world” trades and extreme collateralization triggers.

5.4.5.2 Reliance on Off-Balance-Sheet Vehicles The proper risk measurement of liquidity puts is very similar to the measurement of wrong-way counterparty risk and is addressed in Section 14.3.4.

5.4.5.3 Use of Faulty CDO Models The only point on which I would fault the use of CDO models was that there was too much emphasis on fitting market input and not enough emphasis on modeling that took into account the illiquidity of certain sectors, particularly the super-senior sector. Section 8.4 addresses model risk for illiquid instruments in general, and Section 13.4 addresses this issue specifically as it relates to CDO tranches.

5.4.5.4 Reliance on External Ratings In Section 13.2.1.1 we discuss the proper use of rating agency input in the risk evaluations of a bank. Rating agency evaluations should always be used as a check on internal assessments, but never as a replacement for them. This should apply just as much to credit risk arising through securities holdings as it does to credit risk arising through traditional bank loans. By similar reasoning, risk managers should always rely on their own internal models of credit portfolio risk and not on rating agency models. The credit portfolio models developed to assess bank loans, discussed in Section 13.3, are exactly the same models that are used to evaluate CDOs, as is made clear in Section 13.4. In fact, the CDO models were simply adapted from preexisting loan portfolio models.

5.4.5.5 Overreliance on VaR Measures As is made clear in Sections 6.1.2 and 8.2.6, VaR can play a proper role in the risk management of illiquid instruments, as long as it is clearly understood that what is being represented in the VaR is a liquid proxy and that a separate analysis of the hedging risk of the illiquid instrument by the liquid proxy is an absolute necessity.

5.4.5.6 Failing to Account for the Illiquidity of Super-Senior Tranches One point that is stressed several times in this book is that risk management measures must be arrived at independently, without deference to the way accounting is done for internal business decisions or for reporting to the public. Risk managers need to confirm claims of liquidity for an instrument by looking at trading history (both purchases and sales), as emphasized in the last paragraph of Section 6.1.2. Risk calculations and stress test scenarios for super-seniors required long-term (life of the security) thinking based on the lack of a liquid market. As is emphasized in Section 8.4, illiquid assets require long-term risk measures, even when accounting principles insist on mark-to-market treatment.

5.4.5.7 Inadequate Stress Tests The key point here is that the illiquidity of the positions required longer-term stress tests of the type discussed in Sections 8.4.3, 13.3.2, and 13.4.3.

5.4.5.8 Inadequate Analysis of Statistical Hedging As discussed in detail in Sections 8.2.6 and 8.4, when dealing with illiquid instruments it is vital that risk measures and reserves utilize detailed simulations of potential hedging costs over the life of the instrument to arrive at a conservative estimate. The specific case of hedging illiquid CDO tranches with more liquid CDO tranches is discussed in Section 13.4.3.

5.4.5.9 Too Big to Fail Risk managers should be vigilant in arguing against reasoning that relies on indifference to the size of losses in the event of serious economic downturns. However, many people motivated by this reasoning will not articulate it but will look for other arguments that will disguise their true incentives. One outcome of the crisis is to recognize that the design of trader and executive compensation schemes has an important risk management component. See, for example, Turner Review (2009, 79): “In the past neither the FSA [the British bank regulator] nor bank regulators in other countries paid significant attention to remunerations structures. And within firms, little attention was paid to the implications of incentive structures for risk taking, as against the implications for firm competitiveness in the labour market and for firm profitability. In retrospect this lack of focus, by both firms and regulators, was a mistake. There is a strong *prima facie* case that inappropriate incentive structures played a role in encouraging behaviour which contributed to the financial crisis.” See also Financial Stability Forum (2008, Recommendation II.19). To what degree this is workable in practice remains to be seen. Much of the burden for proper controls on asymmetric incentives will probably rest with government legislation and regulation, as we will investigate more closely in Section 5.5.5.

5.4.6 Insurers

The key risk management lesson we can gain from the experience of AIG and the monoline insurers is similar to the lesson we can take from the experience of investors—the need for extreme caution in taking on illiquid risks in an area in which you lack expertise. This lesson is even more pointed for the insurers, since they took on levels of risk that destroyed their franchises, something that few investors did.

As we emphasized in Section 2.3 on adverse selection, risk managers should always be especially vigilant when traders are dealing in transactions for which they do not possess an informational advantage. The temptation to get involved may be great when a plausible case has been made that returns are high relative to risk, but even when this case seems overwhelming, the size of risk must be kept proportional to liquidity. The risk manager’s greatest friend is always the stop-loss limit, as discussed in Section 6.1.1. Even for risks that the firm does not understand well, a limit can be placed on tolerable losses and an exit strategy planned. But when lack of liquidity means you won’t be able to exit as losses mount, no degree of promised return should be allowed to jeopardize a firm’s franchise. If an opportunity just seems too good to pass up, then invest in the expertise to manage it knowledgeably. Relying on regulatory constraints or advisory services, such

as rating agencies, cannot be considered in any way a substitute for this expertise when illiquidity prevents easy exit.

5.4.7 Credit Contagion

Most of the ideas for reducing the risk of credit contagion are being addressed at the regulatory level and are covered in Section 5.5.7. Chapter 14 addresses counterparty credit risk management at the level of the firm.

5.4.8 Market Contagion

The 2007–2008 experience on the degree to which market illiquidity spread and the length and depth of this illiquidity will need to impact the historical measures of VaR risk (Section 7.1) and the severity of stress tests (Section 7.2) that will be utilized going forward.

5.5 LESSONS FROM THE CRISIS FOR REGULATORS

The numbering of subsections, as in Section 5.4, is designed to allow easy reference to the discussion of the mechanics of the crisis in Sections 5.2 and 5.3. Sections 5.5.1 through 5.5.6 correspond to Sections 5.2.1 through 5.2.6, respectively, while Section 5.5.7 corresponds to Section 5.3.1, and Section 5.5.8 corresponds to Section 5.3.2.

Throughout this section, I have relied as much as possible on recommendations from the Financial Stability Board and its predecessor organization, the Financial Stability Forum. This organization is a joint effort of finance ministers and central bankers from the G-20 countries, as well as major global public institutions such as the International Monetary Fund, the World Bank, and the Bank for International Settlements. It was established to coordinate financial regulation and standards setting globally. Many of its recommendations carry the endorsement of the G-20, the group of 20 leading economies that account for over 80 percent of global gross domestic product (GDP). The G-20 fosters cooperation and consultation on matters relating to the international financial system. As such, I believe it represents the broadest consensus views of the regulatory community. The Financial Stability Forum's 2008 recommendations for enhancing market and institutional stability will be referred to as FSF (2008). I will bring in views of other regulatory bodies and of academics where there are significant disagreements or where a particular document has expressed a view with particular clarity. Two sources that this section utilizes often are the Turner Review of 2009, a broad review of regulatory policy authorized

by the British government, and the 2009 report on financial reform by the Group of Thirty, the same private, nonprofit group of leading representatives of the international business, regulatory, and academic communities whose influential report on derivatives risks I make heavy use of in Section 6.1.1.

5.5.1 Mortgage Originators

There has not been as much focus on mortgage originators in the recommendations arising from the crisis as there has been on CDO creators and rating agencies. Davidson (2007) does have a persuasive suggestion: “There needs to be capital at the origination end of the process. Without capital, representations and warranties have no value. To achieve this, brokers (or whoever has direct contact with the borrower) should be licensed and bonded and firms in the chain of reps and warrants need to maintain sufficient reserves to support their financial promises. This capital would be available to assess damages in the case of fraudulent or predatory practices that hurt borrowers and homeowners.”

One other interesting recommendation is to restructure mortgages to avoid the impact of negative home equity on homeowner defaults. Shiller (2008, Chapter 6) has several interesting suggestions along this line, including real estate derivatives, home equity insurance, and continuous-workout mortgages. A mortgage market that builds in this protection exists in Denmark. George Soros, in a *Wall Street Journal* article on October 10, 2008, explained the Danish mortgages as follows: “Every mortgage is instantly converted into a security of the same amount and the two remain interchangeable at all times. Homeowners can retire mortgages not only by paying them off, but also by buying an equivalent face amount of bonds at market price. Because the value of homes and the associated mortgage bonds tend to move in the same direction, homeowners should not end up with negative equity in their homes. To state it more clearly, as home prices decline, the amount that a homeowner must spend to retire his mortgage decreases because he can buy the bonds at lower prices.”

5.5.2 CDO Creators

Recommendations have focused on trying to better align incentives between CDO creators and investors. For example, the Group of Thirty (2009, Recommendation 13) states that “regulators should require regulated financial institutions to retain a meaningful portion of the credit risk they are packaging into securitized and other structured credit products.” A requirement of this type is a part of both the Dodd-Frank legislation in the United States and rules proposed by European authorities (see Global Legal Group

2011, Chapter 3: “EU and US Securitization Risk Retention and Disclosure Rules—A Comparison”). All such proposals face three large challenges:

1. How to measure credit risk retention given all of the ways that are now available for offsetting risk through credit derivatives, including the use of credit indexes.
2. How to avoid the situation discussed in Section 5.2.2 in which CDO creators viewed the total package as so lucrative that they could regard the retained equity as a “free good” to whose credit performance they were indifferent. Proposals to deal with this are combinations of raising the portion of risk retained and of requiring that a portion of risk be retained in all tranches sold, not just in a single tranche whose losses might not be well correlated with the tranches sold.
3. How to avoid making retention requirements so onerous that they discourage or raise the costs of securitization that is considered beneficial to the general public (e.g., homeowners).

5.5.3 Rating Agencies

FSF (2008, Section IV) contains several proposals to deal with the rating agency issues that were raised by the crisis. I would highlight the following:

- Rating agencies “should clearly differentiate, either with a different rating scale or with additional symbols, the ratings used for structured products from those for corporate bonds.” This would clarify the greater reliance of structured product ratings on models and economic assumptions and their “potential for significantly higher ratings volatility.”
- Rating agencies “should enhance their review of the quality of the data input and of the due diligence performed on underlying assets by originators, arrangers and issuers involved in structured products.”
- Regulatory authorities should “review their use of ratings in the regulatory and supervisory framework” to encourage investors to “make independent judgment of risks and perform their own due diligence” and reduce “uncritical reliance on credit ratings as a substitute for that independent evaluation.” “Investor associations should consider developing standards of due diligence and credit analysis for investing in structured products.”
- Rating agencies should revise their codes of conduct to better deal with conflict of interest issues and “demonstrate that they have the ability to maintain the quality of their service in the face of rapid expansion of their activities.”
- Rating agencies “should disclose past ratings in a more systematic way, and improve the comparability of their track records.”

Many of these points are echoed in the Group of Thirty (2009, Recommendation 14) and Richardson and White (2009). Richardson and White also suggest as an alternative that financial regulations be changed to deemphasize the use of rating agencies. In this approach, “regulated financial institutions would thus be free to take advice from sources that they considered to be most reliable—based on the track record of the advisor, the business model of the advisor (including the possibilities of conflicts of interest), the other activities of the advisor (which might pose potential conflicts), and anything else the institution considered relevant.” But “the institution would have to justify its choice of advisor to its regulator.” This alternative could lead to more competition and new approaches in the ratings advisory market.

5.5.4 Investors

Most regulatory discussion concerning investors has been tied to protecting investors from CDO creators and rating agencies or to controlling the spread of financial crises through credit contagion or market contagion. These issues are addressed in other sections—5.5.2 for CDO creators, 5.5.3 for rating agencies, 5.5.7 for credit contagion, and 5.5.8 for market contagion.

One attempt to address regulatory concerns for investors directly is in the Squam Lake Report (Squam Lake Group 2010, Chapter 4: “Regulation of Retirement Savings”). Given that many of the investors who had losses in the 2007–2008 disaster were pension funds reaching for excess returns in exchange for risks that may have been very poorly understood by the employees who were the ultimate recipients of these losses, this is a timely concern. Among the recommendations in this chapter are:

- Requiring simple standardized disclosure for products offered in defined contribution retirement plans.
- Requiring simple and meaningful standardized disclosure of measures of long-term risk and of investment costs. Any advertisement of average prior returns should also include a standardized measure of uncertainty.
- The standard part of a defined contribution plan should be restricted to well-diversified products with low fees.

5.5.5 Investment Banks

The regulatory responses to the default and near default of several major investment banks in the 2007–2008 crisis can be roughly divided into four categories. The first consists of measures to require tightening of internal risk management procedures, combined with greater scrutiny of these

procedures by regulatory authorities. The second is a new focus on regulatory oversight of compensation policy. The third is significantly higher requirements for bank capital to serve as a buffer against losses before they impact depositors, and therefore governments and taxpayers. And the fourth consists of proposed restrictions on the size and range of activities of investment banks. We consider each in turn.

5.5.5.1 Tightened Internal Risk Management Procedures It is natural for part of the regulatory response to the crisis to be to call for stronger internal risk management procedures within investment banks. But without either very specific guidance on how procedures should change or new ongoing regulatory scrutiny to ensure compliance, this will just be empty exhortation.

This book addresses specific regulatory guidance on particular risk management issues at those points where it is most relevant: new guidance on model review in Chapter 8; new guidance on oversight of compensation policy later in this section; new guidance on stress tests as part of the discussion of increased capital later in this section; new guidance on counterparty risk in Section 5.5.7 on stemming credit contagion; new guidance on valuation in Section 5.5.8 on stemming market contagion. This subsection addresses proposals for increased regulatory scrutiny to ensure compliance.

The most comprehensive set of recommendations for changes in regulatory scrutiny of investment bank internal controls is in the Group of Thirty (2009) Recommendations 1, 2b, 6, 7, and 8. Some of the key points in these recommendations are:

- At a national level, countries should “eliminate unnecessary overlaps and gaps in coverage . . . removing the potential for regulatory arbitrage, and improving regulatory coordination.” The Squam Lake Report (Squam Lake Group 2010, Chapter 2) goes further, calling for a single regulatory authority in each country to be “responsible for overseeing the health and stability of the overall financial system.”
- At the international level, national regulatory authorities should “better coordinate oversight of the largest international banking organizations” and “move beyond coordinated rule making and standard setting” to convergence in application and enforcement of standards and closing of regulatory gaps.
- In countries where the central bank is not the primary regulator of banks, the central bank needs to become more involved in regulation, particularly with regard to the largest systemically significant firms and critical payment and clearing systems.

Saunders, Smith, and Walter (2009) call for a dedicated regulator in each country for “large complex financial institutions (LCFIs),” arguing that these firms are different in character and pose a greater threat to the global financial system than smaller and more specialized firms. “Most importantly, the regulator would have the power and the obligation to ensure that LCFIs operate consistently with priority attention to the institution’s safety and soundness, even if this can only be achieved at the cost of reduced growth and profitability.”

5.5.5.2 Compensation Policy The quotation from the Turner Review (2009) in Section 5.4.5.9 indicates the change in regulatory attitude toward the compensation structures of investment banks. The 2007–2008 crisis has led regulatory authorities to switch from regarding compensation as purely an internal matter for banks to regarding it as a key component of risk control. The Financial Stability Board issued a separate report on sound compensation practices, FSF (2009b), which states that the “perverse incentives” of generous bonus payments for high short-term profits “without adequate regard to the longer term risks they imposed on their firms” “amplified the excessive risk-taking that severely threatened the global financial system and left firms with fewer resources to absorb losses as risks materialized.”

Some key principles elucidated in FSF (2009b) are:

- Compensation must take into account both profit generated and risk entailed.
- The firm’s board of directors must actively oversee the design and operation of compensation policy and must ensure that the compensation policy addresses the balance between profit and risk.
- “Compensation must be adjusted for all types of risk,” including difficult-to-measure risks such as liquidity risk and reputational risk. This necessitates that both quantitative measures and human judgment play a role in determining risk adjustments. (This is consistent with this book’s emphasis on the need for subjective judgment in risk management; see Sections 1.3 and 6.1.1.)
- Compensation outcomes should be symmetric with risk outcomes, with bonuses diminishing or disappearing in the event of poor firm, divisional, or business unit performance.
- “Compensation payout schedules must be sensitive to the time horizon of risks,” with compensation deferred when risks are realized over long periods. (This is consistent with the distinction made on the differing risk management approaches for liquid and illiquid positions in Section 1.2 and 6.1.1. The mix of cash and equity in compensation also needs to be consistent with the nature and time horizon of risks generated.)

- “Firms should disclose clear, consistent and timely information about their compensation practices” to make sure that all stakeholders, including customers, creditors, and regulators as well as stockholders, can make informed decisions.
- Regulators must include review of compensation practices as part of their supervisory role and be prepared to take prompt action when compensation practices are deemed deficient. “Compensation is an incentive system, not simply a market wage” and so must be subject to regulatory review. Given the competitive nature of the labor market for financial institutions, “Market participants are pessimistic about the effectiveness of change unless it is industry-wide and global. . . . Changing compensation practice will be challenging, time-consuming and involve material costs. Therefore, in the absence of sustained external pressure, firms may fail to carry through on originally good intentions. Although some market participants are wary of regulatory pressure, many believe that a widespread change in practice can be achieved only with the help of supervisory and regulatory agencies, which should coordinate at the global level.”

Other regulatory publications in response to the crisis are quite consistent with FSF (2009b). See, for example, Turner Review (2009, Section 2.5 (ii)). Clementi et al. (2009) provide an academic analysis very supportive of this approach.

The FSF (2009b) compensation proposals are primarily aimed at those directly involved in the creation and management of risk positions—traders, marketers, and structurers. The Squam Lake Report (Squam Lake Group 2010, Chapter 6) suggests an interesting approach aimed at the compensation of senior management of financial institutions. Since “governments will bail out financial firms during a crisis,” “the stakeholders in financial firms—executives, creditors, and shareholders—do not face the full cost of their failure. This in turn increases the likelihood of bank failures, the potential for systematic risk, and expected taxpayer costs.” Along with other measures, a “mechanism for inducing financial firms to internalize the costs of their actions” would be holdbacks of a fixed dollar amount of compensation that would be forfeited “if the firm goes bankrupt or receives extraordinary government assistance” over some defined future time period. Some further points raised about this proposal are:

- “More familiar forms of deferred compensation, such as stock awards and options” help align manager incentives with stockholders’ interests but do not align them with taxpayers’ interests.
- “Resignation from the firm should not accelerate payment of an employee’s holdbacks,” since this would “weaken their concern about the

long-term consequences of their actions. . . . In the same spirit, managers should not be rewarded for taking their firms into bankruptcy. If a firm declares bankruptcy, its managers should receive their holdbacks only after its other creditors have been made whole.”

- “[D]eferred compensation leans against management’s incentive to pursue risky strategies that might result in government bailouts. Similarly, rather than wait for a bailout during a financial crisis, the management of a troubled firm would have a powerful incentive to find a private solution, perhaps by boosting the firm’s liquidity to prevent a run, raising new capital, or facilitating a takeover by another firm.”

Rajan (2010) Chapter 7 and the section on “Reducing the Search for Tail Risk” in Chapter 8 offer strong supporting arguments for both the FSF and the Squam Lake proposals.

5.5.5.3 Capital Requirements Regulators have certainly taken many steps since the crisis to raise the levels of capital required. They have taken steps both to increase the level of risk-weighted assets that will be calculated against trading positions and to increase the capital required for a given level of risk-weighted assets. A good summary of the steps taken by the BIS is PricewaterhouseCoopers (2011), in which Chapter 4 covers increases in risk-weighted asset calculations and Chapter 3 covers increases in capital.

While the direction of the regulatory response is clearly correct, the specifics of the approach are troubling. Though regulators have enhanced stress-testing requirements (PricewaterhouseCoopers 2011, Sections 11.1 and 11.3), there is no direct tie between these stress tests and capital requirements. It is true that there is now a stressed VaR calculation that impacts capital (PricewaterhouseCoopers 2011, Sections 4.6.3.3 and 11.2.3), but this just stresses VaR parameters, a form of stress testing that has been shown to be inadequate (see Section 7.2.1 and the discussion of the use of stress tests by Long-Term Capital Management (LTCM) in Section 4.2.1; as I state there, LTCM “did run stress versions of VaR based on a higher than historical level of correlations, but it is doubtful that this offers the same degree of conservatism as a set of fully worked-through scenarios”).

What I find most troubling is that the degree of complexity of capital computations has grown to the point that there is great danger of risk managers and regulators losing sight of the largest risk exposures through the distraction of enlarged reporting requirements. PricewaterhouseCoopers (2011) expresses a similar apprehension in Section 4.7: “A particular area of concern is the introduction of many systems within the market risk process. Previously market risk departments have been reliant on one regulatory risk system, but they can now have up to five systems to manage. This in itself

is likely to increase the operational risk associated with market risk. Even though these measures are being introduced to ensure more comprehensive measurement, their complexity may cause banks to miss positions and, as always, there will be loopholes in the systems, harder to find but also harder to catch.”

By contrast, capital requirements directly tied to stress-test scenarios would focus management and regulatory attention in exactly the right place: on the impact of large moves in major economic variables, exactly the types of events that have led to the events that challenge the health of financial firms and of the financial system. It is this type of event, significant drops in price of important asset classes, for which capital cushions are needed. My arguments supporting this approach can be found in Sections 7.2.2 and 7.3, particularly toward the end of 7.3 where I discuss the reasons that Chase Manhattan moved to basing internal capital requirements on stress tests in the late 1990s.

There may be two objections to basing regulatory capital on stress-test scenarios. The first is that the limited number of individually tailored scenarios that can be considered may allow some risks that avoid attracting capital. This can be dealt with either by having some part of the capital requirement based on VaR or by utilizing statistically driven stress tests as supplements to individually tailored ones, as described in Section 7.2.3. The second objection is the inevitable subjectivity of stress-test scenarios. Some element of subjectivity is unavoidable and, in fact, welcome, as emphasized in Sections 6.1.1 and 7.2.2. U.S. and European regulators have had no trouble specifying stress-test scenarios in the capital adequacy tests mandated in the wake of the crisis (see, for example, Federal Reserve Board 2009). I advocate utilizing these crisis tests as a precedent for an ongoing process, for the following reasons:

- Whatever level of possible stress market move corresponds to the capital requirement, there will always be some possibility that an extreme market move will exceed this level and require some absorption of loss by taxpayers on behalf of depositors. Since it is the regulatory authorities that represent the taxpayers’ interests, they should be the ones to determine the level of protection. There are inevitable trade-offs between capital requirements that are too high and hurt economic activity and capital requirements that are too low and create too high a risk of potential crisis. It is the regulatory authorities acting on behalf of government that should be weighing these consequences and deciding on the correct balance.
- Regulatory authorities could signal their willingness to support certain markets in a liquidity crunch by differentiating by instrument between

the time periods over which stress tests need to be run. For example, a two-week stress event might be considered adequate for government bond and spot foreign exchange markets, signaling government readiness to intervene to quickly restore liquidity in these markets, but a three-month stress event might be required for structured securities in which the government wished to indicate less urgency to intervene to restore liquidity.

- The risk managers of a firm should possess specialized knowledge regarding the trading positions and activities of that firm. There is no reason to think they would possess any specialized expertise about the probability of macroeconomic events, such as large moves in a stock indexes, government bond rates, or housing prices. So I do not see any comparative advantage argument in favor of having these stress levels be set by firm risk management as opposed to government regulators.
- When firm risk managers set the stress levels, there is an inevitable competitive pressure to set levels lower to free up capital and improve returns. A common level set by regulatory authorities would eliminate the competitive advantage a firm could get by hiring more optimistic risk managers. If political pressures prevent regulatory authorities from setting these levels on a regular basis, I would urge financial institutions to seek a way that a common level could be set by an industry association.

5.5.5.4 Limitations on Size and Allowable Activities The other regulatory proposals considered in this section, regarding tightened risk management procedures, capital requirements, and compensation, do not seek to fundamentally change the structure of the financial industry. Some suggested actions do address fundamental structure directly, trying to eliminate a “too big to fail” mentality either by placing limits on the size of financial firms or by creating a strong separation between firms that can engage in certain types of activities and firms that receive any kind of government support.

No proposals of this type were part of the FSF (2008) recommendations, and the Turner Review explicitly rejected proposals for separation of activities, stating in Section 2.9 that “It does not therefore seem practical to work on the assumption that we can or should achieve the complete institutional separation of ‘utility banks’ from ‘investment banks.’ . . . Large complex banks spanning a wide range of activities are likely to remain a feature of the world’s financial system.” Points in support of this view offered by the Turner Review (2009) are:

- A reimposition of Glass-Steagall type separation between commercial and investment banking is impractical, given that many activities that used to be conducted solely by investment banks, such as the

underwriting of corporate bonds, are now “core elements within an integrated service to corporate customers in a world where a significant element of debt is securitized.”

- Many so-called narrow banks that focused almost entirely on traditional commercial and retail banking activities, such as Northern Rock, Washington Mutual, and IndyMac, also failed during the crisis.
- The international integration of financial markets would make it difficult to achieve such a separation without a broad consensus among governments, which is unlikely to be achieved.

Another point in support of this view comes from Rajan (2010, 173): “Proprietary trading . . . is another activity that has come in for censure. . . . Critics argue that proprietary trading is risky. It is hard to see this as an important cause of the crisis: banks did not get into trouble because of large losses made on trading positions. They failed because they held mortgage-backed securities to maturity, not because they traded them.” Rajan’s analysis of the causes of bank failure in the crisis is certainly supported by Section 5.2.5 of this book.

A major proponent of at least considering fundamental changes to industry structure is the Group of Thirty (2009), which in its Recommendation 1 proposes that “Large, systemically important banking institutions should be restricted in undertaking proprietary activities that present high risks and serious conflict of interest” and states that “nation-wide limits on deposit concentration should be considered.” It is perhaps not coincidental that the steering committee for this report was chaired by Paul Volcker, whose “Volcker rule” (see McLean and Nocera 2010, 366) for the ban on much proprietary trading activity by deposit-taking banks has been one of the principal legislative efforts in this direction. In support of its proposal for restrictions on proprietary trading, the Group of Thirty report states that “What is at issue is the extent to which these approaches can sensibly be combined in a single institution, and particularly in those highly protected banking institutions at the core of the financial system. Almost inevitably, the complexity of much proprietary capital market activity, and the perceived need for confidentiality in such activities, limits transparency for investors and creditors alike. In concept, the risks involved might be reduced by limiting leverage and attaching high capital standards and exceptionally close supervision. Some members of the G30 feel such an approach could be sufficient to deal with these risks. . . . Experience demonstrates that under stress, capital and credit resources will be diverted to cover losses, weakening protection of client interests. . . . Moreover, to the extent that these proprietary activities are carried out by firms supervised by government and protected

from the full force of potential failure, there is a strong element of unfair competition with ‘free-standing’ institutions.”

Other proponents of limits on industry structure are Roubini and Mihm (2011), who advocate both limits on size (223–230) and a reimposition of a (greatly expanded) Glass-Steagall (230–233), and Stiglitz (2010, 164–168), who quotes former Bank of England governor Mervyn King: “If some banks are thought to be too big to fail . . . then they are too big.”

Rajan (2010, 169–176) provides a very incisive analysis of these proposals. I would highly recommend this to anyone interested in this topic. While Rajan is skeptical of most of the value of most of these suggestions, he is sympathetic to the idea of limiting proprietary trading, not because it will reduce risk of bank failure, but because of the inherent conflict of interest between banks’ proprietary trading and the interests of their customers. Rajan argues that “Banks that are involved in many businesses obtain an enormous amount of private information from them. This information should be used to help clients, not trade against them.” But Rajan does clarify that he supports *limiting* bank proprietary trading, not eliminating it, because “some legitimate activities, including hedging and market making, could be hard to distinguish from proprietary trading.” My own experience supports Rajan on this point; see my account of market making in Section 9.1.

5.5.6 Insurers

FSF (2008, Recommendation II.8) calls for insurance regulators to strengthen the regulatory and capital framework for monoline insurers in relation to structured credit.

5.5.7 Credit Contagion

In response to the large role that counterparty credit risk on over-the-counter (OTC) derivatives played in credit contagion in the 2007–2008 crisis, it is natural that a major focus of regulatory concern has been to attempt to minimize future use of OTC derivatives and increase the use of exchange-traded derivatives. Section 14.2 will review many of the advantages that exchange-traded derivatives have relative to OTC derivatives in minimizing credit exposure:

- The elimination of credit exposure between counterparties, with all credit exposure centralized with the exchange (or associated clearinghouse).
- The relatively automatic mechanisms for margining, posting of collateral, and closing out of positions that minimize the credit exposure of the exchange.

- The mutualized sharing of the residual counterparty risk among all members of the exchange.
- The ease with which counterparties can extinguish existing positions, reducing credit exposure levels.
- The greater transparency and information-sharing that are encouraged by the exchange's lack of any market exposure.

The Squam Lake Report (Squam Lake Group 2010, Chapter 9) does an excellent job of laying out these arguments concisely in the context of reducing the risk of credit contagion in a crisis. A particular point the Squam Lake Report raises relative to crises is that the ease with which counterparties can extinguish existing positions also reduces demand for collateral, “a precious resource, especially during a financial crisis.”

In response to these arguments, regulatory bodies have been highly motivated to push regulated institutions in the direction of reducing their use of OTC derivatives relative to exchange-traded derivatives. While almost all observers agree that this is a move in the right direction, some cautions have been sounded on two grounds: (1) In the process, some of the advantages to customers of OTC derivatives relative to exchange-traded derivatives, detailed in Section 14.3, will be lost, and (2) as more trading volume is funneled to exchanges, the exchanges may grow to the point where they will become a potential source of systemic risk that could trigger or exacerbate a crisis.

Before looking at these warnings, let's first summarize the actions being contemplated. In November 2009, the G-20 summit issued a recommendation that “All standardised OTC derivative contracts should be traded on exchanges or electronic trading platforms, where appropriate, and cleared through central counterparties by end-2012 at the latest” (see Financial Stability Board 2010). Obviously much of the force of this recommendation will turn on exactly how the word “standardised” is interpreted. In addition to those contracts that are being mandated to be traded on exchanges, powerful incentives are being put in place to encourage the replacement of OTC derivatives by exchange-traded derivatives, by mandating more stringent capital requirements on OTC derivatives. These actions are covered in PricewaterhouseCoopers (2011, Chapter 5). Section 5.3.1.7 of PricewaterhouseCoopers explains that “new rules provide banks with strong incentives to move trades to a central counterparty clearing house (‘CCP’) with exposures to CCPs assigned fairly low risk weights. To complement this, the [Basel] Committee supports enhanced capital standards and rigorous risk management for CCPs. It has therefore specified that the favourable treatment of exposures to CCPs applies only where the CCP complies” with regulatory standards.

Now let's turn to possible objections. The first is the possible loss of the advantages of OTC derivatives over exchange-traded derivatives for some contracts. As detailed in Section 14.3, these are principally the ability to more closely customize OTC derivatives to client needs, less stringent operational requirements, and the willingness of OTC market makers to extend credit beyond what exchanges offer, along with occasional restrictions on trading that disadvantage some customers, mentioned in Section 14.2. How these concerns will be dealt with depends very much on implementation. For example, if the term "standardised" in the G-20 recommendation of the previous paragraph is interpreted narrowly, it will not hamper customization much, but will leave a substantial portion of OTC derivatives outside clearinghouses.

The second possible objection is that concentrating more derivatives trading in exchanges will increase the risk that the exchanges themselves will become a potential source of systemic risk. The clearest exposition of this argument is contained in Pirrong (2011). While exchanges have well-developed mechanisms for containing credit risk, these are not perfect. As explained in Section 14.2, exchanges are exposed to counterparty risk in between margin calls, and their protection against this is much the same type of VaR and stress-test calculations that have failed to prevent banks from being a source of systemic risk. While exchanges have avoided exposure to the illiquid instruments that have frequently been the source of problems for banks, this has been achieved by limiting exchange trading to the most liquid contracts; the price of concentrating more derivatives trading in exchanges may be to expose exchanges to more illiquid instruments.

A balanced approach to the trade-off between reduction of credit risk on OTC derivatives and avoiding the potential for systemic risk at exchanges is the Federal Reserve Bank of New York staff report, Duffie, Li, and Lubke (2010). While calling for measures that will increase the use of exchange trading for more liquid derivatives, a number of measures short of exchange trading are proposed to reduce the systemic risk of less liquid OTC derivatives. These include:

- Increased capital requirements reflecting not just a bank's exposure to counterparty default but also "the risks that it imposes on others" by its own risk of default.
- Increased public transparency of aggregate price and volume information and "going prices," closer to the level of transparency available for exchange-traded derivatives.
- Aggressive trade compression, along the lines discussed in Section 14.3.5 of this book.

One further area that regulators have considered for containing credit contagion is regulation of money market funds. The Group of Thirty (2009, Recommendation 3) calls for “Money market mutual funds wishing to continue to offer bank-like services, such as transaction-account services, withdrawal on demand at par, and assurances of maintaining a net asset value (NAV) at par . . . to reorganize as special-purpose banks, with appropriate prudential regulation and supervision, government insurance, and access to central bank lender-of-last-resort facilities.” Any money market fund not willing to subject itself to these requirements would not be permitted to offer “explicit or implicit assurances to investors that funds can be withdrawn on demand at a stable NAV.”

5.5.8 Market Contagion

Three types of measures have been proposed to limit the spread of problems for any one firm to other firms through market contagion. The first is to limit the pressures on financial firms facing difficulties to quickly shrink balance sheets, thereby reducing downward pressure on markets from distressed selling. These measures are classified as ones to reduce procyclicality. The second is to provide for a more orderly process for placing a firm in bankruptcy, allowing more time for positions to be unwound. The third is to provide regulatory oversight for financial entities that might be impacted by financial contagion, to provide regulators with greater knowledge about positions that could be impacted through market contagion. We’ll consider each in turn.

5.5.8.1 Reducing Procyclicality The primary regulatory effort in this direction has been to require capital buffers that should be built up in periods of good profitability and drawn down in periods of stress. By having some portion of required capital that it is permissible to draw upon in a crisis, the intention is to relieve the pressure on banks to sell off assets in response to a sharp fall in market valuation. FSF (2009a, Section III) calls for “the capital framework . . . [to] be enhanced to provide stronger capital buffers during strong economic conditions that can be drawn down to a credible minimum requirement during periods of economic and financial stress.” Group of Thirty (2009, Recommendation 10) calls for mandated capital ratios to “be expressed as a broad range . . . with the expectation that as part of supervisory guidance, firms will operate in the upper end of such a range in periods when the market is exuberant and tendencies for underestimating and underpricing risk are great.” These recommendations have been acted on by the Basel regulators through requirements for capital buffers that can be drawn down during periods of economic stress. Details

of these requirements can be found in PricewaterhouseCoopers (2011, Sections 10.3.3 and 10.3.4).

While capital buffers have been the focus of the regulatory response to procyclicality, some thought has also been given to reducing the cyclicality of accounting rules. With regard to provisions for loan losses, the Group of Thirty (2009, Recommendation 12(c)) calls for accounting principles that are “more flexible in regard to the prudential need for regulated institutions to maintain adequate credit-loss reserves sufficient to cover expected losses across their portfolios over the life of the assets in those portfolios,” while maintaining transparent disclosure of reserve methodology. This recommendation runs counter to much of the past decade’s tendencies in accounting for loan loss provisions, which have emphasized provisioning only when loss potential on specific loans starts to become apparent (the “incurred loss” model). FSF (2009a, Section IV) also recommends reconsideration of the “incurred loss model by analyzing alternative approaches for recognizing and measuring loan losses that incorporate a broader range of available credit information.” The FSF states that such alternative approaches might have identified loan losses earlier in the credit cycle and potentially reduced procyclicality. The Basel regulators have begun promoting a longer-run approach toward accounting for loan loss provisions, based on long-term data series for default probabilities and historically conservative assumptions for loss given default. These actions are detailed in PricewaterhouseCoopers (2011, Sections 10.3.1 and 10.3.2).

It would be consistent with this longer-term approach to loan-loss provisioning to move to a longer-term approach to valuation of illiquid securities and derivative positions. This would have the same impact of building up reserves during buoyant markets that would reduce the pressure to liquidate assets in times of stress (for a more detailed discussion, see Section 8.4.4). The Group of Thirty (2009, Recommendation 12) calls for a move in this direction. In the supporting discussion, the Group of Thirty argues for “more realistic guidelines for addressing valuation issues for illiquid investments.” FSF (2008, Section III.3) also contemplates changes in this direction.

Zandi (2009, 258–259) makes a similar suggestion: that to keep banks’ survival from being threatened in financial crises, “mark-to-market accounting rules could be tweaked most importantly for securities that financial institutions don’t ever plan on selling. . . . It is reasonable for institutions to value these securities based on expectations of any losses they might eventually suffer, but it isn’t reasonable to value these securities using prices they would get if they sold them today.” Where my proposal would differ from Zandi’s is that my criteria would be liquidity and not intention of sale, and for illiquid securities I would replace the prices at which they could be sold today with *very* conservative estimates of losses, as opposed to expected

losses. I believe large reserves are needed against illiquid instruments, but conservatism should make for relatively stable reserve levels that would only rarely need to be increased in a crisis.

5.5.8.2 More Orderly Bankruptcy The Group of Thirty (2009, Recommendation 16) calls for legislation to give regulators greater authority to provide for “orderly closings of regulated banking organizations, and other systemically significant regulated financial institutions.” The reasoning behind this recommendation states that “Market discipline works best in a system in which failures can happen without being a source of major disruption and contagion.” “To be fully effective, the legal regimes that operate once a failure is triggered should be modified, with a view to placing primary importance on the capacity of the authorities to take actions to protect the health of the system.”

PricewaterhouseCoopers (2011, Chapter 15) provides a summary of actions that have been taken by international regulators along these lines, particularly with regard to requiring each large financial institution to prepare a “resolution plan” for the firm’s orderly liquidation in the event of insolvency. The Squam Lake Report (Squam Lake Group 2010, Chapter 8) provides specific recommendations for preparing resolution plans. Huertas (2011, Chapter 7) provides detailed analysis by a senior member of Britain’s financial regulatory agency of how to improve the resolution process.

5.5.8.3 Broader Regulatory Oversight The Group of Thirty (2009, Recommendation 4) calls for “managers of private pools of capital that employ substantial borrowed funds” (i.e., hedge funds and private equity funds) above some minimum size to register with and provide periodic reports to banking regulators. These reports should include information on “size, investment style, borrowing, and performance of the funds under management.” The regulators should also “have authority to establish appropriate standards for capital, liquidity, and risk management” for those funds “above a size judged to be potentially systemically significant.” This is recognized as being a clear break from the prevailing approach to fund regulation, which has primarily focused on regulation of lenders to hedge funds, an approach that has been justified by the fact that hedge funds do not employ any sort of government guarantee, whereas their creditors do. But the Group of Thirty notes that “the increased emphasis on financial stability in the mandates” of regulators “points to the need for greater, more systemic access to information crucial to understanding the growing risk imbalances in the system.” Strong academic arguments supporting the approach of this recommendation have been supplied by Lo (2008).

5.6 BROADER LESSONS FROM THE CRISIS

When as widely respected a figure in the financial markets as Paul Volcker is moved to say that the single most important contribution of the financial industry in the past 25 years was the automatic teller machine, which at least had proven useful, there is something wrong with the industry that needs to be addressed at levels beyond risk managers and government regulators. (Volcker made this remark addressing an audience of senior finance industry figures in December 2009. It was widely reported—for example, in a *Daily Telegraph* article by Louise Armitstead on December 9.)

I don't doubt that comments like Volcker's overstate the case—many of the innovations in markets, derivatives, and securitization of the past 25 years have genuinely made easier financing, broader investment opportunities, and valuable risk management tools available to firms and people who were worthy recipients; good narratives of these advances, from different perspectives, can be found in Shiller (2003) and Brown (2012). But clearly many reasonable people are starting to feel there is an imbalance—too many innovations that just provide tax and accounting gimmicks or introduce unnecessary complications relative to too few innovations addressing real economic issues. Some suggestive ideas for new directions are:

- The prominent economist Robert Shiller has been focusing on the question of identifying financial innovations that will more closely match genuine social needs (in his words, address “risks that really matter”). Shiller (2008) gives a brief account of these ideas in the context of the 2007–2008 crisis; Shiller (2003) provides a more thorough explication. Some of the innovations he advocates would be ways of hedging the cost of housing, providing home equity insurance, being able to insure against the economic risk of career choice, and hedging against the economic performance of a country.
- Richard Bookstaber, an experienced risk manager, in Bookstaber (2007) advocates a redesign of financial products in the direction of greater simplicity and greater tolerance for survival of disruptions.

Managing Financial Risk

The management of financial risk can be divided into two parts: risk measurement and risk control. In general, the industry agrees more on how risk should be measured than on how it should be controlled.

6.1 RISK MEASUREMENT

6.1.1 General Principles

As stated in Chapter 1, the key characteristic that distinguishes financial risk management from other types of risk management is that financial risk management can take advantage of liquid markets as part of a risk management strategy. In this chapter we examine the structure of financial risk management in more detail, and a good starting point is to consider the hypothetical case in which a market is so liquid that any position can be liquidated instantaneously. While this is obviously an extreme that does not exist in reality, it will still provide an instructive background against which to consider more realistic cases.

With such perfect liquidity, risk management could, in principle, just consist of setting loss limits for each trader and each trading group (the industry jargon for this is a *stop-loss limit*). As soon as a trader reached the limit for a position, the entire position could be liquidated with no further loss. Or if management decided that its risk tolerance had changed because of changes in their view of the economy or the institutional environment, positions could be liquidated with no further losses. Even in such an extreme case, the following rules would be needed.

- **Careful and continuous tracking of market prices of existing positions.** Otherwise, you would not know when a trader was through a stop-loss limit. Traders may be tempted to hide the size of their losses, knowing that being through a limit will cause the position to be closed out and

eliminate their chance of making future gains. An optimistic mark of the position could delay the recognition of losses. And traders who know they are through a limit when management does not are very dangerous, since they will be tempted to swing for the fences (as discussed in Section 2.1). So, no matter how liquid the market, correct and independent valuation of current positions is at the heart of all good risk management.

- **Sensible choices of limit size relative to trader expertise and trading strategy.** A good example would be a position being taken that will benefit from a policy change, such as the lifting of governmental foreign exchange (FX) controls. Such positions often have predictable daily losses for as long as the current policy remains in place but have a large profit potential if the policy is changed. If management is convinced that this is a sensible gamble, or has sufficient trust in a given trader's judgment to allow her to make that decision, it would be self-defeating to implement it with a very small stop-loss limit that would, with high probability, cause the position to be closed out before the policy change occurs. Positions that require patience to make money should be undertaken only if the firm has the risk appetite to allow for that patience.
- **Good procedures for review of request to exceed limits.** When a trader reaches (or is approaching) a stop-loss limit, there is an excellent chance that he will want to make a case to his management for a temporary expansion of the limit. He may believe that a market shift in his favor is "just around the corner." A strict and firm policy to close out all positions that reach a stop-loss limit with no possibility for review would be foolish—the trader may have excellent information and research to back up his belief and automatic closing of the position would mean passing on a profit opportunity without the ability to review the limit in the light of the latest information.

I've rarely seen trading managers make this type of error, but an equally serious error in the other direction is unfortunately more common. Requests for temporary stop-loss limit increases by a trader reaching or approaching the limit may get approved without serious thought by a busy manager. They may be treated as bureaucratic box-checking exercises, particularly when the request comes from a respected trader with a good track record, rather than a genuine decision point. But this renders the stop-loss limit useless, as it will never actually be enforced (of course, there may be some limit beyond which even the most blasé manager will stop approving increases, but then it would be better to acknowledge this as the true limit in advance, since this will lead to better recognition of the actual maximum losses the firm faces).

A genuinely productive stop-loss limit review requires thorough discussion between traders and their managers of the factors that have led to the existing loss and the latest information on prospects for the position. Sometimes even an experienced trader with a great track record needs time away from the market to consider whether new factors have come into play that require a change in approach. Considerations of moral hazard, as discussed in Section 2.1, will certainly influence the discussion. Traders own more of the upside than the downside of their positions and so have an incentive to argue for raising the limit, and they can take advantage of their intimate knowledge of the market to cherry-pick data and arguments with which to make a persuasive case. Managers need to be aware of this informational asymmetry and employ a reasonable degree of skepticism while drawing on their experience of similar past situations and their outcomes. It also helps if the manager has been getting regular independent analyses of the causes of large gains and losses in the trading positions. This brings us to our next point.

- **Analysis of reasons for large losses and large gains** to put the manager in a good position to understand the logic of the trading strategy and to be able to review extension requests intelligently. In Sections 3.1, 3.2, and 4.1.6, we have already discussed the advantages for control of fraud and reporting errors of having control personnel develop thorough explanations of large moves in profit and loss (P&L), whether gains or losses. Here I want to emphasize how a robust P&L explanation process can also serve as excellent input for a manager who may need to review requests for stop-loss limit extensions. Since decision making on stop-loss limit extension requests must often be done under tight time limits in a stressful market environment, time that can be devoted beforehand to giving management deeper insight about the drivers of P&L in a trading book can have significant return on invested effort.
- **Financing plans.** Even when trading losses are well within stop-loss limits, management still needs to be concerned that it has adequate financing for the cash needs of maintaining positions, as the Metallgesellschaft case illustrates (see Section 4.2.2). There is thus a need to understand and forecast funding needs and plan for their financing.

It is now time to drop our unrealistic assumption that positions can be liquidated instantaneously. In virtually every case, when positions need to be liquidated, there will be some lapse of time between the decision to liquidate being made and the execution of the liquidation during which market prices can move. Stop-loss limits need to be set in light of the knowledge of such possible market moves. For example, if you want to be sure that you

don't lose more than \$100 million on a given position and you estimate that you could lose \$20 million in the course of liquidation, you need to set \$80 million as the trigger point for the stop-loss limit.

All five of the points just made about stop-loss limits under conditions of instantaneous liquidation continue to apply, perhaps even more strongly, but other risk control measures will be needed as well. The points made already are still needed to make the stop-loss limits effective, but, with less liquidity, failure to know current market prices can be even more damaging. To deal with the additional costs of liquidation, an estimate of liquidation costs will need to be available to managers, in the form of both a statistical probability analysis of likely market moves during a period of liquidation—called *value at risk* (VaR) analysis—and of stress scenarios to measure potential liquidation costs during periods of unusual illiquidity.

The risk control requirements we have outlined here are very close to the recommendations for managing derivatives risk that were issued by the Group of Thirty (G-30) in July 1993. These recommendations have proved very influential, not just for the management of derivatives risk, but for all trading risk. The Group of Thirty is a private, nonprofit organization that studies international economic and financial issues and is headed by 30 senior representatives of the international business, regulatory, and academic communities. The recommendations that relate most directly to the measurement of trading risk are shown in the box, with the original numbering they had in the G-30 report.

While the G-30 requirements and the approach being outlined here were developed in conjunction with market-making trading operations, they have much wider scope and should be used for any type of financial risk management—that is, any type of risk management that relies on liquid instruments to help manage risk. If you are planning to use liquid instruments to limit your losses, you need to estimate the likelihood that (and degree to which) the liquidity will be there when needed. So if you are managing credit exposure by having counterparties post margin, you need to make estimates of how effective that margin will be in limiting losses (see Section 14.3.3 for details). If you run a hedge fund and hedge positions with liquid instruments or you run a pension fund and are counting on the ability to liquidate positions to assure not dipping below funding requirements for future payouts (a *contingent immunization* strategy), you need to take possible limitations on liquidity into account (see Section 6.1.7 for details).

The rest of this chapter deals with how these recommendations should be put into practice, with many references to detailed discussion in subsequent chapters. But before getting to these specifics, I want to first lay out what I think are the essential components of any risk management

GROUP OF 30 RECOMMENDATIONS RELATING TO THE MEASUREMENT OF TRADING RISK

Here we review select recommendations by the Group of 30 on trading risk.

Recommendation 2: Marking to Market

Dealers should mark their derivatives positions to market, on at least a daily basis, for risk management purposes.

Recommendation 3: Market Valuation Methods

Derivatives portfolios of dealers should be valued based on mid-market levels less specific adjustments, or on appropriate bid or offer levels. Mid-market valuation adjustments should allow for expected future costs such as unearned credit spread, close-out costs, investing and funding costs, and administrative costs.

Recommendation 4: Identifying Revenue Sources

Dealers should measure the components of revenue regularly and in sufficient detail to understand the sources of risk.

Recommendation 5: Measuring Market Risk

Dealers should use a consistent measure to calculate daily the market risk of their derivatives positions and compare it to market risk limits.

- Market risk is best measured as “value at risk” using probability analysis based upon a common confidence interval (e.g., two standard deviations) and time horizon (e.g., a one-day exposure).
- Components of market risk that should be considered across the term structure include: absolute price or rate change (delta); convexity (gamma); volatility (vega); time decay (theta); basis or correlation; and discount rate (rho).

Recommendation 6: Stress Simulations

Dealers should regularly perform simulations to determine how their portfolios would perform under stress conditions.

Recommendation 7: Investing and Funding Forecasts

Dealers should periodically forecast the cash investing and funding requirements arising from their derivative portfolios.

Source: Group of Thirty, *Global Derivatives and Principles* (1993).

framework that will meet the needs identified earlier. I believe there are seven key principles that need to be considered:

- 1. Recognition of the nonnormal distribution of financial variables.** It is an empirical fact that nearly every financial data series exhibits fat tails (see the **Ratios** worksheet of the **VaR** spreadsheet on the website and Exercise 7.3 based on this worksheet for illustrative examples). Part of the explanation for this is the psychology of markets—a tendency for a big move to create panic that exacerbates the size of the move. Part of the explanation is that financial variables are mostly human creations rather than natural phenomena. As Nassim Taleb says in *The Black Swan*, “Money in a bank account is something important, but certainly *not physical*. As such it can take any value without necessitating the expenditure of energy. It is just a number!” (Taleb 2010, 33). To put it another way, when the world’s tallest man walks into a room full of many people, he will change the average height of the people in the room by only a small amount. But when Bill Gates walks into a room full of many people, he will change the average income of the people in the room by a large amount.

Whatever the explanation, risk managers need to recognize that financial data series are most likely fat tailed. They also need to recognize that large market moves in one financial variable often occur at the same time as large market moves in other financial variables, probably because investors will spread panic in one market to other markets. Therefore linear correlations are often very poor representations of the relationship between financial data series. *Any risk management process chosen must allow for handling fat-tailed series that have clustering of large moves.*

- 2. The need for simulation.** The need to handle fat-tailed series and clustering of large moves, as emphasized in the previous point, virtually dictates the need for using computer simulation to generate estimates of potential liquidation costs. More detail can be found in Section 7.1, but the basic argument is that simulation handles fat tails and clustering of large moves in a simple and transparent fashion, while other statistical estimation techniques are far clumsier and more opaque in how they handle these features of financial series.

Simulation consists of an initial specification of the distribution of underlying financial variables, followed by a calculation of the earnings impact of each instance of the distribution. The distribution of liquidations costs is then simply computed as the aggregation across individual cases. Since the step in which the distribution of the variables is specified is separate from the step in which earnings impact is calculated, there is

complete freedom to specify the distribution in the most accurate way possible. Furthermore, simulation offers many other advantages, some of which will be elaborated on in Sections 7.1 and 7.3:

- Many financial products, such as options, involve nonlinear returns. Simulation can handle this easily, since each path of the simulation computes the earnings impact independently from the computations on other paths and separately from the initial specification of the distribution of variables. Statistical techniques that mix together the specification of variables distribution and the calculation of earnings impact are much more vulnerable to error.
- Simulation makes it easy to generate a rich set of statistics on the distribution of liquidation losses, by aggregation of results across paths.
- Simulation makes it easy to attribute risk of potential liquidation losses to individual trading desks and individual positions.
- Simulation methodology can easily handle a range of desired calculations in addition to the basic calculation of liquidation costs. For example, consider the point made earlier in this section concerning the desirability of fitting stop-loss limits to trading strategies. In advance of setting a stop-loss limit, a manager should get some idea of the probability that the stop-loss limit will be activated by a particular trading strategy. This is straightforward for a simulation, since each individual case can be followed over a simulated time period, keeping track of whether a stop-loss limit has been hit along the way.
- Simulation methodology makes design of computations easy. Since each individual case of the simulation calculates earnings based on a single specification of the underlying variables, earnings calculations could be performed on each individual transaction by the exact same production models the firm uses for its official mark-to-market computations. Where this is computationally infeasible, due to a large number of individual simulation cases, approximations to production models are relatively easy to design and check against the official calculations. It is also easy to break up earnings calculations in each individual case by trading desk or product type. Since the earnings distribution is just a simple summation across individual cases, it is now easy to calculate the risk contributions of individual trading desks and product types.
- Checking is made easy by the separation of specification of distributions and the calculation of earnings into two separate stages. Control personnel and front-office personnel who may not be knowledgeable about probability distributions can focus on checking the earnings calculations, using the firm's mark-to-market models for each

individual transaction, as discussed in the preceding bullet point. By parallel reasoning, the specification of probability distributions can be easily checked by economists and statisticians who may not be knowledgeable about earnings calculations.

All of the advantages of simulation apply not just to the value-at-risk computations for relatively liquid positions discussed in Chapter 7, but also to the modeling of relatively nonliquid positions discussed in Chapter 8. This point will be elaborated in Section 8.4.2.

- 3. The need to consider subjective probabilities as well as objective frequencies.** As was discussed in Section 1.2, assessments of risk cannot afford to rely solely on historical frequencies. Subjective assessments of probabilities by the risk managers must be allowed to play a role. Even in computing historical frequencies, the risk managers must rely on some degree of subjective judgment regarding the length of historical period to use and the weight that should be placed on more recent historical experience relative to a longer period of history. These issues will be discussed in more detail in Sections 7.1.1 and 7.2.1.

The need to utilize subjective judgment causes concern for many risk managers. Without anything objective such as a historical data set to point to, how can they count on their recommendations carrying conviction? Why will they be accepted as having expert opinions on a subjective matter? These are questions that must be confronted—when subjective judgment is required, it is best to be frank about it.

The only way for a risk manager's subjective judgments to be accepted is to have well-researched and well-reasoned arguments backing them up. For a good example of what such an argument looks like, see Section 5.2.5.7 for the articles in which to find the arguments presented by the *Economist* magazine in 2004 through 2006 to support a belief that there was a good case to be made for a large drop in real estate prices. It is very important in presenting such an argument to explain carefully that a belief that there is a significant probability that an event will occur is very different from, and requires a very different type of evidence from, a belief that an event is highly probable to occur or is the most likely outcome.

Particularly when it comes to subjective judgments, ultimate decisions will rest with management. It is extremely rare for risk managers to carry enough political clout to be able to force acceptance of their subjective views. But it is surely not acceptable for risk managers to just state their views, have management disagree, and then shrug their shoulders and say nothing more. Risk managers must make their arguments forcefully and, if they believe that management is being unreasonable in its judgments, consider options such as taking their concerns to the risk committee of the board of directors or to regulators. A firm that does

not allow a senior risk manager freedom to do this (on occasional significant points of disagreement) without damaging her career prospects is not a healthy working environment. And, when large disagreements occur more than occasionally or where the freedom to appeal is not part of the culture, risk managers must consider “voting with their feet,” to protect their reputations and integrity.

I will relay one anecdote with respect to voting with your feet, though it preceded the days of dedicated risk management departments. An economist with whom I had worked closely at Chase in connection with the introduction of options products had left to take a good offer at a smaller firm. He later relayed his experience in joining that firm: When he asked for some orientation on how they measured their options risk, management responded by saying they had no need for such measures and could not be persuaded that such measures were needed—they just made “holistic” judgments about the positions they wanted to take. My former colleague reported thinking to himself, “I see—have a hunch, bet a bunch,” and immediately decided to start seeking other employment.

- 4. The distinction between diversifiable and nondiversifiable risk.** The difference between systematic, diversifiable risks and idiosyncratic, non-diversifiable risks is one of the cornerstones of modern finance theory, as developed by Harry Markowitz, William Sharpe, Stephen Ross, and others. Discussion of this critical distinction can be found—in the context of expositions of portfolio theory, the capital asset pricing model (CAPM), and arbitrage pricing theory (APT)—in any textbook on investment theory (see, for example, Bodie, Kane, and Marcus 2009, Chapters 8 and 9); on corporate finance (see, for example, Brealey, Myers, and Allen 2011, Chapter 8); or on asset pricing theory (see, for example, Cochrane 2001, Section 1.4).

A diversifiable risk position can be reduced in several ways, by direct hedging but also by diversification through investing in other positions that have low correlation with it. A nondiversifiable risk position, such as exposure to the Standard & Poor’s S&P 500 index or to interest rate levels, needs to rely almost entirely on direct hedging to reduce risk. Finance theory emphasizes the resulting market demand for high returns on nondiversifiable risks.

Risk managers need to ensure that trading management is especially aware of sizable exposures to nondiversifiable risk, both because it may be more difficult to reduce such positions and because management will want to ensure it is receiving adequate returns for taking on these risks. This was a particularly important issue in the 2007–2008 financial crisis (see Sections 13.4.4, 5.2.4, and 5.2.5.3).

Diversifiable risk can be eliminated through hedging; nondiversifiable risk cannot be eliminated but can only be transferred to someone else. Risk managers need to be very sure they understand this risk transfer process, to make certain that the risk is truly being transferred and not just reappearing elsewhere on the firm's books. We will discuss this further in Section 14.3.4 on wrong-way counterparty risk.

5. **The use of arbitrage theory to decompose risks.** Suppose that you have some exposure to euro interest rates through interest rate derivatives and also some exposure to euro rates through forward U.S. dollar-euro foreign exchange contracts. If these two positions were treated as completely different types of exposure, you might miss offsetting exposures to euro interest rates or, even worse, fail to measure a dangerous risk concentration by not adding together the euro interest rate exposures in the two positions. (This is not simply a hypothetical example—this treatment of interest rate exposures from interest rate swaps and foreign exchange forwards as separate exposures was often encountered in the 1980s in the risk management of major institutions.) Arbitrage theory for derivatives, which has been well developed over the past 40 years, as exemplified in the material in Hull (2012) and similar textbooks, has provided a valuable tool kit for unifying such positions. I will be drawing heavily on arbitrage theory for unifying positions throughout Chapters 10 through 14.
6. **The need to consider periods of reduced liquidity.** All of us who participate in financial markets have experienced several periods of severely reduced liquidity when the ability to trade at anything other than fire sale prices dries up for a prolonged period of time. Estimation of potential liquidation costs associated with stop-loss limits must account for the possibility that liquidation will be required during such a stressful period. In fact, it is often during such periods that stop-loss limits are breached, since lowered liquidity is often associated with sharp price swings and managements may need to cut risk limits in a crisis period. I argue in Section 7.2.1 that detailed scenarios based on subjective judgment must play the key role in this analysis but that there is still room for using simulation based on historical data as a supplement.
7. **The need to distinguish degrees of illiquidity with different tools to handle each type.** Given that projecting possible liquidation costs of a position are such an important part of risk management, it is natural that different tools are required based on the degree of liquidity of a position.

In teaching classes on this topic, I like to use a variation of a quotation from Shakespeare, who said, “Some men are born great, some achieve greatness, and some have greatness thrust upon them” (*Twelfth Night*, Act 2,

Scene 5). My variant is “Some positions are born illiquid, some achieve illiquidity, and some have illiquidity thrust upon them,” and each of these three types of positions requires a different type of risk management.

Positions that have illiquidity thrust upon them are positions in instruments for which frequent liquid market quotations are available, and they are *not* of such large size that liquidation of the position will significantly impact market price. These positions will become illiquid only under conditions of extreme and unusual market stress. These are the type of positions well handled by normal mark-to-market pricing (Section 6.1.3) and VaR calculations (Section 7.1), supplemented by stress tests (Section 7.2) to consider the possible impact of an unusual market stress that causes a normally liquid position to become illiquid over a period of a few weeks during a large market move.

Those positions that achieve illiquidity are also positions in instruments for which frequent liquid market quotations are available, but where position size has grown to the point that liquidation *will* significantly impact the price. In such cases, the risk tools just referred to must be supplemented by a way of measuring the impact of this larger position size. Note that an illiquid position size should impact both VaR calculations (since liquidation even in normal market conditions will come with added cost) and stress-test calculations (since once a period of unusual market stress has been weathered and more normal liquidity has returned to the market there will still be added costs to liquidating a large position). My suggested approach for handling large positions is a separate and supplementary simulation of the distribution of possible costs to be incurred, explained in Section 6.1.4.

Positions that are born illiquid are positions in instruments that lack liquidity even during the best of market conditions. They are the positions referred to in Section 1.1 as having actuarial risk. This could be a transaction that completely lacks a market component; we’ll discuss an example at the beginning of the next section. Or it could be an instrument with very limited liquidity; a good example is a position in a one-way market, as discussed in Section 6.1.3. It could be a position that is so large relative to the size of daily trading that it cannot be liquidated even over an extended period; an example would be the loan books of most banks, as discussed in the introduction to Chapter 13. It could be a position that can only be sold under certain conditions, such as restricted stock. It could be an instrument that is so complex that liquidation in any reasonable time period is unlikely. (One of the first consulting assignments I had in the wake of the 2007–2008 crisis was for a large bank looking to find a methodology for valuation of collateralized debt obligation (CDO) tranches it was holding. When I asked whether the valuation was just for accounting purposes or was meant to drive decision making on possible sales, the somewhat testy

response was: “I couldn’t possibly sell any of these tranches—it would take me six months just to explain all the cash flows to a potential buyer.” Even allowing for hyperbole born of frustration, there is likely to enough truth in this comment to serve as a warning about assuming liquidity on highly complex positions.)

It is the management of positions in instruments that lack liquidity that presents one of the greatest challenges to financial risk management, as has been confirmed by the 2007–2008 crisis that was largely due to inadequate risk management of illiquid CDO tranches (see the discussion in Section 5.2.5). I will therefore address a recommended approach to this issue in a separate section.

6.1.2 Risk Management of Instruments That Lack Liquidity

When a market component is completely lacking, financial risk management techniques may be wholly inappropriate and it may be proper to manage risk utilizing the type of actuarial techniques we discussed in Section 1.1. A good example of how to identify instruments for which a market component is completely lacking and how to manage this type of risk can be found in the excellent discussion of weather derivative options in Jewson, Brix, and Ziehmann (2005, Section 1.4). They declare that “for locations where the [weather] swap is not traded, and which are not highly correlated with locations on which swaps are traded, actuarial valuation of the options is the only choice.” They specify that actuarial valuation is “fundamental analysis” of the type used in pricing insurance contracts, based on “historical meteorological data and meteorological forecasts to predict the distribution of possible outcomes.” When weather swaps are traded for the location to which the option is tied, or on locations whose weather is highly correlated with this location, then they advocate valuation based on market prices and arbitrage pricing models (e.g., Black-Scholes).

This is a good illustration of the approach I support. Positions that have no liquidly traded instruments that can meaningfully be used as hedges should be evaluated and managed just as if they were positions of a conventional insurance company. The general principles of risk management referred to in Chapter 1 apply; the financial risk management principles that are the subject of this book are irrelevant. But when liquidly traded instruments can be used to hedge a meaningful portion of the risk, then we can utilize financial risk management.

For this latter case, the approach I strongly favor is to (1) set up a liquid proxy that allows the total risk to be split into liquid risk and illiquid risk; (2) use the liquid proxy in all standard risk reports and limits (e.g., position reports, VaR, stress tests); and (3) use a separate simulation to manage the

risk of the mismatch. As an illustration, consider the example discussed in Section 10.2.2, a 40-year interest rate swap in a market that has interest rate swap liquidity out to only 30 years. A 30-year swap would be assigned as the liquid proxy and used in all standard risk reports, while a separate simulation would be used to assess the risk of using a 30-year swap as a hedge against a 40-year swap.

The reasons I favor the use of a liquid proxy to represent positions in illiquid instruments are:

- Some of the risk in an illiquid instrument can be managed by liquid instruments, and the use of the liquid proxy ensures that this possibility can be exploited. To continue with our example of the 40-year swap, the booking of 40-year swaps certainly exposes the firm to interest rate risk over the first 30 years of the swap in addition to the exposure for years 31 through 40. Use of the liquid proxy assures that this exposure to the first 30 years shows up in position reports and limit calculations properly added in to liquid instrument exposures taken in the same direction. This will alert management to concentrated exposures and give traders and marketers proper incentives for hedging that portion of the risk that can be hedged through liquid instruments. If 40-year swaps were treated as a completely separate category from swaps of 30 years or less, these goals might still be accomplished, but it would require the building of a completely separate reporting structure and there is always the possibility of gaps occurring in the design of new reporting structures. The simple act of insisting on a liquid proxy takes advantage of all existing reporting structures, such as mark-to-market, VaR, stress tests, and position limits (e.g., maturity bucket limits), with no further effort beyond the calculations currently in place for these reports.
- Basing reserves and limits for the illiquid risk just on the variability of the mismatch between the illiquid position and its liquid proxy should often lead to reducing the need for reserves and limits. Continuing with our example, managing the risk on the difference between a 40-year swap and a 30-year liquid proxy over the 10-year period that you must wait until the 40-year swap becomes liquid (after 10 years it has only 30 years left to maturity) will almost certainly be computed as significantly lower than the variability of return on an unhedged 40-year swap. But this latter computation would be an overstatement of the risk of the trade, since use of a hedge involving the liquid proxy is always a choice the trading desk can make. Computation using the liquid proxy does not in any way require trading desks to make use of this actual hedge—but, if they don't, the additional risk will show up as a use of their trading limits for VaR, stress tests, and positions.

- Less compelling, but still of some weight, is that making sure that even illiquid positions make an appearance in standard reports, such as VaR and stress tests, makes it less likely that managers will forget about these risk positions. It serves as a reminder. But true measurement of the risk of an illiquid position cannot be accomplished solely through standard reports designed for liquid positions. There must also be a separate and well-thought-through report on the potential cost of the mismatch between the illiquid position and the liquid proxy. This will be our next point of discussion.

Modeling the potential impact of the difference between the actual trade and the liquid proxy should use simulation for the similar reasons as given for using simulations in Section 6.1.1—to reflect a full range of possible outcomes and to generate a statistical distribution that can be used in assessing issues such as capital adequacy. But simulations of the differences between the actual trade and its liquid proxy cannot just be for the short periods used in VaR calculations; they must go all the way to final payout or to when the trade becomes liquid. Simulations must reflect the possibility that the model used for pricing and trading the product may be wrong. All of these issues in simulation of the difference between the actual trade and the liquid proxy will be discussed in detail in Section 8.4.

More controversially, I do not believe in using mark-to-market pricing on the difference between the actual illiquid trade and the liquid proxy. Reserve levels should be adequate to protect against extreme events, and it is extremely rare that short-term market changes reveal new information about the potential depth of an extreme event. Mark-to-market pricing is designed to measure prices at which risk can be exited in the near future. Since an illiquid trade cannot be exited in the near future, mark-to-market pricing is not truly reflecting changes that impact the position.

I know that to many people this will seem as if I am trying to go easy on illiquid positions, which would be particularly foolish in light of all the havoc overindulgence in illiquid products by financial firms caused in the recent crisis. But I do not believe this proposal is moving in the direction of easier treatment of illiquid trades. The size of reserves I want to keep is quite substantial, and my experience with this reserve methodology leads me to believe it would lead to less use and more cautious use of illiquid products than the wholly inadequate reserving processes that appear to have been operating in the run-up to the recent crisis. I will ask readers to withhold judgment until they can see my argument in detail in Section 8.4.4. I discuss how my proposal might have mitigated the spread of the crisis in the discussion of reducing procyclicality in Section 5.5.8.1.

Given the very different treatment I am advocating for illiquid instruments relative to liquid instruments, it is vital to have good tests available to distinguish between illiquid and liquid instruments. When trading desks make suspect claims of liquidity, independent risk managers need to insist on evidence from reliable external sources or from a history of actual trading tickets. Trading history that is overwhelmingly in one direction, extremely sporadic, or concentrated with just one or two counterparties, or that has been executed at prices substantially different from internal valuations needs to be regarded with extreme wariness. Trading desks that want to overcome such suspicion should be prepared to demonstrate the ability to liquidate significant blocks of inventory at prices close to internal valuations.

6.1.3 Market Valuation

The policy of marking to market all trading positions, at least as often as the close of business each day, as per the G-30's Recommendation 2, constitutes the essential foundation for measuring trading risk because of three primary reasons. First, without a nearly continuous marking to market, it would be possible that ineffective hedging strategies would not be recognized until long after being put in place. Second, the analysis of revenue will yield insight only if the revenue figures being analyzed are tied to genuine changes in value. Third, in measuring the risk exposure to market moves, it is far easier to make good judgments about possible short-term moves than it is about longer-term moves. But if trades are not revalued frequently, it becomes necessary to measure risk exposure over longer periods.

When highly liquid external prices are available for marking a position to market, then the issues involved in performing the mark are largely operational. An example might be a position in spot foreign exchange (FX) for the dollar versus Japanese yen. This is a market for which quotations are readily available on trading screens, with market conventions that ensure that firms posting prices are prepared to actually deal in reasonable size at these prices. Quotations for mark-to-market purposes can be captured electronically from trading screens or entered by hand and later checked against printouts from screens—the choice should be based on the operational cost versus error rate and the cost of correcting errors. Another example would be a position in a well-traded stock or exchange-traded futures option for which the last price at which an actual trade occurred is readily available from an exchange ticker.

For many positions, mark-to-market pricing is not this straightforward. Either the market itself does not have this type of liquid quote available or the size of the position held is so large that closing it out

might impact the market. The price at which the position can be exited will be uncertain to some degree. In such cases, two interrelated questions must be asked:

1. How should a most likely exit price be arrived at?
2. Should some markdown of the price be used to account for the uncertainty and, if so, how should the amount of reserve be determined?

Establishing the most likely exit price may require a model to create a mark based on more readily available prices of other instruments. Models can range from very simple computations, such as the interpolation of an illiquid two-and-a-half-year bond from prices on more liquid two- and three-year bonds, to complex theoretical constructions. A discussion of how to use models in the marking process and how to establish reserves against the associated uncertainty can be found in Chapter 8.

What if price quotes are available, but are not sufficiently liquid for a readily agreed-upon external valuation? This implies that deriving the most likely exit price from these quotes requires an understanding of the relative quality of available quotes. For each quote, questions like the following need to be answered: Is the quote one at which the firm or broker providing the quote is offering to do business, or is the quote just provided as a service to indicate where the market is believed to be today? If the quote is an offer to do business, how large a transaction is it good for? What is the track record of the quotation provider in supplying reliable information? Are there possible motivations to provide misleading information in an attempt to influence pricing to move in a direction that favors a quote provider's position? How frequently are quotes updated?

With such a multiplicity of information bearing on the issue, there is no doubt that traders of an instrument have the best judgment on determining this valuation. Their continuous contact with other firms' traders and brokers enables them to build the experience to make these judgments. The ability to make such judgments is a major factor in determining a trader's success, so traders who have built a successful earnings track record can make a strong claim of having the expertise to determine most likely exit prices.

Unfortunately, reliance on traders' judgment raises moral hazard concerns. As discussed in Section 2.1, traders are often tempted to mislead management about position exit prices in order to inflate reported profits or to increase flexibility in the positions they are allowed to hold. Outsiders, from corporate risk management, corporate finance, or the middle office, need to be involved in making these judgments to preserve independence. However,

designing mechanisms for resolving disputes between traders and control personnel raises many difficult issues:

- How can control personnel obtain a sufficient knowledge base to challenge traders' judgments? At a minimum, traders should be required to make public the information on which judgments are made. This can be accomplished by insisting that quotes be sent to the firm in writing (whether through trading screens, e-mail, or fax). Alternatively, control personnel should have the right to selectively listen in on phone conversations in which quotes are made.
- Ideally, control personnel should have a range of experience that enables them to arrive at independent conclusions regarding quotations, perhaps even prior trading experience.
- Records should be kept of prior experience with the reliability of a particular trader's valuations by tracking the path of internal marks leading up to an actual purchase or sale price and noting suspicious patterns. Control personnel should adjust their deference to trader valuations by the degree of proven reliability.
- A trader's ultimate weapon for bringing credibility to a valuation is to actually exit part of the position. A recorded price narrows disputes down to the single question of whether the size of the trade relative to the retained position is large enough to be a reliable indicator of the exit price for the remainder.

Despite best efforts to design dispute resolution processes that balance power between traders and control personnel, traders inevitably retain a strong advantage based on information asymmetry. They can utilize their knowledge of a wide variety of sources of price quotes to selectively present only those that are the most advantageous to their case. They sometimes use friendships and exchanges of favors to influence other market participants to provide quotes biased toward their valuations. Traders also often rely on an aggressive personal style and internal political power based on their profitability to prevail through intimidation.

To remedy this power asymmetry, some firms prefer to rely on more objective computations for determining valuations, even where this reduces accuracy by lessening the role of judgment. A typical approach would be to average the quotes obtained from a set panel of other firms or brokers, perhaps discarding outliers before averaging (discarding outliers is a possible protection against a few quotes that have been biased by friendship or favor). Changes in panel membership should be difficult to make and require agreement between traders and control personnel.

A promising development toward more objective valuations for less liquid instruments is the Totem Market Valuations service. This service is designed to share information among firms making markets in less liquid products. Firms can obtain access to quotes on only those products for which they are willing to provide quotes. Their access to quotes can be cut off if the quotes they provide are frequently outliers, indicating either a lack of expertise or an attempt to bias quotations. Although the extensive machinery of this process means it can make quotes available only once a month and with a lag of a few days, it still provides a valuable check on the valuations of a firm's traders.

The following are some pitfalls to be wary of when setting up a procedure for deriving valuations from less than fully liquid market quotes:

- **Model-derived quotes.** Here is an illustration of a frequently encountered problem. You need a valuation for a particular bond and you have a choice: either use a model to compute the value based on observed prices of more liquid bonds with similar maturities and credit ratings or use price quotes for the particular bond obtained from brokers. Before choosing the latter, ask the following question: Are the brokers providing a quote specific to this bond or are they just providing the output of their own model based on prices of more liquid bonds? If your external source is model based, might you be better off using your own model? The following are some advantages to using a model-based external quote:
 - You may be able to get model-based quotes from several sources with the hope that errors will average out.
 - The external models are being tested by the use of the quotes by many different firms, so it is more likely that objections will be raised if the model is missing something.
 - It is less likely that traders will influence the outcome when an external source is being used.
 - The quotes may become so widely used as to be a good indicator of where the market is trading.
 - The primary disadvantage to using a model-based external quote is that you may not be able to obtain details of the model used, so it is harder to estimate potential error and build adequate reserves for uncertainty than when using your own model.
- **Revealing positions.** When quotes are not available on regularly displayed screens or reports, firms seeking quotes may need to make specific inquiries to obtain quotes. Their inquiries reveal information about the positions the firm holds that can be used to the firm's disadvantage by other market participants. This is particularly true if the

conventions of the market require an indication of either buy or sell interest to obtain a quotation, as opposed to obtaining a bid-ask quote. Even when you do not need to reveal the direction of your interest, in some markets the direction of a firm's position is well known to other participants and the expression of interest in a particular instrument is highly revealing of holdings. It is always possible to disguise positions by requesting quotes for a range of instruments, including instruments held and not held. However, the quality of the response may suffer as efforts to provide quotes get diffused over too many instruments. Market conventions concerning the tolerated ratio of inquiries to actual transactions also limit the amount of information that can be obtained. If trader reluctance to reveal positions limits the extent of the external quotes obtained, models may need to be relied on more heavily to infer valuations.

- **One-way markets.** You not only need to worry whether the size of transaction for which an obtained quote is valid, but you must also worry about whether the quote is valid for your firm. Markets that tend to be one-way, with customer demand strongly on one side and market maker supply on the other, may lead to quotations that are good for customers only. A typical example would be an options market in which almost all customer interest in options beyond five years was to sell options, not buy them. A market maker, in such circumstances, might supply reasonably liquid quotes for the purchase of long-term options to customers, but be unwilling to buy on the same terms from other makers. The principle is to reserve the limited capacity to take on risk to encourage customer relationships, not to help competitors for this customer business by allowing them to distribute some of their risk. This is not to say that market makers will never buy longer-term options from one another in such circumstances, but they may do so only on a negotiated basis, with no actionable quotes available, even through brokers.

A market maker may still succeed in finding out the prices that other market makers are paying customers for longer-term options, since customers often let them know what bids they are seeing from other firms. It will be a definite source of comfort to know that the firm's prices are in the same range as their competitors' prices, since this is an indication that the firm's models and trading strategies are not suffering from some major error, such as overlooking a source of risk. Equivalently, a firm derives comfort from seeing that it wins its fair share of deals in a given category, neither too many nor too few.

Although this comfort is genuine, it should not be confused with obtaining a price at which the firm can exit its risk positions. In the absence of

quoted prices at which the firm itself can transact, it is prudent to anticipate the need to hold risk longer and to utilize models to estimate longer-term profit and loss (P&L) and reserves, and limits to control the associated risk, as discussed in Section 8.4.

6.1.4 Valuation Reserves

When there is substantial doubt about the price at which a position can be exited, a safety margin can be provided by calculating a *valuation reserve* that can be subtracted from the most likely exit price.

The issue of how large reserves should be for valuation uncertainty is probably the single issue that leads to the greatest conflicts between traders and corporate risk managers. Based on their experience and knowledge of the motivations of the creators of market quotes, traders tend to believe they know the price at which positions can be exited with a fair degree of certainty. With some justice, they will point out that the uncertainty is mostly on the part of the outsiders, such as corporate risk managers and the corporate finance function, who do not have the traders' access to information. Reserves lower the reported P&L, which is the ultimate scorecard for the traders, determining bonuses, promotions, the size of positions management will allow, and, ultimately, continued employment. Understandably, traders will push to minimize reserves. (The one universal exception to this tendency is a trader who inherits a book from another trader. Invariably, the new trader will want to increase reserves for the inherited positions. I call this the principle that no profit should fund only a single bonus.)

Occasionally, though, one encounters a trader who claims to be a proponent of large reserves. I came across one when a trading book of exotic options was being established for which I was to be responsible for the risk management. The head trader expounded on his philosophy of avoiding any appearance of claiming too much P&L before achieving certainty of the results. He wanted reserve levels to be generously high. Here, I thought, was someone I could get along with well. And so I did, through many months in which both P&L and reserve levels were high, with easy agreement between the two of us on the reserves.

Then came the unfortunate day when an operations error in booking a trade was discovered several months after the trade had been booked. Rebooking the trade correctly would lead to a large loss, large enough that the trading desk would show its first negative P&L for a month. The head trader, although duly upset by the operations failure, was unfazed by the P&L consequence. Now, he informed me, was the time to release some of that reserve that had been accumulating—just enough to make P&L for the month come out positive. I protested. First of all, the reserves had been

created for valuation uncertainty, not as a hedge against possible operating errors. Second, the amount of uncertainty in the valuation was exactly the same on the day after the error was discovered as it was on the day before it was discovered. So how could a lowering of reserves be justified? The era of good feelings had come to an abrupt end.

This experience illustrates why a great deal of suspicion exists around valuation reserves, which is often expressed by regulators, such as the Securities and Exchange Commission (SEC) and auditors. Aren't reserves just a cushion to allow reported earnings of a trading book to be smoothed, creating an illusion of less uncertainty of return than actually exists? To avoid this, a definite principle must be in place that reserves are strictly for uncertainty concerning current valuation and never for uncertainty concerning future market variation. As an example, take a position in a liquid instrument, such as the dollar versus Japanese yen spot FX we previously cited. The future movements of this highly volatile exchange rate (and hence the P&L) may be surrounded by great uncertainty, but there should be no reserve, since the position can be exited at short notice at a known price. A reserve should be considered only if the position reaches a size that places a limit on this freedom of exit and therefore calls into question the valuations of the current position.

To make sure that this principle of using reserves only for the uncertainty of current valuation and not for the uncertainty of future market variation is followed, clear independence of reserve determination from the control of insiders with a motivation to show smoothed earnings must be demonstrated. This requires the final decision authority to be with an independent business unit, such as corporate risk management or corporate finance, and relatively objective standards for determining reserves to be utilized.

The uncertainty of current valuation could be due to the illiquidity of available price quotes or it could be due to reliance on a model to obtain a valuation. Section 8.4 discusses how to establish objective standards for reserves against model uncertainty. We now focus on how to establish objective standards for reserves against positions for which only illiquid price quotations are available.

The most direct method for reserving against an illiquid position is to estimate the degree to which exiting this position in the market might cause prices to move. This can accommodate fairly objective standards by using *haircut* tables on valuation. These have set percentage discounts tied to the size of the position held relative to some measure of the market size, such as the average amount of daily trades. This method takes proper account of both types of possible illiquidity, since this ratio could be high based on either a small denominator (indicating an illiquid market) or a large

numerator (indicating a big position in a liquid market). The downside to this method is that it may be difficult to establish reasonable haircut percentages to use. Rarely do firms keep good historical records of the impact of exiting large positions, and it will, in any case, be very difficult to sort out such impacts from other effects on market prices. This leaves the determination of haircut percentages to a subjective debate in which the traders' greater experience will be difficult for outsiders to question.

A method that lends itself to a more evenly matched debate is to first estimate the amount of time it will take to exit a position without substantially moving prices and then reserve against a possible market move over this time period. This exit time estimate will also be based on a ratio of size of position held to daily trading volume. It thus shares the previous method's advantages of taking proper account of both types of possible illiquidity and also the previous method's disadvantage of making it difficult for outsiders to debate trader judgment. However, the potential price move estimate allows for outsider objectivity, since it is very similar to the sort of calculation that goes into VaR. It also enables reserve levels to be calibrated to management-determined levels of uncertainty that should be reserved against. A uniform uncertainty level used for different trading desks can help to ensure the comparability of results across the firm.

For example, consider a \$500 million position in a stock in which the amount that can be transacted in one day without adversely impacting prices is estimated to be \$50 million. So \$500 million/\$50 million = 10 days of price moves should be reserved against, which implies that on average there will be $10/2 = 5$ days of price moves prior to sale. If the daily standard deviation of price moves is 1.5%, and if management decides on a reserve to a 95% confidence level, which is equivalent to 1.65 standard deviations of a normal distribution, then the reserve level should be:

$$\$500 \text{ million} \times 1.65 \times 1.5\% \times \sqrt{10/2} = \$27.7 \text{ million} \quad (6.1)$$

It should be reiterated that despite the appearance of a term that is tied to the uncertainty of future market variation, this remains a reserve methodology based on current valuation uncertainty. Future market variation is being reserved only to the extent it is outside management control, due to a large position size preventing exit at the desired time.

A third method, which can be used as a complement to the other two, is to create a reserve against aged positions. This method establishes a formula that marks a position down by a certain percentage the longer it is held. This can only be used as a complement to one of the other two methods, since it will not establish any reserve against a large illiquid position recently entered into.

Why should there be uncertainty about position valuation just because a position has been held for a long time? It is based on the observation that traders may delay exiting a position when they suspect that it will cause a decline in value from the level they are currently marking it at. Although I have heard much anecdotal evidence supporting this observation, it would be intriguing to perform a statistical study on the correlation between the length that positions are held and the size and direction of the price move between the last mark and actual sale. An aging reserve policy can also be justified on the pragmatic grounds that it is providing traders with the right incentives—to realize profits and cut losses in a reasonably short time period.

As was stated at the beginning of this section, reserving against valuation uncertainty is probably the leading cause of the greatest conflicts between traders and corporate risk managers. The risk managers need to provide a degree of conservatism that will assure investors, lenders, and government regulators that P&L is not being overstated and must provide a degree of independence to allay suspicions that reserves are being used as a means for smoothing earnings results. However, this leads traders to suspect that too much conservatism is being applied to protect risk managers against any possibility of criticism. A reserve that is too conservative hurts not only the trader, but also the ultimate profitability of the firm by limiting the amount of business that can be transacted.

In my experience, traders often misunderstand the need for conservatism and independence. One argument I've frequently encountered when specifying the reserve I think needs to be placed on an illiquid position goes something like this: "If you want the firm to value the asset at that low a value, then you would be happy if I went massively short the asset at that price." However, this mistakes conservatism for a view on fair price—if the trader was to go short this illiquid asset, I would want reserves to establish a conservatively high value for the short position. In other words, reserves are used to establish a bid-ask spread on an illiquid position, and the greater the illiquidity, the wider the spread. I've also encountered the argument from traders that they have excellent inside information as to where a position will trade, but they don't currently want to enter into the trade at that price. I need to point out that unless they can find some means of translating inside information into something publicly verifiable, we cannot ask the firm's shareholders and depositors to bear the risk that they are wrong.

Of course, my dialogue with traders is far from a one-way street. Often it is a case of their educating me on sources of information or aspects of hedges that cause me to change my initial view. Over time, with almost all of the traders I've dealt with, we've come to an accommodation of mutual respect, but with a realization that our interests sometimes differ. However,

I still wonder at times whether other risk managers have found better ways to avoid initial contentiousness. I was therefore a bit amused at some dialogue I overheard.

I was meeting with the head of market risk at a major investment bank, one of the most respected individuals in the industry. Our conversation was interrupted by an urgent phone call from one of his staff. I heard only his side of the phone conversation, which went something like this: “Well, certainly you need to put a reserve on a trade like that. . . . I don’t care whether the trader likes it. If he doesn’t, let him sell some of the position and show us where it should be priced. . . . You can’t accept a statement like that from him. The fact that the reserve you’ve calculated would make him book an up-front loss doesn’t prove that your reserve is stupid. Tell him that your reserve calculation shows that his price is stupid.”

6.1.5 Analysis of Revenue

The G-30 study states, in support of Recommendation 4 to identify revenue sources, that “measuring the components of profit helps participants to understand the profitability of various activities over time relative to the risk undertaken, as well as to gain insight into the performance of hedges.” A basic justification of using mark-to-market valuation in the management of risk is that it will lead to an early identification of ineffective hedging strategies, which can trigger experimentation with alternative hedges or changes in the mix of products being offered. This can happen only if an effective and frequent analysis is made of what is causing changes in P&L. In particular:

- P&L must be segregated by product line to identify which products may be encountering hedging difficulties.
- P&L must be broken out into that part attributable to newly booked business versus that part attributable to hedging activity on existing business. This ensures that hedging problems will not be masked by the offset of profits from new business, leading to a Ponzi scheme, as discussed in Section 2.2. A persistent pattern of profitable new business offset by hedging losses is an indication that either traders have chosen to take positions that (at least temporarily) have had bad results or valuation reserves have been inadequate.
- To distinguish between these two cases, it is important to identify what portion of hedging profits is due to movements against specific risk factors, such as delta, gamma, vega, and theta. In this way, losses stemming from deliberately taken positions can be distinguished from those that arise from risks such as correlation exposure, which the trader

cannot completely hedge. This analysis is also important in confirming that risk positions are reported correctly. If daily P&L swings cannot be accounted for by the reported size of risk positions and the daily changes in market variables, it is a warning that the reported risk measurements may be incorrect. This should lead to investigations of whether some transactions have been misrepresented in the reporting systems or whether additional or more detailed risk measures are required. Particular attention should be paid to unexplained P&L swings that take place around a date on which a payment is made or determined. If a model is not properly valuing a payment that has already been determined or is very close to determination, the probability is very high that the trade has been misrepresented. More detail on this point will be found in Section 8.2.7.1.

- It is extremely important to highlight any P&L changes due to changes in those assumptions that cannot be directly tested against available market prices or changes in models. This eliminates the possibility that P&L due to such changes will mask the results of ineffective hedging strategies.
- Significant differences between official P&L changes and the informal trading desk estimation of these changes should be investigated. These differences can be indicators of hedges that are not performing as expected.

6.1.6 Exposure to Changes in Market Prices

The need for measuring exposure to market changes is emphasized in G-30 Recommendations 5, 6, and 7. Proper daily mark-to-market valuation, as discussed in Section 6.1.3, is the key to properly measuring the exposure to changes in market prices. The correct daily valuation ensures that exposure is being evaluated from the correct starting point and also serves as a basis for translating changes in observable market prices into changes in portfolio valuation. Since the daily mark-to-market needs to relate valuation to some observable external prices, possibly through the use of models, this same relationship can be used to take a change in market price and convert it into a change in instrument value.

To take a concrete example, consider an option position on the Standard & Poor's S&P 500 index with an expiry in five months. When considering how to value it, decisions must be made about what model to use and what the inputs to the model should be. Let us say a Black-Scholes model is chosen, requiring input for the price of the underlying and an implied volatility. For the underlying price, we might decide to use an average of one-third of the closing three-month S&P futures price and two-thirds of the closing

six-month S&P futures price. For the implied volatility, we might decide to use an average consisting of one-third of the implied volatility of the closing three-month S&P option price and two-thirds of the implied volatility of the closing six-month S&P option price. These choices will be made based on trade-offs between basis risk and liquidity risk and could include reserve adjustments for lack of liquidity. However, once the choices are made for valuation, they become simple recipes for translating changes in market prices of the three-month S&P futures, six-month S&P futures, three-month S&P implied volatility, and six-month S&P implied volatility into a change in the five-month option price, utilizing the Black-Scholes model.

Once these pieces have been established, the remaining task is to decide on the market price shifts on which to calculate exposure. Three primary types of shifts are used:

1. **Standard shifts such as a 1 basis point interest rate move, a 1 percent stock price move, or a 1 percent implied volatility move.** The advantage of standard shifts is that they easily convey a precise meaning to a wide group of users. The main issue to be decided when using standard shifts is which market prices to group together—do you want to report exposure to each individual stock price moving, all stock prices moving together, or a particular industry shifting relative to all others? These detailed decisions are best examined in the context of specific risks. We address these decisions more closely in subsequent chapters, particularly Section 7.1, Section 8.4, and Section 9.4.
2. **Shifts based on the statistical analysis of the probability of the size of the change.** The advantage of statistically based shifts is that they make it easier to compare the size of exposures in different risk classes. For example, it's hard to say whether a \$5 million loss for a 1 percent change in stock prices is more of a danger or less of a danger than a \$2 million loss for a 1 percent change in implied interest rate volatilities. However, a \$5 million loss for a stock price change that has a 5 percent probability of occurring is clearly more worrisome than a \$2 million loss for an implied interest rate volatility change that has a 5 percent probability. Probability distributions also make it possible to combine shifts in unrelated asset classes into a single measure, such as the 95th percentile VaR, defined as the amount of loss that will be exceeded only 5 percent of the time, based on all of the positions within a portfolio. The difficult issue with statistically based measures of risk is how to determine the probability distributions. These measures and the means of deciding on distributions are discussed in Section 11.1.
3. **Shifts based on scenarios determined by economic insight into the potential size of different shifts and the relation between them.** An example

would be a stress scenario for the impact of the debt default of a particular large developing economy, which might be judged to result in, say, a 5 percent decline in all stock prices, a larger decline in the stock of companies with large investment in that economy, a 10 percent decline in all emerging market FX rates, a 15 percent increase in the credit spread of all emerging market debt, and so on. We study alternative approaches to defining such shifts in Section 11.2. Scenario analysis is needed for cash flow as well as for P&L to anticipate funding liquidity problems, which is consistent with the G-30 Recommendation 7, as discussed in Sections 3.5 and 4.2.2.

6.1.7 Risk Measurement for Position Taking

It can be argued that the G-30 recommendations should apply to the market-making function of trading with an emphasis on keeping position holdings to a minimum, but not to the position-taking function of trading, where positions may be held for very long time periods based on fundamental views of where market prices are headed (refer to Section 2.5 for the distinction between position taking and market making). Is it really important to measure short-term price fluctuations in positions being held for the long term? In this context, it is interesting to note an SEC letter (December 8, 1999) that emphasized the obligation of mutual funds to value assets based on *fair value*, the amount an arm's-length buyer would currently pay for a security. The SEC letter specifically states that fair value cannot be based on "what a buyer might pay at some later time, such as when the market ultimately recognizes the security's true value as currently perceived by the portfolio manager" or "prices that are not achievable on a current basis on the belief that the fund would not currently need to sell these securities." These views reflect the G-30 principles.

Arguments for applying current market valuation and short-term price exposure measures to positions being held for the long term include:

- The desire to hold positions for the long term may reflect the motivation of fund or proprietary position managers, but they may not be the only constituency for valuation information on the fund. Fund investors, lenders to the fund, senior managers of the firm of which the proprietary position managers are a part, and regulators may all have an interest in knowing prices at which the positions may be exited in the near future. Investors may want to exit the fund. Lenders may need assurance that margin calls can be met. Senior managers could decide that they want to reduce the amount of risk-taking authority being allocated to the position takers. Senior firm management will also want to view

integrated risk reports for the entire firm, which will cover both market making and proprietary positioning functions. Regulators may be seeking assurance that fund withdrawals can take place in an orderly manner. All these points were particularly emphasized by the Long-Term Capital Management (LTCM) experience discussed in Section 4.2.1.

- It is possible to find anecdotal evidence of successful fund managers and proprietary traders who do not desire any feedback from market price changes. They view themselves as investing for the long run, and they see short-term price changes as distracting noise that does not reflect changes in fundamental values, but only short-lived shifts caused by supply and demand imbalances. However, it is possible to counter this with anecdotal evidence of successful fund managers and proprietary traders who want to receive constant feedback from the market. Even though they are investing for the long term, they want to be constantly aware of the price at which risk positions can be unwound. They attempt to make money by having a few positions on which they are right and earn a large amount, and avoid having any positions on which they lose a large amount. The constant feedback of market prices at which positions can be exited provides both a means to ensure that a limit is placed on the amount that may be lost on any one position and a signal that markets are moving in ways they do not fully understand. In such circumstances, they seek to exit the market and wait until they can gain a better understanding before reentry.

An alternative but related argument would be that fund managers do not need to be concerned with tail risk but only with the trade-off between expected return and standard deviation of return (the Sharpe ratio), since prudent investors will utilize a fund as just one small part of their overall investments and that it is up to each investor to manage individual tail risk. This is essentially the argument we considered in Section 4.2.1 on the LTCM disaster: “Nor is there a major difference in consequences between bankruptcy and a large loss short of bankruptcy for an investment fund. It shouldn’t matter to investors whether a fund in which they have invested \$10 million goes bankrupt or a fund in which they have invested \$30 million loses a third of its value.” And, as we saw in Section 1.3, if all we need to be concerned about is the Sharpe ratio, many of the elements of financial risk management, such as the inclusion of subjective judgment, are not as strongly needed.

While there is some truth to this argument, it also has some deficiencies. Firms that have credit risk to the fund, often through counterparty risk on derivatives, may care very much about a fund’s tail risk. Regulators are showing increasing concern about potential destabilizing effects of

investment fund bankruptcy. And investors may be concerned that they are not receiving adequate return for tail risks the fund is taking, if these tail risks are unmonitored. There is a strong argument for at least measuring a fund's tail risks, even if it is a greater tolerance for an investment fund taking recognized and adequately compensated tail risks. Thus we are seeing more use of financial risk management techniques such as VaR and stress testing for investment funds (for example, Duc and Schorderet 2008).

Investment funds in which investors are expected to have concentrated risk, most particularly pension funds, are certainly expected to be concerned with tail risk. This is especially true when they pursue a strategy that explicitly depends on liquidity. A good example is contingent immunization, a strategy developed by Leibowitz and Weinberger (1981, 1982, 1983). In contingent immunization, you constantly monitor how much excess you have in fund assets relative to the amount needed to invest in perfectly safe assets, such as zero coupon Treasury bonds, to meet the minimum payout requirements of the fund. As long as an excess exists, fund managers are free to engage in any investment strategy they wish and still be able to assure meeting the minimum payout requirements—as soon as the surplus reduces to zero, they would switch the fund assets into the safe portfolio. But calculation of this surplus requires constant monitoring, both as to the amount of safe assets needed to meet the minimum payout, which will change as zero coupon Treasury rates fluctuate, and as to the liquidation value of the current portfolio in the event this switch needs to be made. Constant calculation of liquidation value requires all of the machinery of financial risk management: accurate, independent, and continuous marking to market; VaR and stress test computations to assess potential loss in liquidation; and valuation reserves against illiquid positions. O'Kane (2008, Section 22.2) has a good discussion of the gap risk in constant proportional portfolio insurance, a product whose design is very similar to a contingent immunization strategy.

6.2 RISK CONTROL

Once an adequate measurement of risk is available, the next logical question is how to control it. Two fundamental and complementary approaches are available. The first is for higher levels of management to place detailed limits on the amount and type of risk that lower levels of management can take—limits on VaR, position size, vega, gamma, and so on. The second is for higher levels of management to provide incentives to lower levels of management to optimize the trade-off between return and risk. The latter approach, based on incentives, gives lower levels of management, which are

closer to the information required to make informed trade-off decisions, the flexibility to find combinations of risks that can maximize the return for a total risk level approved by senior management. However, the incentive approach can also lead to unacceptable risks in the aggregate if too many traders decide to take a similar position, pointing toward a mixed use of both approaches. This is the pattern that can be found at almost all investment banks and will be the approach followed in this book for discussions of control techniques for specific risk classes.

The most extreme form of an incentive-based approach is to restrict controls to assigning to each trading desk a maximum amount of trading losses the traders will be allowed to take before their positions are closed out. This gives the trading desk maximum flexibility in deciding what positions to put on and gives complete freedom as long as unacceptable losses are avoided. Everyone concerned—the traders, senior managers, and risk managers—can agree on such stop-loss limits as a bare minimum for risk control. If all positions could be instantaneously liquidated at any time at the values reflected on the firm's books, it could be argued that this is an adequate limit structure. However, there have been too many instances where a trader has built up a large risk exposure that proved costly to exit when management decided to stop out losses. The time that traders exceed loss limits is often also the time when markets are moving wildly, decreasing liquidity and subjecting positions to large P&L moves even when closeout can be accomplished in the relatively short time of a day or two. At least some additional form of risk control is needed.

Historically, added risk controls have most often been quite detailed limits on the sizes of specific exposures that could be taken, with limit sizes closely tied to both the liquidity of exiting the exposure and the degree of management confidence in the trader.

When the VaR measure was first introduced, it was initially seen as a possible supplement for limiting risk. However, soon traders came to see it as a tool for gaining added flexibility, since it treats all risk as fungible in arriving at a single risk number. Since this risk number is a statistical estimate of the loss that could occur during the period in which a position is being closed down, an argument could then be made for using this as the only supplement to a stop-loss limit, allowing control on the loss that can occur after management has decided to close out a risk position without the need to place detailed controls on particular exposures. This control can take the form of a limit on the total VaR exposure that a trading desk can take and/or a measure of risk in a calculation of return on risk or risk-adjusted return that can be used to compare the performance of different trading desks against targets and against one another to decide on compensation, promotion, and continued employment.

The following are arguments favoring an incentive approach, with senior management input reduced to broad measures such as stop-loss limits and VaR limits or risk-return targets, giving great flexibility in deciding on the risk-taking profile to the business:

- An incentive approach enables trading desks to respond quickly to new opportunities without slowing down decision making by needing to make their case to senior management.
- By not restricting a given trading desk to positions in a particular asset class, an incentive approach encourages broad thinking across asset classes, searching for interrelationships.
- When a trading desk is confined to a particular market at a time when there is a shortage of good trading opportunities in that market, traders are often tempted to pursue riskier opportunities in that market as the only hope for earning a bonus. Giving trading desks the flexibility to trade in other markets when the one they specialize in is less promising is a way to avoid this temptation.
- It is less risky to have many traders with the ability to take positions in a given market than to restrict position taking to a single trading desk. In most circumstances, positions taken by one desk will be offset by positions taken by a desk with a different opinion. When enough desks all line up in the same direction to create a sizable net position, it is a good indication of particularly favorable return-on-risk circumstances.

The following are arguments favoring a more detailed limit approach:

- A more detailed limit approach enables management to restrict position taking in a particular market to only those trading desks possessing sufficient knowledge and expertise in the market to be able to make reasoned judgments.
- As a corollary to the previous argument, it forces trading desks to focus their attention on those areas in which they are expert without having this focus distracted by trying to find opportunities in other markets.
- The real danger is that a trading desk that does not have a successful strategy in its primary market can obscure management recognition of this fact by trying to build a profitable trading record in another market. This can be particularly harmful if it helps to perpetuate a Ponzi scheme in which the firm is delayed in recognizing the mispricing of a transaction with long-term consequences. If a desk is allowed to play in another market, it is important to make sure that P&L attribution firmly separates the results for different products for the same reason we have seen that it is important to separate P&L in newly booked

deals from P&L on management of existing deals. As a particular case, if an options trading desk is allowed to take substantial outright positions in the underlying asset, the P&L from underlying positions must be clearly separated from the P&L on management of volatility and convexity risk. Likewise, if an exotic options trading desk is not forced to lay off the substantial part of the vanilla options risk it generates, then the P&L from vanilla options risk must be clearly separated from that on the residual exotic options risk.

- The final argument given in favor of more flexible position taking is actually quite misleading in two directions. First, it underestimates the degree to which opinions can be infectious and create bandwagon effects, particularly among traders who are not experts in a particular market. The risk is that when the trading desk with the most expertise in a particular market puts on a position, other trading desks will pile on to get a piece of the action. As a result, the firm as a whole will wind up with a much larger position than the trading desk with the expertise would have thought prudent. Second, when situations with less certainty arise and trading desks put on offsetting positions, the firm as a whole winds up being arbitrated—it has flat P&L if the market moves in either direction, but must pay a bonus to the winning trading desk in either case. This points to the need for trading management to insist on trading desks utilizing diverse styles to avoid this form of arbitrage, and detailed limits can play an important role in enforcing the diversity of trading style.
- Management may distrust the excessive reliance on statistical measures of risk. Statistics are based on history and may not reflect management judgment about risk. This may be particularly true in markets that tend toward infrequent but large jumps, such as pegged FX rates, which, due to government intervention, may show a long history of very little movement followed by one sharp break. When the government resources are no longer adequate to hold the desired peg, the tendency is for the resulting move to be very large to reflect the market pressures whose reflection in the price has been suppressed by government intervention. In a period when the peg still holds, historically based VaR will show very little risk, but this will not adequately reflect the possibility of a jump move. Instead of historical relationships, VaR can be based on implied volatilities, which reflect a market judgment of future uncertainty, or on management estimates of risk. A more direct approach is to explicitly limit exposures to management-designed stress scenarios.

The issue of whether to permit a trading desk to take positions in instruments outside its primary expertise is not just a question of whether a desk

should be allowed to actively seek such positions. This issue also comes up as a question of whether a desk should be forced to close out positions that result as by-products of its primary product focus.

Consider an FX options market-making desk. Their primary expertise would be on issues such as the proper management of volatility risk. However, outright FX positions arise naturally in the course of its business, as changes in exchange rate levels lead to changing deltas on its option positions. Should the options desk be forced to close out these outright FX positions, leaving the firm's positioning of outright FX to the spot-and-forward FX market makers, the firm's experts at managing these positions? Or should the options desk be allowed to take its own view on these positions? The same arguments, pro and con, that we have presented previously apply here as well, with a particular emphasis on the second argument pro flexibility and the third argument against flexibility.

Those who favor flexibility point to the broader view of economics and the markets that will come from the trading desk looking at its options positions as a whole, rather than trying to break them apart into a position in the underlying and a position in volatility. They will point out that this encourages thinking about correlations between underlying prices and volatility levels that can best be taken advantage of by being able to manage positions in both the underlying and volatility. These are powerful arguments, as discussed in Chapter 11.

Those suspicious of the consequences of flexibility point to cases in which poor pricing of volatility risks and poor management of options positions were delayed in being recognized by profits that came from taking positions in the underlying (perhaps just by copying positions that the primary underlying desk was putting on). This certainly indicates the need to have, at a minimum, risk reporting that clearly breaks out P&L attributable to the underlying position from P&L attributable to volatility positions.

Even if management decides in favor of the less flexible approach with specific limits on options traders taking positions in the underlying, some degree of flexibility should be retained from a pure transactional efficiency viewpoint. For example, if an options trading desk is never allowed any position in underlying assets, it will need to spend too much of its time writing tickets to close out delta shifts arising from underlying price changes and will lose too much of its P&L in bid-ask spreads. (Even if these are only internal and hence not lost to the firm, it will still be demotivating to the traders.)

The arguments we have presented here for options traders and their positions in the underlying apply equally to forward traders and their positions in the spot market, basis traders and their position in legs of the basis, and exotic options traders and their positions in vanilla options that can hedge part of the exotics' risk.

This discussion on risk controls has important implications for the use of risk decomposition techniques throughout the remainder of this book. It explains why I place such a strong emphasis on utilizing risk decomposition to break apart less liquid transactions into constituent parts—usually a more liquid piece and a less liquid residual. Identifying the more liquid constituents enables the separation of P&L attribution and encourages closing out positions with the desk that can create the maximum liquidity for the firm. It also avoids the booking of phantom P&L by having a different valuation technique used for the same position depending on whether it was created directly or created as part of a more complex transaction. Finally, it also avoids the firm's unknowingly building a large position in a particular product. For example, this motivates the use of a formula for vanilla options that does all the pricing and representation of risk in terms of forward prices derived from the trading desk that is the primary market maker in that product (see Chapter 11) and motivates the attempt to price and represent the risk of exotic options to the greatest extent possible as a combination of vanilla options prices derived from the trading desk that is the primary market maker in that product (see Chapter 12).

A closely related question is whether trading books that take positions in a product in which they are not a primary market maker should be forced to do all their transactions through the firm's primary market-making desk for the product. As a concrete example, consider a trading desk specializing in FX options, which will certainly need to transact hedges in underlying spot and forward FX. Should the traders be forced to transact all such hedges with the firm's spot and forward FX trading book, or should they be given the choice of dealing directly in the market?

Note that this issue arises regardless of whether trading limits are used to force the options desk to restrict its outright FX positions to a small size. In either case, the desk will be transacting at least some hedges either internally or with the market.

The arguments for requiring internal hedging are powerful:

- It enables the desk with the greatest expertise and advantage in trading a product to be the one initiating all external transactions.
- It reduces the amount of transaction costs the firm must pay by encouraging trades in opposite directions to be closed out within the firm and enabling internal trades to be crossed with customer transactions. Nothing pleases traders more than to be able to boast of the profits they have made by standing in the middle of trades in opposite directions put on by different desks of a rival firm. Even if positions are not completely offsetting or exactly simultaneous, funneling the trades through a single desk enables that desk to see the total flow of the firm's dealings in the

product. This desk can build on observed patterns of usage to forecast and anticipate flows and minimize transaction costs.

- The use of a common central trading desk forces all desks within the firm to value positions in the same product at a common price. This avoids phantom profits arising from the internal arbitrage that can occur if two desks value their positions in the same trade using different broker quotes or different models. Proper valuation discipline can eliminate this even if a policy of forcing all trades through a single desk is not employed, but this is the easiest mechanism for enforcing this rule.

The argument for permitting several desks to trade the same product directly with other firms is that competition for business will create enough efficiencies to overcome these strong advantages of a common central trading desk. The fear is that creating an internal monopoly in a product will permit the monopolist to try to collect monopoly rents from the other desks trading in the product—that is, to price at excessive bid-ask spreads that will increase profits of the central desk, but decrease the firm's overall profit by discouraging optimal use of the product by other desks. Avoiding this situation may require a difficult internal policing effort (it's not always easy to measure the size of the bid-ask spreads being used, since trades in different directions do not come in simultaneously).

VaR and Stress Testing

In the statement of requirements for robust risk management in Section 6.1.1, the estimation of losses that could result from liquidation of positions figured very prominently. This showed up under the headings “The need for simulation” and “The need to consider periods of reduced liquidity.” The need for simulation, which closely corresponds to the G-30 Recommendation 5, “Measuring Market Risk,” is discussed in detail in this chapter as value at risk (VaR). The need to consider periods of reduced liquidity, which closely corresponds to the G-30 Recommendation 6, “Stress Simulations,” is discussed in this chapter as stress testing.

These two methods for measuring the total risk exposure of a portfolio still need to be supplemented by more detailed nonstatistical risk measures, such as the value of the basis point, delta, or vega, for reasons given in Section 6.2. But measures of total portfolio risk do offer advantages that detailed nonstatistical risk measures do not:

- Nonstatistical measures do not allow senior managers to form conclusions as to which are the largest risks currently facing the firm. It is not possible to meaningfully compare the value of a basis point in two different currencies, since this comparison does not reflect the relative size of potential interest rate moves in the two currencies. Both VaR and stress testing give a measure that combines the size of position and size of potential market move into a potential impact on firm profit and loss (P&L). Moreover, both produce a measure that can compare risks between disparate businesses, such as interest rates and equities.
- Nonstatistical measures do not interact with one another. Should you add up the risks under different measures into some total risk? Clearly this would be wrong because it would ignore the effect of correlation between market factors. Both VaR and stress testing account directly for correlation between market factors.

We will first discuss the methodology of statistical measurement, VaR, and then discuss the methodology for nonstatistical measurement, stress testing.

A book-length treatment of the topics discussed in this chapter is Dowd (2005), which offers a wealth of detail and covers all the methods that I consider best practices in this area. This is a book I recommend highly for those working on implementation of VaR methodology. You will see many references to it in this chapter. What I offer here are the aspects of VaR that are most important for everyone involved with risk management to know, and those methodological considerations for implementation that my experience in the field has shown to be of greatest consequence.

7.1 VAR METHODOLOGY

Strictly speaking, VaR is a measure of the worst loss that can occur at a given confidence level. But the statistical methodology used to determine VaR can also be used to calculate broader measures of the distribution of potential losses. In Section 7.1.1 we'll first look at the methodology for calculating the distribution and in Section 7.1.2 we'll turn to the question of how best to summarize it.

Since statistical risk measures first began to be calculated by financial firms (about 20 years ago), three methods have dominated:

1. Direct measurement of P&L distribution.
2. Calculation of P&L distribution based on historical statistics representing the variance and covariance of market variables and the current size of position exposures to each of these market variables. So if s_i represents the firm's exposure to each market variable, σ_i represents the volatility of each market variable, and $\rho_{i,j}$ represents the correlation coefficient between each pair of market variables, the volatility of overall firm P&L is calculated as:

$$\sqrt{\sum_{i,j} s_i s_j \sigma_i \sigma_j \rho_{i,j}}$$

The P&L distribution can now be calculated from this volatility.

3. Simulation of P&L distributions based on a selected set of possible moves of market variables and the current size of position exposure to each of those market variables. So if s_i represents the firm's exposure to each market variable, $m_{i,j}$ represents the size of move of each market

variable in each considered scenario, and p_j represents the probability assigned to each scenario, with:

$$\sum_j p_j = 1$$

Then the P&L movement in each scenario is calculated by:

$$\sum_i s_i m_{i,j}$$

And the P&L distribution is calculated by multiplying each of these terms by its respective p_j .

We will consider the advantages and disadvantages of each of these three methods.

The direct measurement of P&L distribution is still widely used, as can be seen from the frequent use of histograms of daily P&L distributions published in annual reports of financial firms, of the type illustrated in Figure 7.1. It has the advantage of simplicity of calculation, not having to make any use of models or statistical assumptions. It also has the ability to capture effects of the trading culture, which the other methods do not. For example, does management respond to periods of greater market volatility

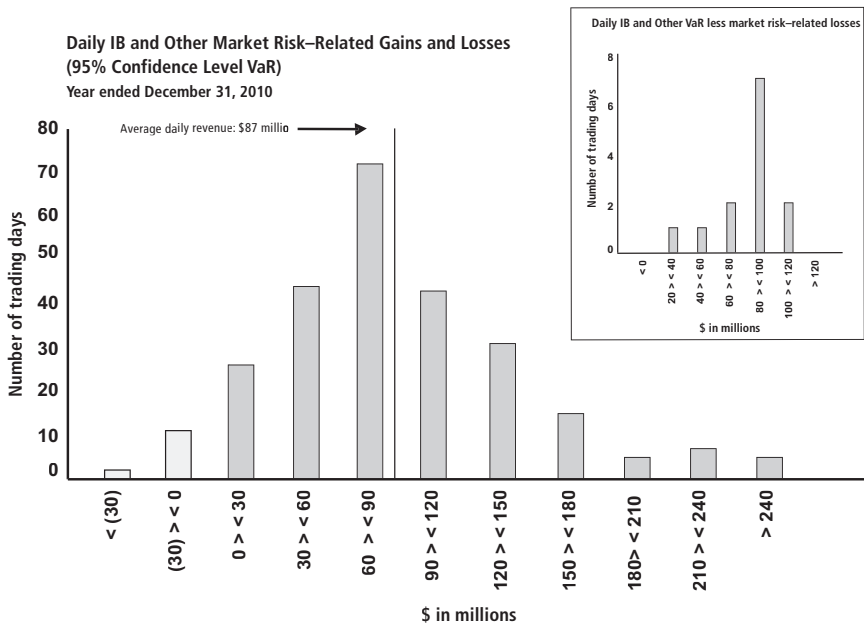


FIGURE 7.1 P&L Histogram from JPMorgan 2011 Annual

by reducing position size? If it does, this will mitigate some of the earnings volatility resulting from market volatility.

Direct measurement of P&L distribution is also the only method that is available for measuring risk when access to details of trading positions is not available. For example, a hedge fund investor probably does not have any access to details of the investment holdings of the hedge fund. To estimate its risk, the investor may need to rely on historical P&L distribution of the fund (for more on risk management of investments in hedge funds, see Section 8.4.1).

However, direct measurement of P&L distributions cannot take into account the possibility that current position taking may be radically different from historical position taking (in the fund management world, this is known as style drift). Corporate risk managers and regulators will insist on risk measures that fully reflect current portfolio composition, whenever available. This renders direct measurement of the P&L distribution close to useless as a stand-alone risk measure, though it is still valuable as a complement to other measures.

The use of the variance-covariance method has now been virtually abandoned by sophisticated financial firms in favor of simulation methods. The primary reason for this is that relative to the simulation method, the variance-covariance method provides very little flexibility in evaluating the contribution of nonlinear positions, notably options positions, to P&L distributions. As we will see, simulation gives the flexibility to tailor the degree of detail used in calculating nonlinear positions to the degree of accuracy required for particular portfolios. Detail can range from simple factor approximations (using delta, gamma, vega, etc.) to full valuation of each individual option, with several gradations available in between. By contrast, variance-covariance can't go beyond factor approximation. Secondary reasons are:

- The greater difficulty that the variance-covariance method has in dealing with the fat-tailed distributions normally encountered in financial markets.
- The inability of variance-covariance to pick up the phenomenon, often observed in financial markets, that the largest changes in variables often cluster together (e.g., the high correlation between stock markets in different countries in the 1987 stock crash) to a greater degree than will be indicated by correlation coefficients (i.e., the joint distribution is not bivariate normal).
- The realization that almost all the benefits of simplicity and speed of computation claimed for variance-covariance relative to simulation were based on fallacious comparisons. As will be seen in our discussion of simulation methodology, the degree of simplicity and speed of computation are largely determined by the choice of the user. Achieving

a level of accuracy similar to that obtained by variance-covariance, simulation is at least as simple and fast to compute as variance-covariance. Simulation offers the flexibility, which variance-covariance does not, of increasing accuracy as a trade-off against simplicity and computation time, but having more flexibility can surely not count as a disadvantage.

Currently, the primary users of variance-covariance are smaller firms that do not hold significant options positions and that wish to outsource the market data component of their VaR computations. For such firms, variance-covariance does offer the distinct advantage that they only need to obtain volatilities and correlations rather than the day-by-day pricing histories required for simulation, a considerable savings in the amount of data to be transferred.

In Exercise 7.1, you will have a chance to see an example of how variance-covariance computes VaR and why a simulation calculation that is as simple computationally and is superior in flexibility is always available. I will therefore not spend any more time on variance-covariance or the various tricks that have been devised to provide capability to approximate option positions and incorporate fat tails within it. For readers who wish to pursue this approach, I recommend Chapters 6 and 10 of Dowd (2005).

7.1.1 Simulation of the P&L Distribution

Remember that the simulation approach consists of determining a number of possible scenarios, to be indexed by j , determining the size of move of each market variable in each scenario $m_{i,j}$, and then calculating:

$$\sum_i s_i m_{i,j}$$

as the firm's total P&L movement in each scenario. The steps in a P&L simulation consist of (1) determining a set of scenarios specified by the size of move in each of a set of underlying market variables and a probability to be assigned to each set and (2) translation from the size of move of underlying market variables to size of move for all market variables. For example, the underlying market variables for a set of bond positions could be interest rates for 10 key tenors, and the full set of market variables could be prices for individual bonds. There are two alternative approaches to the first step—historical simulation and Monte Carlo simulation. The decisions to be made for the second step do not depend on the choice made for the first step. We will discuss each step in some detail.

7.1.1.1 Step 1: Determine Underlying Market Probabilities The historical simulation approach is quite simple; a group of historical periods is chosen and the observed sizes of market moves in each of these historical periods constitute the scenarios. So, for example, you could choose 1,200 scenarios consisting of all the most recent one-business-day changes in market variables—the changes in market variables from 6/7/99 to 6/8/99 would be one scenario, the change from 6/8/99 to 6/9/99 another scenario, and so forth. Or one could choose all the 10-business-day changes.

The most commonly used method for historical simulation assigns equal-probability weights to all of these possible market moves. This makes calculation of VaR very simple, since it is just equivalent to one particular scenario (for example, if you wanted the 99th percentile VaR and you are working with 1,200 scenarios, the 12th worst loss in any of these scenarios is the 99th percentile VaR). When we explain the calculation of measures of P&L distribution in Section 7.1.2, we will discuss the possible advantages of and methodology for assigning unequal probability weights to these market moves.

Historical simulation offers a large advantage in terms of simplicity—simplicity of implementation, simplicity of assumptions, simplicity of explanation. The advantage in terms of assumptions is that no modeling assumption needs to be made beyond the assumption that the immediate future will resemble the past. There is no parameterization of either variance or correlation and no assumptions about distribution shape (e.g., normality). If fat tails or clustering of large moves between variables are present in the historical data, they will be reflected in the simulation.

The advantage in terms of explanation is that any questions raised by traders or managers concerning a VaR that seems too high can be easily traced to a subset of specific historical dates that would show large losses against the current firm holdings. Disagreement can be quickly focused on accuracy of data for a few specific dates or on arguments about the probabilities to be assigned to repetition of particular historical events. By contrast, both the variance-covariance approach and the Monte Carlo simulation approach make it far more difficult to resolve such questions.

This advantage of simplicity of historical simulation also underlies its primary disadvantage—the VaR produced is dominated by market moves on a few specific historical days. If a particular combination of market events did not occur in the historical period being considered, it cannot contribute to VaR. It is difficult to overcome this problem by just expanding the historical period you are considering. Data availability tends to get sparse once you go back more than a few years, because of failure to retain data, because data becomes more difficult to clean the further back you go in time, and because some currently traded instruments may not have histories that go back that far.

This disadvantage of generating scenarios utilizing the historical method is the primary argument in favor of the Monte Carlo method. The Monte Carlo method starts with a specification of the underlying market variables that is similar to that of the variance-covariance approach, but may have a richer specification of each single variable than just a volatility—for example, a multiparameter specification that allows the generation of distributions that are skewed or fat-tailed. Monte Carlo generation of distributions that fit specified parameters can be achieved in several ways:

- By mixing together normal distributions, distributions that are skewed and fat-tailed can be generated. This can be done using the **MixtureOfNormals** spreadsheet we encountered in Chapter 1. Mixing normal distributions with the same mean and different volatilities produces fat tails but no skew. The larger the difference in volatilities, the greater the kurtosis, a measure of how fat-tailed the distribution is. Mixing normal distributions with different means and different volatilities produces both fat tails and skew. More detail can be found in Dowd (2005, Section 6.5.3) and Wang (2001).
- By using stochastic volatility and jump process specifications, similar to those we discuss in Sections 11.6.2 and 12.3.2. See also Hull (2012, Sections 26.1 and 26.2) and Dowd (2005, Sections 6.5.4 and 6.5.5).
- By using processes specially designed to generate Monte Carlo distributions that match given skew and kurtosis parameters, such as those discussed in Shaw (1997). Monte Carlo techniques are then used to generate a set of scenarios that fit the desired statistical specifications. Shaw (1997) discusses building Monte Carlo simulations following the algorithm of Ramberg et al. (1979). An implementation of this algorithm can be found on the website for this book (this algorithm on the website is called “Quasifit”).

Usually, users of Monte Carlo simulation want to take advantage of the flexibility it offers to generate many more scenarios than can be practically generated with historical simulation. This has led to the incorrect assertion that Monte Carlo simulation *requires* more scenarios than historical simulation does. Rather, Monte Carlo simulation offers the flexibility of achieving greater accuracy if the greater expense of running more scenarios is justified by the increase in accuracy. Standard computerized techniques for improving the trade-off between accuracy and speed for Monte Carlo simulation can also be employed (e.g., stratified sampling, importance sampling, low-discrepancy sequences; see Hull 2012, Section 20.7; Dowd 2005, Section 8.4; and Jackel 2002, Chapters 8 and 10).

Advantages that Monte Carlo simulation offers are:

- **Ability to select the most suitable technique to estimate each parameter.** Volatilities and correlations can be forecast using statistical techniques such as weighted moving averages and generalized autoregressive conditional heteroscedasticity (GARCH). For a discussion of the most common statistical methods used in forecasting volatilities and correlations, see Hull (2012, Chapter 22), Dowd (2005, Chapter 5), and Jorion (2007, Chapter 9). Statistical methods for adjusting parameters derived from historical data to be more robust, including random matrix theory and shrinkage estimation, can be found in Fabozzi, Focardi, and Kolm (2006, Chapters 8 and 9). Valuable discussion of this topic can also be found in Meucci (2005, Chapter 4). Swensen (2000, Chapter 5) is a valuable approach with less emphasis on statistical methodology and more on economic insight. Where implied volatilities are available, they can be substituted for or blended with statistical measures. (Should implied volatility always be used when available? We'll examine this question at the end of this subsection.) The choice can be separately made for each variable, though you do need to be careful not to generate impossible or implausible combinations of correlation coefficients; for discussion of how to avoid creating impossible correlation matrices, see Dowd (2005, Section 5.3).
- **Ability to select the most relevant data set for estimating each parameter.** You might have 10 years of good historical data for one variable and only two years for another. Historical simulation would force you to use only two years' worth of data for both. Monte Carlo simulation lets you choose the data set individually for each variable. You can also choose the most appropriate weights to assign to different historical periods for each variable, with more discounting of older historical data for some variables than for others. Historical simulation can only utilize a single weighting scheme that applies equally to all variables (see the discussion of this weighting of historical simulation in Section 7.1.2).
- **Ability to select the most relevant data set for estimating different aspects of a single variable.** For example, volatility could be based on recent data or derived from an implied volatility while higher-order parameters of the distribution are estimated from longer data periods. Recent data is often considered a better predictor of near-term future volatility, but shape parameters, such as fatness of tails, are hard to discern from a small data set.
- **Greater flexibility in handling missing data.** Data for individual dates can be missing because a particular market was closed for a holiday or because of errors in data gathering. In fact, all sources of market data,

whether data vendors, brokers, or databases internal to the firm, are notoriously poor in quality and require major data scrubbing efforts. But some data will not have sufficient duplication of sources to scrub successfully and must be regarded as unavailable. Monte Carlo simulation can exclude periods for which a particular data series is missing from the calculation of each individual variable without excluding this period from the calculation of other variables for which the data are available. Historical simulation lacks this flexibility—it must either completely include or completely exclude a particular day's data.

- **Greater flexibility in handling nonsynchronous data.** Correlations observed between variables that are sampled at different times of the day can be highly misleading and result in significant misstatements of risk. Monte Carlo simulation has the flexibility to measure correlation for each individual pair of variables based on quotations from the best time of day to represent that particular pair, or by basing the correlation on a multiday time interval, which will tend to smooth out nonsynchronous effects. For more detail on statistical methods that can be used in estimating correlations between nonsynchronous data, see RiskMetrics Group (1996, Section 8.5) and Holton (2003, Section 6.3).
- **Ability to combine histories.** Consider a corporate bond held in the firm's portfolio. By historical experience, one knows that some of these bonds may suffer a ratings downgrade and subsequent large fall in price. But it may be that none of the bonds currently held has suffered such a downgrade since the firm avoids holding such bonds. Historical simulation would show no ratings downgrade events for these bonds. But Monte Carlo simulation could be used to combine ratings downgrade possibilities based on the history of a large pool of bonds with specific pricing history of actual bonds held.

Another example would be a foreign exchange (FX) position held in a currency that has been pegged at a fixed exchange rate to the dollar by government intervention. You may have no historical example of this particular currency devaluing, yet want to include some probability of devaluation. Monte Carlo simulation could incorporate a devaluation event, possibly parameterized by devaluation experience in other currencies, as a jump process superimposed on the specific history of this FX rate.

Still another example would be two stocks that have begun trading in a very tightly related fashion since a merger announcement. You would not want to reflect their previous more volatile arrangement as part of the history that determines VaR. So you must generate the price of one stock as a function of the other, but with a random element introduced to represent the risk of a sharp break in the price relationship

if the merger fails to go through. This random element should be based on the price history of a large pool of stock pairs following a merger announcement.

Finally, Monte Carlo simulation allows users great flexibility in deciding on the most effective approach to specifying each individual variable. Consider as an example specifying parameters for credit default swaps (CDSs) (see Section 13.1.1.2 for more details on CDSs). CDS prices for some corporations may have sufficient liquidity that you want to estimate the parameters for this price based solely on the price history of this particular CDS. For other corporations with less liquidity in CDS prices, you can choose to break the CDS price up into a credit spread on a bond issued by that corporation plus a spread between the CDS spread and the bond's credit spread (we'll call this the *CDS basis*). The parameters for the credit spread on the bond might be based solely on the history of credit spreads for this corporation's bonds, while the parameters for the CDS basis might be better estimated from observations of CDS basis history drawn from a larger universe of similar corporations. Even though you are estimating the CDS basis for several different corporations from the same data source, you don't expect them to be perfectly correlated, but you can estimate a correlation coefficient from historical observations of how changes in CDS basis differ between corporations.

It is straightforward to reproduce any desired correlation matrix in a Monte Carlo simulation using the Cholesky decomposition method described in Dowd (2005, Section 8.3). But covariance matrices employ correlations based on multivariate normal distributions and therefore do not capture any relationships that are extremely unlikely under this hypothesis (e.g., the clustering of large changes in variables). Addressing these concerns requires more refined data analyses. For example, different correlation matrices could be used depending on the size of price moves (see Kim and Finger 2000). Days in which price moves are larger would use a correlation matrix derived from a sample of days with large moves. The **MixtureOfNormals** spreadsheet can produce correlations with different degrees of clustering, as shown in Figures 7.2 and 7.3.

Figure 7.2 shows the correlation between two normally distributed variables with 25% correlation, both with mean 2%, standard deviation 5%. Figure 7.3 shows a mixture of 95% of the first distribution and 5% of two normally distributed variables with 60% correlation, both with mean 0, standard deviation 10%. Note the clustering of points with large losses in both variables and large gains in both variables in Figure 7.3. This clustering does not appear in Figure 7.2, which displays a multivariate normal distribution.

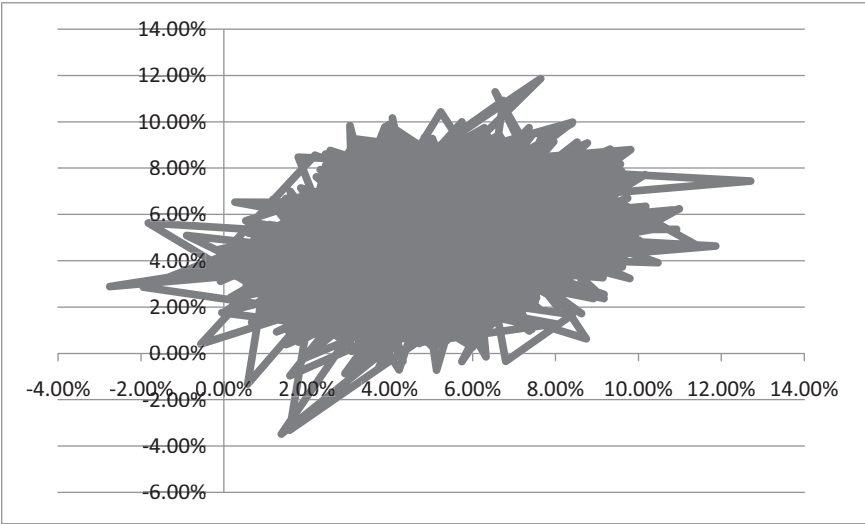


FIGURE 7.2 Correlation Between Two Normally Distributed Variables

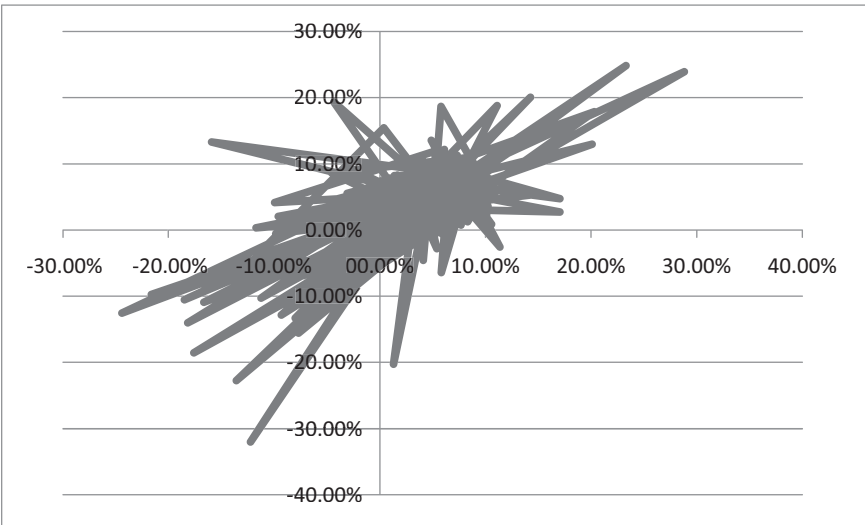


FIGURE 7.3 A Mixture of Two Normal Distributions Shows a Clustering of Points with Large Gains and Large Losses

More general methods for analyzing nonlinear correlations and generating Monte Carlo distributions based on this analysis have been widely studied in recent years. This is known as *copula* methodology; see Dowd (2005, Section 6.8) for details.

Given all these advantages to Monte Carlo simulation in its flexibility to handle data and estimation issues, it is preferable, and sometimes even unavoidable, to still employ some Monte Carlo simulation techniques when you have chosen historical simulation as your primary methodology. Consider these examples:

- A certain stock held in your portfolio has only recently been issued. To develop a past history for the price of this stock for use in historical simulation, you may represent it by some formula based on a selected stock index. But if you are long this stock and short this index, you would measure your position as having no risk during the period when it is represented by the index. To avoid this, you need to introduce a random element into your generation of the stock's back price history, basing the size of the random element on observed changes during the period since the stock began trading. But this is precisely the Monte Carlo approach.
- The ratings downgrade risk case, the FX devaluation risk case, the merger arbitrage risk case, and the CDS basis case discussed in the bullet point headed "Ability to combine histories" under advantages of Monte Carlo simulation are good examples of where a random element needs to be introduced into the historical series.

In cases like these, how should a random element be introduced into historical simulation? One method that is sometimes used is to randomly assign the distribution of the random element among the days of historical data. For example, if there is a 1 in 250 chance of a ratings downgrade for a bond, the price drop that would result from a downgrade would be randomly assigned to four days out of a 1,000-day historical simulation. But this has a large chance of having no impact on the VaR measurement since the four days randomly selected would likely miss the days of largest losses that determine the VaR measure, but a small probability of having a large effect if it happens that one of the four days selected at random corresponds to one of the days of largest loss that determine the VaR measure.

I believe that this randomness in contributing to VaR contributes nothing to the accurate measurement of risk. It is far better to simply accept that this part of the VaR measurement *must* be performed by Monte Carlo simulation, even if you have chosen to do the bulk of your VaR measurement by historical simulation. The historical simulation results for the main body of the portfolio are treated as a single series as input to the Monte Carlo

simulation, with a uniform distribution assigning an equal probability to each day's simulated result (for example, if there are 1,000 simulated days in the historical simulation, each path in the Monte Carlo simulation has a 1 in 1,000 chance of picking each of these days). The elements that cannot be treated by historical simulation would be the remaining series in the Monte Carlo simulation, parameterized as discussed in the section on Monte Carlo simulation. Correlations between factors will be chosen based on best historical evidence and economic intuition.

There are other areas in which historical simulation can usefully borrow Monte Carlo simulation techniques. For example, historical simulations can be modified to choose a volatility for a particular instrument based on any of the techniques mentioned in the first bullet point under advantages of Monte Carlo simulation. All that is required is to multiply each historical observation for the instrument by the ratio between the desired volatility and the volatility over the historical period. This transformation leaves all shape characteristics of the historical distribution, such as fatness of tails and correlation structure, intact. This approach is illustrated in the VaR spreadsheet, using the volatility override input explained in the documentation for the historical simulation portion of the spreadsheet. Dowd (2005, Section 4.4.2) outlines a similar idea.

When it comes to dealing with missing or nonsynchronous data, the options for historical simulation are very limited. Some way needs to be identified for modifying data *before* it is input into the simulation.

For missing data, some type of statistical inference must be used to arrive at a most likely value for the missing data based on the last prior good data point, the next following good data point, and good data points for related data series (for example, if data is missing for an interest rate for a two-year tenor, related data series would be interest rates for the one-year tenor and the three-year tenor). The simplest methods involve averaging between the last prior good point and next following good point, but overlook valuable information from other data series. At a minimum, one should modify simple averaging to follow the pattern of change that took place in related data series between the last prior good data point and next following good data point. Possibly, more advanced statistical models could be used.

For nonsynchronous data, a new data series should be generated of "most likely" synchronous values. For example, suppose that you have available data for closing prices for some stock issues as of Tokyo close of business (COB) and other stock issues as of New York COB. You need to find some series that can bridge the time gap—perhaps a futures contract on a Japanese index that trades in the New York time zone. Then all of the stock quotes as of Tokyo COB can be adjusted for the movement that took place in the Japanese stock index between Tokyo COB and New York

COB, generating a series that approximates what the quotes for these stocks would be as of New York COB. This obviously leaves room for error in estimating where true liquidation of positions will take place, but it is the best you can do if you are not utilizing Monte Carlo simulation.

Just as we can modify historical simulation to include some of the advantages of Monte Carlo, we might want to modify Monte Carlo to include some of the advantages of historical simulation. Beyond simplicity, the primary advantage of historical simulation is the more refined way in which it handles multivariate correlation. By utilizing actual daily simultaneous price moves across the set of all relevant market variables, nonlinear impacts of arbitrarily great complexity are directly incorporated. This points toward a modification of Monte Carlo in which all individual variables are generated by standard Monte Carlo techniques and all correlations between variables are based on historical simulation. This approach, roughly following Shaw (1997), works as follows.

First you perform a standard historical simulation, with equal probabilities assigned to each day's history. Then, each individual variable is re-generated using a Monte Carlo method based on whatever estimation technique is considered most appropriate (e.g., GARCH, implied volatilities, multiparameter specification). Different methods can be individually tailored to different variables. The use of the historical simulation values is to determine which values of the variables occur simultaneously, based on rank order. On the website for this book, you will find my implementation of this procedure in MATLAB, titled "Reorder."

For example, suppose you have a historical simulation with 850 days. Monte Carlo simulation is used to generate 850 values of each variable. If a particular historical data set consisted of the fourth highest value of variable 1, 38th highest value of variable 2, 625th highest value of variable 3, and so on, then you would create a simulation instance with the fourth highest value of the 850 Monte Carlo simulations of variable 1, 38th highest value of the 850 Monte Carlo simulations of variable 2, 625th highest value of the Monte Carlo simulations of variable 3, and so on.

While this approach retains many of the advantages of Monte Carlo simulation, it cannot incorporate them all. It lacks the flexibility to base some correlations on one data set and other correlations on another data set. It requires complete data for every variable in every day to be included in the dates determining the correlation structure. And it has the same problems as historical simulation with nonsynchronous data.

Finally, let's examine the question of whether implied volatility should always be preferred to historical volatility when it is available. In Chapters 11 and 12 of this book, "Managing Vanilla Options Risk" and "Managing Exotic Options Risk," I argue strongly for always valuing

options at volatilities implied from liquid market prices. But this is a pricing argument—we need to determine prices at which longer-term volatility risk can be exited in order to convert longer-term risks into shorter-term risks. VaR is already dealing with shorter-term risks, usually overnight. It is also doubtful that there are liquid prices for options to manage risk over such short time periods. Implied volatilities can be used as indicators of overnight volatility, and there may be arguments for believing they carry a great deal of information. But there are also arguments against giving much weight to implied volatilities—they sometimes have more to do with supply and demand factors than forecasts of price variation, as discussed in Section 11.6.2. The decision must be based on belief about their predictive value, as there is no pricing argument for using them.

7.1.1.2 Step 2: Determine All Market Variables This section discusses various approaches to representation of the firm's exposure to market variables. More details for specific positions can be found in Sections 9.2, 9.3, 9.4, 10.4, 11.4, and 13.1.3. Dowd (2005, Chapter 12) is also a good source for recommendations on this point. Whatever choices are made for how positions should be stored and represented, the most important point with regard to representation of firm position in VaR and stress test calculations is the need for basing all position inputs on data that is entered and controlled by support staff independent of the front office (as per Section 3.1.1). The VaR and stress test reports are key elements for controlling and managing the firm's risk, and it is just as important for position information feeding these systems to be immune from front-office manipulation as it is to have independent P&L reporting. Indeed, some of the most recent major frauds by rogue traders have been perpetrated primarily through manipulation of risk reporting rather than through the manipulation of P&L reporting (as per Section 3.1.1). Falsification of P&L reporting can lead to stop-loss limits being missed, but the equally important limits on buildup of large positions that can be very costly to liquidate depend on VaR and stress test reports.

Nearly as important is to have checks in place to ensure that the position information that feeds P&L calculations and desk-level risk reports is identical to the position information that feeds VaR and stress test reports. Not only is this a vital check on accuracy of risk reporting, but it is also needed to maintain good dialogue between risk managers and front-office personnel. Nothing is as destructive of good dialogue as VaR or stress limit violations that make no sense to traders because they contradict desk-level risk reports.

When computing VaR for spot positions, the translation from underlying market variables to the full set of market variables that you want to

multiply by the firm's positions is quite direct. Spot positions such as spot FX or the holding of an individual stock or stock index or spot gold or spot oil is just directly multiplied by the generated price change from Step 1.

Computation for forward positions is less straightforward. If you are currently holding a Treasury bill maturing one month from now, you don't want to apply to it the price move you observed for that Treasury bill on a date six months ago, since at that point the Treasury bill had seven months to maturity, and you expect seven-month instruments to demonstrate much larger price changes than one-month instruments. So you want to utilize yield curve parameters as underlying market variables and then multiply those yield curve parameters by the appropriate value of a basis point measure of forward position. This has the important added advantage of not having to separately price each interest rate instrument but instead working with a summary description of the entire position.

Issues are most complex for option positions (in which we include any nonlinear payoff positions). The conceptually simplest and most accurate approach would be to value each individual option separately based on the changes in the underlying market variables of forward price and implied volatility. Even such a simple approach has complications, since for each scenario it is necessary to choose a volatility at which to evaluate the option. This requires deciding which point on the implied volatility surface is the right one to apply. Suppose you are repricing an option with one year to expiry, a strike price of 100, and current underlying price of 80. Which implied volatility shift do you use when sampling from a period six months ago when the underlying price was 100? Most practitioners would opt for looking at the shift in options with a one-year expiry and a strike of 125, since that would give the same "moneyness" (i.e., a strike 25 percent above current spot). But this is clearly open to interpretation and a variety of theories on what drives options pricing (see Derman 1999). Very similar considerations apply to option-adjusted spreads on mortgage and mortgage-backed securities, which should be related to the security that had a comparable relationship to the prevailing new mortgage rate. The reasoning is similar, since option-adjusted spreads represent the market pricing of uncertainty in options exercised by homeowners.

While the simplest approach is the most accurate, it is clearly also the most costly, and the heavy expense of doing full individual revaluation of each option position is what was primarily responsible for incorrect claims that simulation methodology for VaR was inherently expensive to perform. In fact, simulation methodology can achieve better accuracy than variance-covariance at no greater cost by the easy trick of representing option portfolios by summary statistics of deltas, gammas, and vegas and multiplying these by the appropriate price change, half the square of change

in price, and the change in implied volatility, respectively. This simplified representation makes options positions no more computationally difficult for simulation than linear positions. So it is a matter of trade-off in desired accuracy versus cost to be determined for each options position.

There are also intermediate approaches. One that can provide quite accurate approximations is to interpolate results based on a price-vol matrix representation of the options portfolio, as per Section 11.4. If a reasonably detailed price-vol matrix is already being calculated as part of the trading desk's own risk reporting, this is a good way of taking advantage of a large number of full revaluation runs that are already being made (since each bucket of the matrix requires all options in the portfolio to receive a full revaluation) without needless duplication of effort. As we note in Section 11.4, the price-vol matrix can potentially capture all higher-order terms in the Taylor series of both the underlying price and the volatility, as well as cross-terms between them. It will not capture impacts such as nonparallel shifts in volatility surface, so these sensitivities will need to be separately accounted for.

Whatever approximations are used should be tested occasionally against a full revaluation by individual option to see if a finer degree of detail is needed. The scenarios involving the very largest shifts should probably always be evaluated by full revaluation by individual option. This is a form of importance sampling (see Dowd 2005, Section 8.4.3). One possible implementation would be to first use a selected approximation technique to simulate all possible shifts, then focus on the ones that produce the highest P&L changes, which will have the greatest influence on the VaR measure, and recalculate these using full revaluation of each option.

Choices as to whether to work with full revaluation of individual option positions, a price-vol matrix, or summary sensitivity statistics should be solely motivated by trade-offs between computation time and expense versus accuracy. In all cases, the ultimate accuracy of P&L simulations rests on the accuracy of the models the firm uses to value transactions. This is true whether the models are used directly in full revaluation or indirectly in supplying the deltas, gammas, vegas, and price-vol matrices, which are multiplied by positions in simulations or in variance-covariance calculations. Reviews of accuracy of the firm's models should always consider their impact on risk calculations such as VaR and stress tests along with their impact on valuations and limit calculations (see Section 8.2.3).

Another important determinant of the cost of calculating simulations and the cost of storing the data needed as input to these simulations is the degree of detail with which positions and market prices are recorded. At one extreme, it would be foolish not to keep separate prices and positions for each different currency for spot FX—there just are not that many

different currencies, and movements between them can be significant. At the other extreme, it would be equally foolish to store market data on forward rates for all possible tenors (i.e., 365 days \times 30 years). Most of these rates are just being produced by interpolation anyway, so you might as well store just the 20 to 50 liquid rates on the curve that all the others are calculated from. In between, there are trade-off decisions to be made. For example, do you want to track individual histories on every stock you hold, or do you want to keep track of just indexes with individual stocks represented through their betas relative to the index? If you choose the latter approach, then a separate estimate needs to be made of the VaR due to idiosyncratic stock risk.

Finally, we note that some of the determinants of exotic derivative prices are not market variables whose price history can be observed and so are not suitable for inclusion in a VaR analysis. Consider an option on a basket of stocks. The impact of changes in the prices of the stocks and in the implied volatilities of each stock in the basket can be computed and included in the VaR. But there will probably be no liquid market quotations for the implied correlations impacting this option. Analysts are occasionally tempted to substitute changes in historical correlation for unobservable changes in implied correlation. I would argue that this is an error.

If the basket option has three years remaining, you should presumably look at the change from one business day to the next of a change in the three-year historical correlation. But since these two three-year periods will share all but one day at the beginning and end in common, the change in correlation that you will measure must be tiny. We know from experience that implied volatility can change far more rapidly than a similarly computed change in historical volatility, and I do not know of any reason why correlations should behave differently. If, on the other hand, you decided to choose a much shorter period for computing the historical correlation in order to increase the potential size of the change from day to day, how would the choice of period be justified? I believe it is better to acknowledge that such nonmarket observables cannot be included in VaR analyses and that their risks should be accounted for separately through reserves and stress tests, as discussed in detail in Section 6.1.2.

Another factor that some risk managers have been trying to incorporate into VaR is liquidity considerations (see Dowd 2005, Chapter 14). Rather than using overnight price moves to represent each instrument, price moves over a longer period will be used to represent less liquid instruments. If this is not handled carefully, it can result in underrepresentation of illiquid risks. For example, you might have a short position in a very liquid government bond and a smaller long position in a less liquid corporate bond. If you compute VaR based on a one-day move for the government bond and

a two-day move for the corporate bond, this could show less risk than a one-day move for both, since the larger moves for the corporate bond have the same effect in the computation as increasing the size of the position. A better approach is to separately calculate a liquidity penalty, as an add-on to VaR, for the cost of exiting less liquid positions, using a formula similar to that proposed for liquidity reserves in Section 6.1.4.

7.1.2 Measures of the P&L Distribution

Simulation is ideally suited to producing full P&L distributions, since individual cases are simulated and probabilities assigned to each case. While the full distribution can be represented graphically, for example by a histogram like that in Figure 7.1, some type of summary statistics are desirable to convey information succinctly. In practice, the primary focus has been on producing a single summary measure, the percentile loss. For example, the VaR at the 99th percentile would be the amount of loss that will be equaled or exceeded only 1 percent of the time. While less well known, another summary measure that is very useful is the *shortfall VaR*, which is the average loss conditional on being beyond a given percentile. For example, the shortfall VaR at the 99th percentile is the probability-weighted average of all losses greater than the VaR at the 99th percentile.

Computation of both VaR and shortfall VaR at any selected percentile is very direct from a simulation. If we have simulated 1,000 equally probable P&Ls, we only need to sort them. The 990th P&L in the sort is the 99th percentile VaR. the average of the 991st P&L through 1,000th P&L in the sort is the 99th percentile VaR shortfall. The VaR spreadsheet on the book's website demonstrates this calculation for both historical and Monte Carlo simulation.

Despite the VaR measure being better known than the shortfall VaR measure, the latter has several advantages that recommend it as a superior summary statistic. The advantages are:

- Shortfall VaR is sensitive to the entire tail of the distribution, whereas VaR will not change even if there are large increases in some of the losses beyond the cutoff percentile at which the VaR is being measured. This can be quite dangerous if it encourages businesses to tailor products to produce risks that escape the VaR measure by being too far out in the tail.
- In practice, shortfall VaR has proved a more stable measure than VaR in showing less sensitivity to data errors and less day-to-day movement due to seemingly irrelevant changes in input data. Presumably, this is due to a greater tendency to average out the noise in the data.

TABLE 7.1 Negative Portfolio Effects

	Portfolio A	Portfolio B	Combined Portfolio A and B
Third worst case for A	-20 million	+10 million	-10 million
Second worst case for A	-25 million	-17 million	-42 million
First worst case for A	-30 million	-10 million	-40 million
Third worst case for B	-7 million	-20 million	-27 million
Second worst case for B	-10 million	-40 million	-50 million
First worst case for B	+5 million	-60 million	-55 million
99th percentile (third worst case)	-20 million	-20 million	-42 million

- With VaR, apparently negative diversification effects can arise, as shown in Table 7.1, in which the 99th percentile of the combined portfolios, a loss of 42 million, is greater than the sum of the 99th percentile losses in the two separate portfolios, 20 million + 20 million = 40 million. Shortfall VaR never displays negative diversification effects.

Negative portfolio effects are undesirable both from the standpoint of clarity of exposition, when explaining risk measures to managers, and from the standpoint of control structure; even if all units of the firm are within allocated VaR risk limits, the firm itself may be outside its risk limits. Negative portfolio effects are associated with risk measures that have been termed *incoherent* in the terminology of Artzner et al. (1997). By contrast, shortfall VaR and stress scenario measures are coherent and so cannot have negative diversification effects. Dowd (2005, Section 2.3) has a good discussion of coherent risk measures in general and shortfall VaR in particular, though Dowd uses the terminology *expected shortfall* (ES) instead of *shortfall VaR*.

Given these drawbacks of VaR, why has it been so widely adopted as a risk measure? The real question senior managers and regulators would like to ask is “What is the worst loss that can possibly occur?” But this is a question that does not admit a concrete answer, so a confidence interval needs to be specified, which presumably leads to questions like “What is the worst loss that will happen no more than 1 percent of the time?” This is the question to which VaR is the answer. But it seems doubtful that management really wishes to convey indifference to the size of the losses beyond this threshold. And my experience confirms that there is a very real danger that traders and product structurers will interpret a fixed VaR threshold as an invitation to hide risk in the tails—deliberately create positions or design products that result in low-probability risks that are just beyond the

threshold and so show up in VaR reports as having no risk. No risk translates to no risk capital charge, and even a small return on a position that attracts no capital charge can look attractive to some front-office personnel. Such extreme tail risks are often quite illiquid and should, in any case, attract a capital charge on grounds of illiquidity, but sending the right signal through VaR is also constructive.

Based on these considerations, I would recommend shortfall VaR as a more desirable summary statistic. If management or regulators still wish to know the VaR, then I would recommend estimating it by a properly selected shortfall VaR. For example, a good estimate of the 99th percentile VaR is the 97.6th percentile shortfall VaR. The two measures are almost exactly equal for normal distributions, and using the 97.6th percentile shortfall VaR as an estimator provides greater stability, avoids negative diversification effects, and eliminates incentives to hide risk in the tails. The VaR spreadsheet illustrates the estimation of VaR by a properly selected shortfall VaR, as detailed in the documentation for the calculation of the historical simulation VaR.

When Monte Carlo simulation is utilized, all simulation runs are assigned equal probability weights, since any differences in weightings of historical data has already been taken into account in the estimation of input parameters to the simulation. But for historical simulations, if you want to assign different weights to different historical periods, you need to do it at the point at which VaR and shortfall VaR are computed, by considering the probabilities that are assigned to each simulation run. Utilizing different weights for different historical periods in historical simulation can help to overcome one of its least attractive features, the way in which the VaR calculation can shift suddenly when a particularly volatile day leaves the data set. For example, if you are using the past 1,000 days of data for your VaR calculations and June 20, 2010, was a very volatile day, VaR calculations on June 20, 2014, might include that day, and VaR calculations on June 21, 2014, and subsequent days might exclude it. If you assign weights to historical periods with a gradual drop in weights as a date becomes more distant, this shift will take place far more smoothly. Dowd (2005, Section 4.4) has a good discussion of this issue and of a variety of reasonable weighting schemes to consider.

If you want to use simulation results to project possible extreme results (i.e., at very large percentiles), then you need to extrapolate beyond the historical data set. For example, if you want to produce a VaR or shortfall VaR at 99.99 percent, you need to forecast what will happen 1 out of every 10,000 days. But you will almost certainly be working with far less than 10,000 days of historical data. We will discuss later, in Section 7.3, the reasonableness of calculating such extreme measures, but for now, let's see how it can be done if needed.

Extrapolation beyond the historical data set requires statistical tools from extreme value theory (EVT). A very brief summary of the principal EVT techniques most often used in VaR analysis appears in the box.

KEY RESULTS FROM EVT

The results from EVT that are most often used in portfolio risk measurement are estimates for VaR and shortfall VaR at percentiles far out on the tail of the distribution. For example, you can find the formulas for these estimates along with derivations as numbers (6) and (10), respectively, in McNeil (2000). I will state them in slightly altered notation, which is designed to make them easier to utilize in a standard VaR framework.

Let VaR_p and ES_p stand for the VaR and shortfall VaR at any given percentile p . Let u be a percentile at which we can directly measure VaR_u by standard simulation. The formulas are:

$$\begin{aligned}\text{VaR}_p &= \text{VaR}_u + (\beta/\xi)\{[(1-p)/(1-u)]^{-\xi} - 1\} \\ \text{ES}_p &= (\text{VaR}_p + \beta - \xi\text{VaR}_u)/(1 - \xi)\end{aligned}$$

The estimation procedure requires a choice of a base percentile u as well as a choice of the parameters β and ξ . A good discussion of the most frequently used methods for determining these parameters and how much confidence may be placed in the estimation procedure can be found in Diebold, Schuermann, and Stroughair (2000). An example using these formulas can be found in the EVT spreadsheet. Dowd (2005) has a good discussion of the application of EVT methods to portfolio risk measurement in Chapter 7, with derivation of these formulas and examination of parameter estimation in Section 7.2. Dowd, in Section 7.1.2, also provides a shortcut version of EVT that can be used as a first approximation. Schachter (2001) also has a good presentation of this material.

There are many issues with the use of EVT, such as the need to make assumptions that are nearly impossible to test and the difficulty in estimating parameters. But its virtue is that, if such data extrapolations need to be made, it provides a smooth and consistent methodology that is superior to the alternative of extrapolating based on empirical curve fitting. A brief and lively discussion of these issues with plentiful references can be found in Embrechts (2000). As Embrechts indicates in this article, EVT is even more

problematic when used with high-dimensional data, which combines in a nonlinear fashion. This is a good description of VaR of a large firm's portfolio, with options valuation providing the nonlinearity. So direct application of EVT to the VaR measure for the portfolio is highly questionable; for a similar critique of applying EVT to VaR, see the section in Schachter (2001) titled "EVT Is No Panacea Either." More reasonable is application of EVT to each individual input variable in a Monte Carlo simulation, combined with as much structural modeling of correlation as possible. This approach will be discussed in Section 7.2.3.

As with any model, a VaR model needs to have its predictions tested against real results to see if it is sufficiently accurate. This process is sometimes known as *back-testing*, since you are looking back to see how the model would have performed in the recent past. It has been particularly emphasized for VaR models, owing to insistence by regulators that if firms are to be allowed to use internally built models for calculation of regulatory capital, they must be able to demonstrate that the models fit real results. The suggested regulatory back-test is a straightforward comparison between the 99th percentile produced by a VaR model on each day during a specified period (since it is this percentile that determines regulatory capital) and the actual P&L on each day. The model is considered satisfactory (or at least erring acceptably on the side of too much capital) if the number of days on which P&L exceeds the predicted 99th percentile is not statistically significantly greater than 1 percent. While this approach has the virtue of simplicity, it is statistically quite a blunt instrument. Much more information can be extracted by comparing VaR projections to actual results at many different percentiles. More sophisticated methods for back-testing are very well presented in Chapter 15 of Dowd (2005). Chapter 6 of Jorion (2007) also covers some alternative back-testing methods, with particular emphasis on how VaR interacts with the Basel capital rules.

A methodological question is whether to back-test against actual reported P&L or against P&L that has been adjusted for components that the VaR cannot reasonably be expected to pick up. Such components are revenue from newly booked transactions, revenue from intraday or (when running VaR for periods longer than a day) intraperiod trading, and gains or losses due to operational error (e.g., trades incorrectly booked). The argument in favor of using unadjusted P&L in the comparison, besides simplicity of computation, is that these are all real components of P&L that can be quite difficult to identify, so it is better to be aware of the extent to which your model is underpredicting actual reported loss events. An argument in favor of making at least the largest adjustments is that without getting the target data to line up with the forecasting process, you are working with a suboptimal diagnostic tool.

7.2 STRESS TESTING

7.2.1 Overview

As stated in Section 6.1.1, risk assessment must include an evaluation of the potential impact of a period of severely reduced liquidity, *stress tests* for short. There are two fundamental approaches that have been proposed to performing stress tests: reliance on historical data and reliance on economic insight. I will argue that strict reliance on historical data is not a viable option—economic insight must be utilized. But I will also argue that economic insight can usefully be supplemented by historical data.

From a computational standpoint, stress testing is simply another variant of simulation; it just uses a different method to generate the scenarios of underlying market variables. But after that, the other two steps in simulation analysis—translation to all market variables and calculation of firm P&L—can be carried out exactly as per simulation VaR; indeed, the exact same system can be used for both.

As we will see, the use of economic insight requires a great deal of extra effort and introduces a substantial amount of subjective judgment. So why bother departing from statistics? Couldn't we just rely on Monte Carlo simulation based on historical data to generate highly unlikely but still plausible scenarios? The answer is clearly no, for several reasons:

- The distribution of market moves in a crisis event may not resemble the distribution of market moves in normal market circumstances. Experience indicates that you cannot safely assume that market moves in a crisis event simply represent extreme values of ordinary market distributions. In particular, correlations often swing to extreme values in a crisis (Dowd 2005, introduction to Chapter 13). For example, in a flight to quality triggered by a major credit scare, otherwise uncorrelated asset prices may move sharply down at the same time.
- Some scenarios represent such sharp breaks with history that no analysis of past experience can offer a complete story. Economic forecasting based on hard-to-quantify judgment is required. When firms were worried in 1999 about the potential impact on the financial markets of the Y2K systems bug, no purely historical analysis could offer any guidance. When many of the nations of Europe adopted a common currency, a scenario based on the possible collapse of that currency could not be based on any clear historical precedents.
- Some scenarios do not relate to public price observations at all, so cannot be based on historical records of price changes. If a firm has an inventory of options on stock baskets whose pricing depends on

long-term correlations for which no liquid public prices exist, a scenario for a market event that would cause the portfolio to be revalued must be formed based on market knowledge. For example, a wave of mergers might drive up the input level of correlations used in valuations. Both the judgments about how plausible a given level of merger activity might be and how much this might impact the firm's internal valuation policies must be based on the knowledge and experience of individuals.

- Many scenarios require judgment about the impact of large declines in market liquidity that often accompany extreme price moves. Record keeping on price liquidity is extremely sparse relative to record keeping on price levels, so it is doubtful that any such scenario could be constructed based on historical statistics. To deal with this limitation, it is generally necessary to estimate the length of time it will take to liquidate a position in a crisis. Since this bears little resemblance to the time it takes to liquidate a position in normal circumstances, it requires an analysis completely independent from that which goes into VaR calculations.
- Some scenarios focus on the plausibility of contagion (chain reactions of changes in one market spilling over into other markets through investor behavior). An example may be fear that a stock market crash will spur sales of bonds by firms needing to meet margin calls. Refer back to the discussion in connection with Long-Term Capital Management in Section 4.2.1. Such scenarios must be constructed based on knowledge of the current composition of investor portfolios. Historical statistical analysis is likely to be of limited value.
- Some scenarios need to emphasize the interactions among market risk, credit risk, funding liquidity risk, and reputational risk (see Basel Committee on Banking Supervision 2009a, Principles for Banks 10 and 14). Historical data will be of little use here. What is required is economic insight based on thorough examination of previous stress periods and creative thinking about similarities between what has occurred in the past and current economic and institutional circumstances.
- As emphasized in Sections 1.3 and 6.1.1, when attempting to estimate low-probability events, it is important to include subjective judgment. Estimating the impact of infrequent episodes of diminished liquidity is a paradigm of estimating low-probability events. Use of scenarios based on economic insight is a systematic way to ensure that subjective judgment is utilized.

7.2.2 Economic Scenario Stress Tests

The use of economic insight may be necessary for stress testing, but it does pose difficulties.

Working out plausible combinations of the entire set of underlying variables that can impact a large firm's trading position is hard work and requires a lot of attention to detail.

While in principle subjective probability judgments could be used to specify probabilities for scenarios, once we leave the realm of historical distributions, different people are likely to have wide differences in subjective probabilities that are difficult to reconcile. In practice, a standard of plausibility is substituted for one of probability, and plausibility is a very subjective notion. But, however subjective, plausibility must still be insisted upon. Without such a standard, stress testing becomes equivalent to the child's (and childish) game, "Who can name the largest number?" No one ever wins, because one can always be added to the last number. And you can always specify a stress test that is one shade more extreme than the last one specified.

Here are some points that should be considered in scenario generation to try to deal with both the amount of effort involved and the degree of subjectivity.

- One aid is to split the work up between a senior group that determines a global scenario for the most important variables and specialist groups that work out the consequences of that global scenario for less important variables. Global scenarios generally reflect major shifts in economic conditions: a stock market crash, an oil embargo, a series of large credit defaults.
- It is important to be sure that splitting the work among specialist groups does not allow inconsistent relationships to develop in the overall scenario. For example, if one group develops the government bond yield curve and another group develops the AAA-rated corporate bond curve, you don't want there to be any tenors at which the government bond yield is higher than the AAA corporate bond yield. This can be avoided by having the second group develop a curve for the spreads between AAA corporate bonds and government bonds, rather than developing a curve for the absolute level of AAA corporate bond yields. Schachter (2001), in the section on "Implementing Useful Stress Tests," has many valuable suggestions along these lines, including:
 - Using proportional shocks rather than absolute shocks for volatilities, to avoid the possibility of specifying negative volatilities.
 - Specifying shocks to yield curve shape and to volatility surface shape, rather than individual shocks to each interest rate and volatility, to avoid unreasonable shapes.
 - Checking that arbitrage relationships, such as cost of carry relationships between cash and futures prices, are maintained.

- Given the difficulty of developing hypothetical scenarios, it is unreasonable to think that more than a handful (say between 10 and 20) can be in use at any one time. Given all the potential combinations of events in markets, it is important to focus on those possibilities that are most significant to the types of positions your firm generally holds.
- Anchoring the assumptions for the move of a particular variable to the largest move previously observed historically is a good preventative against playing the “Who can name the largest number?” game and overcoming some of the inherent subjectivity. But care should be taken to consider a broad enough range of evidence. For example, if the largest previous daily decline in one country’s broad stock market index has been 10 percent and that of the stock index in another country with a similar level of economic development has been 15 percent, there is a presumption in favor of using 15 percent as a historical worst case for both.
- Acknowledging the need for subjectivity and plausibility rather than probability must never be used as an excuse for just utilizing the opinions of a narrowly drawn group. In fact, subjectivity and plausibility are strong markers of the desirability and necessity of considering a wide range of viewpoints. When you encounter (or, even better, seek out) a view with which you strongly disagree but that is backed by reasonable arguments, you need to take it into account. If you were just producing a most likely scenario or deciding on expected return, you would need to finally rely on your best judgment and not on views you strongly disagreed with. But a search for plausibility must cast a wider net, and you can easily include views you don’t agree with as being improbable but still having a small probability of occurring, and so worthy of consideration when degree of protection is being measured. See Section 5.2.5.7 for a specific illustration.
- The most important choices are always about which variables can plausibly move together, not about the size of moves. History can be some guide, particularly experience in prior large moves; history of statistical correlations is virtually worthless. It is important to consider linkages that are caused by investors as well as linkages caused by economics. For example, consider the correlations experienced between seemingly unrelated markets when Long-Term Capital Management was forced to begin liquidating its holdings. Building in such correlations requires market intelligence on the type of holdings that large institutional players may have accumulated.
- Large moves in variables are closely associated with market illiquidity. The size of variable moves chosen should correspond to moves that occur from the time a liquidity crisis begins to the time it ends; prices

recorded in between these times often have little meaning, since you can't really do any significant size of business at those prices. Since record keeping related to market liquidity is usually sparse, choice of the starting and ending points for a liquidity crisis usually depends on the institutional memory of people involved in the trading business.

- One point of contention between traders on one side and risk managers and regulators on the other side is the assumption that no delta rehedging of options positions will take place during the unfolding of a stress scenario (there is a parallel contention about the same assumption when used for the largest moves seen in VaR simulation). Traders rightly point out that they often have firm rules and limits that would require them to perform a delta hedge when underlying prices move sufficiently. However, the reason that risk managers and regulators often insist on assuming no rehedging is the fear that leak of market liquidity in a crisis will prevent rehedging from being executed successfully.
- Creating linkages between large market moves and related losses due to credit risk, funding liquidity risk, and reputational risk is difficult; for some guidance, Basel Committee on Banking Supervision (2009a) is a good source. A starting point could be an internal database of difficult-to-quantify risk factors, as discussed in Section 8.2.6.5. Particular focus should be on situations in which:
 - Credit exposure (most usually counterparty credit exposure) is highly correlated with market prices, such as stock market levels, interest rate levels, or foreign exchange rates; or credit exposure will be impacted by a change in a counterparty's credit rating (see further discussion in Section 14.3.4).
 - The firm's ability to hold positions through a liquidity crisis may be impacted by actions of the firm's creditors or by changes in accounting treatment.
 - Reputational concerns combined with large market moves may cause the firm to voluntarily take losses on positions for which the firm has no legal responsibility.

It has been my experience that some of the time and effort that goes into the generation of a scenario produces little benefit and may even decrease the value of the results. Too much attention to trying to produce values for every market variable that comprises a particular scenario can be self-defeating. For example, suppose you start with an assumption that there will be a big drop in stock markets globally. Both historical experience with previous stock market crashes and economic insight about the responses of central banks to such events may lead to incorporating a large drop in short-term government bond rates with this event. But historical experience

with previous crashes may show mixed results about the direction of foreign exchange rate changes, and economic insight may not offer clear guidance.

To spend a lot of time arguing over which direction of exchange rate move is more plausible given the main characteristics of the scenario is unproductive. It may actually reduce the value of the scenario by choosing a direction that, given the firm's portfolio, reduces the size of the overall P&L impact when in fact it is just as likely that exchange rates would move in the opposite direction and exacerbate the firm's losses. One possible remedy would be to split the scenario in two: one that has exchange rates going up and one with exchange rates going down. But there may be several such choices to make, and multiplication of scenarios may quickly get out of hand. A better solution is to utilize Monte Carlo simulation on some variables to supplement the scenario analysis of the major variables. For example, the decision could be made that the stress loss would be considered the worst 16th percentile loss (roughly one standard deviation) or all cases that consist of the specified scenario levels for the major variables and a normal VaR-type Monte Carlo simulation of the other variables.

7.2.3 Stress Tests Relying on Historical Data

Supplementing hypothetical scenarios with those developed primarily on historical data is desirable for a few reasons. The intensity of effort that goes into developing a hypothetical scenario limits the number that can be used at any given time, which leaves open the possibility that some plausible large risks have been ignored. While exposures to systematic risk factors, such as a large change in stock market prices or a large shift in interest rate levels, will be captured, large exposures to idiosyncratic risk factors, such as a long position in one set of stocks and a short position in another set of stocks, are likely to show no stress exposure in generated scenarios (review Section 6.1.1 for the definition of systematic and idiosyncratic risk as used here). But such positions are subject to losses in some periods of extreme reduction in liquidity. Also, having a more methodical process in place for searching for plausible extreme events may lessen some of the concern about the subjective nature of scenario generation.

We can distinguish two general approaches to forming hypothetical scenarios based on historical data:

1. A complete replay of a previous stressful event, like the 1987 stock market crash or the 1997 Asian crisis. The fact that such an event has actually occurred is a strong argument for the plausibility of a similar event occurring in the future. While there are always some arguments

along the lines of circumstances having changed so much since the time of the event to make a similar event unlikely, it should be remembered that the standard is plausibility, not probability, so arguments against reoccurrence should be fairly overwhelming in order to rule it out. The simulation process for a prior event is pretty simple: select the proper start and end dates based on when market liquidity was restored, make sure you've stored or have researched the historical values of the market variables, and do some artful creation of values for variables for which you don't have historical values. For example, there was no significant liquid emerging market debt in 1987, so you have to create prices based on how emerging market debt fared in subsequent large stock market downturns.

But even utilizing specific past historical events is very resource intensive in researching the needed historical data, determining appropriate start and end dates, and creating values for some variables, so the number of separate scenarios that can be considered will not be large. Idiosyncratic risk positions, such as the long-short stock position described earlier, will still probably not have their vulnerability to liquidity crises properly measured. This shows the need for some reliance on computation methods, which will be our next topic.

2. Use of a computational approach in which a large number of scenarios is generated. This approach is much closer in spirit to VaR calculations, but focuses on trying to determine large moves outside the range of standard VaR. The rest of this section is devoted to different ideas for implementing this computational approach.

It is not difficult to specify plausible large moves for individual parameters. Often these have already been specified as part of stress scenarios based on economic insight. Even when they haven't, similar techniques to those recommended for economic scenarios can be used, looking at a long run of past historical data, but alert to larger moves that may have occurred for similar variables. This is also a good place to apply the extreme value theory (EVT) techniques outlined in Section 7.1.2, since EVT is most appropriate when applied to individual parameters. The difficult question is how to combine plausible large moves for individual variables into plausible large moves for combinations of variables.

One approach is to use historical data to determine a correlation matrix, apply Monte Carlo simulations to generate a distribution of returns, and establish some probability threshold as a quantitative measure of plausibility. Another approach is to find a more mechanical rule for determining which combinations will be considered plausible. The most popular of these mechanical rules is the "factor-push" methodology, which starts by defining

any possible combination of plausible large moves of individual variables as a plausible large move for the combination of factors.

The major drawback for the Monte Carlo approach is the discomfort many risk managers feel for translating the notion of plausibility into a specific probability threshold. The major drawback of the factor-push methodology is that assuming that all variables make a worst-case type of move simultaneously may strain the limits of what is legitimately considered plausible. And both approaches entail significant computational challenges. In the remainder of this section we look at the specifics of these two approaches, along with some suggested variants, and see how these drawbacks might be mitigated and how the computational challenges might be met. We consider the more mechanical factor-push approach first.

7.2.3.1 Factor-Push Stress Tests Factor-push stress testing involves determining a plausible maximum up move and down move for each variable, and then evaluating all possible combinations of these up and down moves. Those that produce the largest negative P&Ls become plausible stress scenarios. The advantage of this approach is that it investigates a large number of possible scenarios (2^f where f is the number of factors) while requiring decision making or statistical analysis around a small number of inputs, the plausibility ranges for each factor. Dowd (2005, Section 13.3.1) provides a useful analysis of factor-push stress testing.

Two principal criticisms of factor-push methodology have been offered. The first is that it does not follow from each individual factor move being plausible that each combination of these factor moves is plausible. This would be particularly true for closely related factors—it would be totally implausible for the two-year Treasury rate to make its largest plausible up move while the three-year Treasury rate is making its largest plausible down move.

The second criticism of factor-push methodology is that it assumes that worst-case P&L always occurs at the extremes of the factor range. While true for linear products, it may not be true once options are involved (e.g., Dowd's example of a long straddle option position where the greater the move, up or down, the greater the gain, so maximum loss occurs far from the extremes).

The second criticism is easier to overcome than the first. Mechanically, it would be easy to design a Monte Carlo simulation that uniformly takes samples from all possible moves of the individual variables between the agreed plausible up and down extremes. Since all possible combinations of plausible individual moves are regarded as plausible, whatever combination shows up with the worst P&L of all these runs is regarded as a plausible worst case. In practice, this involves a very large number of runs, so a number of methods have been proposed for finding the worst case

with fewer runs under certain conditions; see Dowd (2005, Sections 13.3.2 and 13.3.3) for an introduction to maximum loss optimization and crash metrics and Breuer and Krenn (2000, Sections 2.3.2 and 2.3.3) for implementation details.

Attempts to deal with the criticism that not all combinations of plausible individual moves are plausible combinations has fostered a variety of suggested approaches for selecting some combinations as plausible without relying on probabilities. For example, a simple approach would be to sum up the severity (measured by the percentage of the largest plausible moves) of all individual variable moves and create a boundary on this total beyond which a combination is considered implausible. Approaches along this line are discussed in Breuer and Csiszar (2010).

7.2.3.2 Monte Carlo Stress Tests The alternative approach is to accept the identification of plausibility with some type of probability measure. Part of the resistance to this identification is the idea that correlation relations are totally destroyed in crisis events. But, as pointed out by Kim and Finger (2000), “The well-known tendency of correlations to change abruptly in stress events is no valid argument against the inclusion of correlations in the formulation of plausibility standards. For the plausibility standards can be based on crisis correlations as well as correlations in calmer periods.”

The greater resistance is to how to identify plausibility with a specific probability level. Risk managers have good reason to resist attempts to identify numerical probability estimates with a standard of plausibility; given the lack of historical data to support estimation of such low-probability events, it would be easy to try to override sensible caution by ridiculing the low probabilities of the events that are being guarded against. So I would suggest substituting a standard of “relative plausibility.” For example, suppose that risk managers have agreed to some plausible economic scenarios that, judged by historical data, have a 0.005% chance of occurring (since many economic scenarios are tied to large moves in a single key variable, such as a stock market crash or a spike in oil prices, it is not an unreasonable task to estimate such a probability). Then accept as plausible any losses generated by Monte Carlo simulation that have that degree of probability or greater. No one needs to concede the accuracy of the probability estimate; it is quite probable that historical data leads to severe underestimates of the true probability given the number of times markets are hit with what were declared “once in 10,000 years” events. But we are hypothesizing that events that come out with the same measure of probability based on historical data have roughly similar degrees of plausibility.

Operationally, this methodology works similarly to a Monte Carlo simulation of VaR, with the exception that the parameters for individual

variables has been specified so as to include large plausible moves within the probability range that has been agreed on as the cutoff for plausibility (for example, this is easy to accomplish with a mixture of normal approaches). Correlation matrices are specified based on historical data, probably weighted toward data from crisis periods. Many cases will need to be run in order to be able to make a reasonable estimate of losses at an extreme probability level, so some form of importance sampling will be needed to keep to a reasonable usage of resources, possibly by first using quick estimates for P&L and then making more detailed estimates for only cases that have the highest preliminary loss estimate. A paper by Andrea Rafael that illustrates this methodology can be found on the website for this book.

The advantage of this approach is that it does not have any of the absurd combinations of the factor-push methodology (use of correlation matrices, even ones drawn from crisis periods, won't allow extreme up moves in the two-year Treasury rate along with extreme down moves in the three-year Treasury rate). But all positions will get stressed, including positions like one long in some stocks and short in others, at roughly similar levels of severity. It should produce a loss level as severe as or greater than most of the economic stress scenarios, since the plausibility level is directly derived from these scenarios.

7.3 USES OF OVERALL MEASURES OF FIRM POSITION RISK

In an excellent article, Wilson (1998) distinguishes several possible uses of VaR: preventing embarrassing losses, setting operational risk limits, risk comparability, determination of capital adequacy, and performance measurement (see Section 3.2 of Wilson's article). I will use Wilson's framework, stating my own opinions on the usefulness of both VaR and stress testing for these purposes, and comparing my views to his.

Certainly a major concern that firms have looked to VaR and stress testing to help mitigate is the risk of embarrassing losses such as those discussed in Chapter 4, "Financial Disasters." I would agree with Wilson that many of those disasters are due to issues of improper controls (e.g., Barings, Allied Irish Bank) or improper valuation (e.g., Kidder Peabody, UBS) that cannot be controlled by VaR or stress testing. Improper controls and valuation lead to positions being incorrectly reported, and VaR and stress testing cannot overcome issues of deliberate or inadvertent errors in input. If you look at the disasters covered in Chapter 4, only two resulted from unexpectedly large market moves interacting with correctly reported positions: Long-Term Capital Management and Metallgesellschaft. Even for cases like these, I share Wilson's skepticism about the usefulness of standard VaR as

a controlling mechanism since market moves that cause losses of sufficient size to threaten a firm's stability are generally radical departures from recent historical experience.

This still leaves the possibility of using stress testing or an extreme value version of VaR as a good controlling mechanism for those embarrassing losses that are based on large market moves. For the reasons I have given in Section 7.2, I believe stress tests based on economic insight are far more likely than statistical methods to produce useful measures for controlling extreme market moves.

When it comes to risk comparability, both VaR and stress offer the advantages I emphasized at the beginning of this chapter—allowing meaningful comparison and aggregation between different businesses. As Wilson states, traditional risk measures, such as value of a basis point or vega, “provide little guidance when trying to interpret the relative importance of each individual risk factor to the portfolio's bottom line or for aggregating the different risk categories to a business unit or institution level.” The ability that VaR and stress provide to make such comparisons and aggregation, Wilson says,

correctly allows an institution to gain a deeper understanding of the relative importance of its different risk positions and to gauge better its aggregate risk exposure relative to its aggregate risk appetite. VaR accomplishes these objectives by defining a common metric that can be applied universally across all risk positions or portfolios: the maximum possible loss within a known confidence interval over a given holding period. Besides being able to be applied universally across all risk categories, including market, credit, operational, and insurance risks, this metric is also expressed in units that are (or should be) meaningful at all levels of management: dollars (or pounds, francs, etc.). It therefore serves as a relevant focal point for discussing risks at all levels within the institution, creating a risk dialogue and culture that is otherwise difficult to achieve given the otherwise technical nature of the issues.

Wilson's words on this issue square very closely with my own experience. From the very first VaR runs and stress test runs our risk management group performed for Chase Manhattan, management interest was as strong or stronger in what they revealed about the relative risk of individual positions as it was in the measurement of total firm risk. Of particular interest were positions that management had regarded as relatively insignificant contributors to the firm's risk that showed up as among the largest contributors to VaR and stress tests—small absolute position size was outweighed by large price volatility. It's the ability of VaR and stress tests to combine position size, price volatility, and correlation with the rest of the firm's portfolio

into a single measure, comparable across all business lines, that makes them valuable tools in conveying risk information to management.

This information on relative risk of positions has many potential uses. It can provide input for management discussions with trading desks on the proper size of stop-loss limits. It identifies business lines and positions that require extra management attention. It can be used in calculations of risk versus return in performance measurement. When there is a need to reduce risk because limits are being breached, it helps identify actions that will have the quickest impact.

Given the importance of reports on the contributions of risk positions to VaR and stress tests, careful attention to the design of these reports will have large payoffs in better management processes and in appreciation of the value of the risk function. The classic work in this area remains the Goldman Sachs report “Hot Spots and Hedges” by Litterman (1997a, 1997b). Section 11.2.2 of Dowd (2005) provides a succinct précis of these ideas. Here is my take on the main points to consider:

- Reporting is needed for several different types of decomposition—business lines and trading desks for performance measurement, trading positions that may go across trading desks for understanding of the firm’s risk structure, and to identify targets for risk reduction.
- Reporting needs to be able to accommodate both organization structure and highlighting of critical risks. Some reports will need to be organized in a hierarchal fashion, so that reporting matches the way management is used to thinking of the businesses. But other reports should be organized in a largest to smallest risk fashion to be sure that there is sufficient awareness of the largest risks and to facilitate risk reduction.
- All reports should be designed with drill-down capability, so that risks that need extra attention can be further broken down.
- Ability to take quick actions to reduce risk and management understanding of risk are both enhanced by reporting risks using categorization that is meaningful to businesses and to management. The same guidance that will be given in Sections 9.2, 9.3, 9.4, 10.4, 11.4, and 13.1 for informative reporting of nonstatistical positions should be followed here. For example, VaR and stress test risk of interest rate positions should be reported by exposure to parallel shifts of the yield curve and exposure to changes in steepness of the curve as in Section 11.4.
- Design of optimization procedures to identify small portfolios of a few instruments that can replicate a large portion of the VaR or stress test risk is useful both as a design for a quick hedge and as a way to convey an intuitive understanding of the major components of the firm’s position.

In reporting the contribution of product lines, trading desks, and risk components to overall firm risk, several approaches must be considered:

- Each component can be represented by the scenario risk measure it would have as a stand-alone portfolio. This is the easiest approach to implement and certainly gives a good indicator of relative risk, but fails to capture any correlation effects with other risk components that contribute to overall firm risk.
- Each component can be represented by the impact on total firm risk the full elimination of that risk component would have. This captures correlation effects, but may be unrealistic in that full elimination of a business line may not be a feasible alternative.
- Each component can be represented by its marginal impact on total firm risk. This captures correlation effects and gives a good measure of the immediate impact on firm risk of adding to or offsetting some of a component's risk, but it is very dependent on the current mixture of risk components. A very risky business line may get represented as having a small contribution to risk just because it has low correlation with the current mix of risk for the firm. It may be best to use a stand-alone risk measure in conjunction with a marginal impact measure to make sure that components that can potentially make large contributions to risk receive timely management focus.

The marginal impact measure has a nice side benefit—when you take the weighted sum of marginal impact, weighted by current positions, you get the total risk measure for the firm. Compare the discussion here with Dowd (2005, Section 11.2.1)—note that Dowd uses the terminology *component VaR* for what I am calling marginal impact. This makes the marginal impact a convenient tool for exercises such as allocation to business line of firm capital where you need the sum of the parts to equal the whole. In order to have this property, a risk measure need only satisfy the condition that it scales directly with position size; that is, a position that has the same composition but is k times as large has a risk measure k times as large as the original position. This homogeneity condition is clearly met by both VaR and stress testing measures.

To see that the weighted by position sum of marginal impacts equals total risk, first write the risk measure of the portfolio as $R(x_1, x_2, \dots, x_n)$ where x_i is a component of the portfolio. By hypothesis, $R(kx_1, kx_2, \dots, kx_n) = kR(x_1, x_2, \dots, x_n)$. Taking the derivative of both sides with respect to k , the left-hand side by the chain rule, we obtain:

$$\sum_i x_i \frac{\partial R(kx_1, kx_2, \dots, kx_n)}{\partial kx_i} = R(x_1, x_2, \dots, x_n)$$

Setting $k = 1$,

$$\sum_i x_i \frac{\partial R(x_1, x_2, \dots, x_n)}{\partial x_i} = R(x_1, x_2, \dots, x_n)$$

which states that the sum of the marginal impacts weighted by position equals total risk.

Given this ability to place different risks on a common footing, it is quite natural to want to place limits on businesses based on VaR and stress scenario losses. Stress scenario losses offer the added benefit of controlling against at least some forms of financial disaster. However, this does not provide a complete solution to control of a trading business, and other (nonstatistical) limits are needed as well. Wilson emphasizes speed of calculation and ease of understanding and communication as the reasons for needing other limits besides VaR. I would emphasize, as in Section 6.2, the need to match position taking to expertise and to assure adequate diversity of trading style.

A supplement to the use of limits to control risk is the provision of an adequate capital cushion against potential losses. This cushion is required for both earnings volatility and market moves. Earnings volatility measurement aligns well with VaR, while the impact of large market moves is a risk better measured by stress scenarios. While I believe this to be a sound argument for basing internal measures of capital adequacy on both VaR and stress loss, regulators have strongly favored VaR as the measure on which to base capital required for regulatory purposes. Since capital required for regulatory purposes can have a direct impact on the firm's stock price performance, regulators have been wary of any tie to a measure such as stress, which directly relies on human judgment, for fear that management will manipulate it. VaR has been viewed as preferable based on the relative difficulty of manipulating a statistical measure. VaR is viewed as at least capturing relative differences in level of risk. Translation into a required capital cushion against large, unexpected moves is then approximated through multiplication by an essentially arbitrary constant. For a more detailed discussion of the regulatory capital standards revolving around VaR, see Chapter 3 of Jorion (2007).

For performance measurement, the critical objective is to have a means of adjusting the P&L performance of the firm and of business units for the level of risk taken in achieving this performance. As with the capital cushion, the risk taken is both a function of earnings volatility and of vulnerability to unexpectedly large market moves, arguing for using a mix of VaR and stress loss in developing this measure. But the subjectivity of stress scenarios, combined with the sole reliance of regulatory capital on VaR, has led almost all firms to the decision to base this risk measure completely on VaR. The

firm where I have worked for the past several years, Chase Manhattan (now JPMorgan Chase), has been very unusual in utilizing both VaR and stress in this measure. I will relate some of the history that led Chase management to conclude that stress loss was worth utilizing despite the disputes between the central risk management group and business units, which are inevitable when experience and judgment are significant determinants of a performance measure.

When the Asian credit crisis of the fall of 1997 started to spread to other emerging market economies, we noticed that the losses being experienced by Chase trading desks very closely matched the projections of the hypothetical flight-to-quality stress scenario we had constructed. The match was not just for the firm as a whole but for individual business units. This experience persuaded management to experiment with tying the risk adjustment of business units relative to stress losses, as an incentive to reduce vulnerability to large market shocks. As business adjusted to the new performance measure in early 1998, we noticed a significant impact in terms of strategies to continue to meet P&L targets with less reliance on positions that were vulnerable to these shocks. The result was that Chase weathered the fall 1998 market shock due to the Russian default and the unraveling of Long-Term Capital Management with much smaller losses than in the fall 1997 crisis and smaller losses than almost all of our largest competitors (see O'Brien 1999). Continued experience with the impact of this decision since then has continued to confirm its value.

The mechanisms for adjusting P&L return for risk, which include calculating risk-adjusted return on capital (RAROC) and shareholder value added (SVA), are not topics addressed in this book. Interested readers are referred to Chapters 20 and 21 of Culp (2001) and Chapter 16 of Jorion (2007).

EXERCISES

7.1 Vaule-at-risk computations

Using the data in the VaR spreadsheet (with equal weights on all days) and a 10 percent position in each of the 10 variables, calculate the 99th percentile VaR using the following five methods:

1. Variance-covariance.
2. Historical simulation using a single-point estimate of the 99th percentile.

3. Historical simulation using $2.33 \times$ the standard deviation of the daily total portfolio valuations as the 99th percentile.
4. Historical simulation using a single point estimator of the 99th percentile and substituting the historical volatility over the most recent 100 business days for the historical volatility over the full data set, but using the full data set to simulate results.
5. A Monte Carlo simulation.

Your answers to 1, 3, and 5 should be very close to equal. Why? What does this tell you about the relative ease of implementation of the three methods?

7.2 Maximizing diversification

Try the same exercise as in 7.1 with a combination of investment percentages that you choose yourself. Can you find a combination without any short positions (all investment percentages positive) that gives a high diversification benefit (cell D24 of the **Var-Cov VaR** worksheet in the **VaR** spreadsheet)?

7.3 Measuring fat tails in historical data

Look at the **Ratios** worksheet in the **VaR** spreadsheet. What does it tell you about how fat tailed the time series used in these calculations is? At what percentile level do you begin to see a significant impact of the fat tails?

7.4 Generating fat tails in Monte Carlo simulations

Experiment with the **MixtureOfNormals** spreadsheet and see how different selections of input parameters produce different degrees of kurtosis and clustering of large changes.

Model Risk

Any book on financial risk management needs to address the subject of *model risk*, the risk that theoretical models used in pricing, trading, hedging, and estimating risk will turn out to produce misleading results. This book, which emphasizes quantitative reasoning in risk management, pays particularly close attention to how models can be used and misused in the risk management process.

Since the publication of the first edition of this book, the financial risk management focus on model risk has intensified. In the wake of the 2007–2008 crisis, as we discuss in Sections 5.1 and 5.2.5.3, there have been accusations that model failure was one of the root causes of the meltdown. When a widely discussed article has the title “Recipe for Disaster: The Formula That Killed Wall Street” (Salmon 2009), it is clear that model risk needs to be addressed with a sense of urgency.

Fortunately, in addition to this sense of crisis surrounding model risk, the past several years have witnessed greater attention to analysis of how model risk can be controlled. Concise, excellent articles by Derman (2001) and Rebonato (2003) are now recognized as touchstones for the analysis of model risk. Morini (2011) is the first thorough book-length treatment of model risk. The Federal Reserve and Office of the Comptroller of the Currency joint document for “Supervisory Guidance on Model Risk Management,” which I will reference as FRB (2011), and the Basel Committee on Banking Supervision’s “Supervisory Guidance for Assessing Banks’ Financial Instrument Fair Value Practices,” which I will reference as Basel (2009b), provide regulatory responses to the lesson of the 2008 events for model risk. I find the joint Federal Reserve/Comptroller of the Currency document to be particularly thorough and persuasive in its analysis of the many aspects of model risk.

This chapter begins, in Section 8.1, with an overview focusing on the variety of opinions that have been expressed about the importance of models, or their unimportance, in managing financial risk. Section 8.2 examines

the procedures that ought to be used for risk evaluation and control for models of all types. The following three sections give a more detailed analysis of model review standards, distinguishing among three types of models: those used for valuation and risk measurement of liquid instruments in Section 8.3, those used for valuation and risk measurement of illiquid instruments in Section 8.4, and those used for making trading decisions in Section 8.5.

8.1 HOW IMPORTANT IS MODEL RISK?

When examining model risk, one immediately encounters a very wide range of views on the role that models can play in controlling risk and creating new risks. These vary all the way from viewing model error as the primary cause of financial risk to viewing models as largely irrelevant to risk.

The view that models are largely irrelevant to risk can often be encountered among traders who view models as just convenient mathematical shorthand with no real meaning. All that really matters are the prices the shorthand stands for. A good example is the yield of bonds as calculated by Securities Industry Association standards. This includes many detailed calculations that have no theoretical justification, but can only be explained historically (for example, some parts of the calculation use linear approximations, which made sense before calculations were done on computers). No one would claim that this yield has a precise meaning—you don't necessarily prefer owning a bond yielding 7 percent to one yielding 6.90 percent. However, you can translate between yield and precise price given the industry standard rules. It is convenient shorthand to convey approximate values. The degree to which these calculations give misleading yields hurts intuitive understanding, but does not result in mispricing.

Those who view models as playing no real role in pricing and risk management view almost all models used in financial firms as playing a similar role to that of bond yield calculation. A typical claim would be that the Black-Scholes option model, probably the model most frequently used in the financial industry, is just a mathematical convenience that provides shorthand for quoting options prices as implied volatilities rather than as cash prices. In this view, implied volatilities are an attractive way of providing quotations, both because of common usage and because they provide more intuitive comparisons than a cash price, but they should not be regarded as having any meaning beyond representing the price that they translate to using the Black-Scholes formula.

If this viewpoint is correct, models would play an extremely minimal role in controlling risk, and model testing would consist of little more

than rote checking to see if industry-standard formulas have been properly implemented. However, this extreme a view cannot explain all the ways in which trading firms use models such as Black-Scholes. The valuation of unquoted options is derived by interpolating the implied volatilities of quoted options. The Black-Scholes model is used to translate prices to implied volatilities for the quoted options and implied volatilities to prices for the unquoted options. The risk reports of position exposures use the Black-Scholes model to compute the expected impact of changes in underlying prices on option prices. Scenario analyses presented to senior management quantify the impact of changes to the implied volatility surface. For more details, see Chapter 11 on managing vanilla options risk. This behavior is inconsistent with a claim that the model is being used purely to provide convenient terminology. By contrast, the industry standard bond yield formulas are not used in comparable calculations—interpolations and risk reports are based on a more sophisticated model of separately discounting the individual cash flows that constitute a bond, with a different yield applied to each cash flow. In this computation, none of the linear approximations of the industry standard formulas are utilized. For more details on these calculations, see Chapter 10 on managing forward risk.

The view that models are the primary cause of financial risk is often encountered in articles describing major trading losses, which are frequently ascribed to the firm having the wrong model. What is often unclear in these claims is whether “having the wrong model” just means making incorrect forecasts about the future direction of market prices or if it means misleading the firm’s traders and managers about the nature of positions being taken. A good illustration is the discussion in Section 4.2.1 of whether the reliance by Long-Term Capital Management (LTCM) on models should be viewed as a primary cause of the collapse of the fund. And after the 2007–2008 crisis in mortgage collateralized debt obligations (CDOs), one began to encounter claims such as “[David] Li’s Gaussian copula formula will go down in history as instrumental in causing the unfathomable losses that brought the world financial system to its knees” (Salmon 2009).

Of course, once products start encountering losses, modelers who had been promoting a view of the importance of models may now wish to take the opposing view. Morini (2011, Preface) quotes modelers, speaking after the crisis, telling him “Models were not a problem. The problem was in the data and the parameters! The problem was in the application!” Morini’s response is that “Models in finance are tools to quantify prices or risks. This includes mathematical relations, a way to use data or judgment to compute the parameters, and indications on how to apply them to practical issues.

Only by taking all these things together can we talk of ‘a model.’ Modellers should stay away from the temptation to reduce models to a set of mathematical functions that can be thought of separately from the way they are specified and from the way they are applied. If this were the case, models would really be only blank mathematical boxes and people would be right to consider them useless, when not outright dangerous.” I would add that any modelers who want to separate their work from choices on data or parameters are basically saying that they are programmers. There’s nothing wrong with being a programmer—it’s a highly demanding profession. But with rare exceptions (a few people who are able to pioneer an extraordinary speedup of existing calculations), programmers are not compensated at the level modelers are and do not have the degree of influence in making decisions about innovations in products that modelers do.

In the final analysis, whether model builders take the responsibility or the traders and risk managers who use them take the responsibility, models play a key role in managing risk and we must develop clear guidelines to see that the role they play is to clarify issues rather than to obscure them. This is the task to which we now turn.

8.2 MODEL RISK EVALUATION AND CONTROL

In this section, we look at those procedures that ought to be used for risk evaluation and control for all types of models—those used for making trading decisions as well as those used for valuation and risk measurement. In Section 8.2.1, we discuss the scope of model review and in 8.2.2 the proper roles and responsibilities that need to be established around model review and control. In Section 8.2.3, we look at those procedures that check whether the model selected has been correctly implemented—whether the model actually performs as specified; Morini (2011) calls this *model verification*. In Sections 8.2.4 and 8.2.5, we examine two particularly important pieces of model verification, the verification that contractual arrangements have been correctly specified in the model and the evaluation of approximations. In Section 8.2.6, we turn to procedures that check whether the model selected is appropriate for the product or trading strategy being modeled; Morini (2011) calls this *model validation*.

The procedures in Sections 8.2.3, 8.2.4, 8.2.5, and 8.2.6 are primarily designed for the initial evaluation of models leading up to the decision whether the model should be approved for use, and what restrictions, if any, should be placed on its use. In 8.2.7 and 8.2.8, we look at those aspects of model evaluation and control that should take place continuously or periodically during the life of the model’s use to see if any new information is

available to change the initial conclusions about the model approval or to suggest model modification or replacement.

8.2.1 Scope of Model Review and Control

The first point that needs to be established is what determines that something is a model that requires review and control. FRB (2011, Section III) casts the net very wide, stating that “For the purposes of this document, the term *model* refers to a quantitative method, system, or approach that applies statistical, economic, financial, or mathematical theories, techniques, and assumptions to process input data into quantitative estimates. . . . Models meeting this definition might be used for analyzing business strategies, informing business decisions, identifying and measuring risks, valuing exposures, instruments or positions, conducting stress testing, assessing adequacy of capital, managing client assets, measuring compliance with internal limits, maintaining the formal control apparatus of the bank, or meeting financial or regulatory reporting requirements and issuing public disclosures.”

It is important that a definition this broad be used. A computation may be made by a simple formula in a spreadsheet and still give rise to as great a danger of incorrect estimation as a computation requiring a complex mathematical derivation and a supercomputer churning away for hours to produce the result. Simply averaging observed two- and three-year interest rates to obtain a two-and-a-half-year interest rate already entails an assumption that requires review and control (we’ll discuss this example further in Section 8.2.6.1). The mental image many of us have of a model as a complex piece of mathematics and computer engineering can create blinders when we are looking for potential sources of model risk.

A second point made in FRB (2011, Section V) is that “Vendor products should nevertheless be incorporated into a bank’s broader model risk management framework following the same principle as applied to in-house models, although the process may be somewhat modified.” Whether a model has been created in-house or by a vendor, the consequences of the model being incorrect still affect the profit and loss (P&L) of the firm using the model, so there should be no variation in the standards applied for model review and control. The Federal Reserve goes on to point out the challenges of reviewing vendor models since they “may not allow full access to computer coding and implementation detail” and there is a need for “contingency plans for instances when the vendor model is no longer available or cannot be supported by the vendor.” The model review procedures of this chapter can be used for vendor models, but I have encountered instances where a vendor model is so opaque that I have needed to insist that it be replaced

by an in-house model or by another vendor model that permitted more transparency.

After establishing the scope of the definition of a model, the next step is to agree on what needs to be included in a review. Some key points from FRB (2011, Section III):

- “Models are of necessity simplified representations of real-world relationships and so can never be perfect.”
- As a result, model use invariably results in model risk, which can be defined as “financial loss, poor business and strategic decision making, or damage to a bank’s reputation” based on “incorrect or misused model outputs and reports.”
- Model risk can result from either fundamental errors in the model or inappropriate use of a model, particularly the use of a model outside the environment for which it was designed.
- “Model risk should be managed like other types of risk. Banks should identify the source of risk and assess the magnitude. Model risk increases with greater model complexity, higher uncertainty about inputs and assumptions, broader use, and larger potential impact.” The intensity and rigor of model reviews need to be matched to the degree of model risk identified.
- Model risk cannot be eliminated, so it needs to be controlled through limits on model use, monitoring of model performance, adjusting or revising models over time, and informed conservatism in inputs, design, and outputs. But while conservatism “can be an effective tool” it cannot be “an excuse to avoid improving models.”

8.2.2 Roles and Responsibilities for Model Review and Control

I have been involved with the design and approval of several firmwide model review policies. In every case, I have insisted on a prominent statement that “Risk management serves as a second set of eyes for model review.” This means that the business unit that develops and utilizes the model has first responsibility for reviewing the model and assessing its risks. The role of the risk management function is very important in ensuring that an independent unit without insider incentives reviews the model and in creating a uniform model review environment throughout the firm. But the knowledge that an independent review will be performed by risk management cannot be used by the business unit as an excuse for not performing its own thorough review. Having two sets of eyes reviewing the model is important both for providing an extra layer of security and in obtaining the benefit of insider

product expertise to complement outsider independence and model review process expertise.

FRB (2011) supports this viewpoint on business unit responsibility. In Section VI it states: “Business units are generally responsible for the model risk associated with their business strategies. The role of model owner involves ultimate accountability for model use and performance. . . . Model owners should be responsible for ensuring that models are properly developed, implemented, and used . . . [and] have undergone appropriate validation and approval processes.”

Just stating that business units have responsibility for model review is only the first step. Incentives need to be properly aligned to make sure this responsibility is taken seriously. Steps to assure this include:

- The business unit responsibility for model review is just as much about clear communication as it is about clear thinking. As Morini (2011, Section 1.4.1) emphasizes strongly: “The choice of a valuation model must be based on an analysis . . . reported to senior management in an *aggregated and understandable form*.” “Quants, traders, and other technically strong practitioners” must find ways to communicate technical ideas in nontechnical language “comprehensible for senior management.” Technically strong practitioners who have difficulty in doing this should seek help from colleagues who have stronger communication skills or from corporate risk management personnel who have more experience in this aspect of model review. But no one should be under the illusion that they will escape responsibility for consequences because “the senior guys just weren’t capable of understanding what we were doing.” If you truly can’t get senior managers to understand the potential consequences, even with renewed effort at clever communications, then this is a product your business unit should *not* be trading.
- A clear distinction should be made between losses due to market uncertainties that were clearly identified and advertised as part of the business unit’s model review and losses due to market uncertainties that were ignored in the model review. The latter should have more serious consequences for performance review and compensation than the former, and this policy should be widely advertised within the firm.
- To make sure that the policy in the previous bullet point is successfully implemented, an analysis of significant trading losses needs to be conducted by control personnel independent of the business unit to determine how losses are related to model reviews.
- The independent model review conducted by the risk management area should include identification of weaknesses in the business unit

model review. Patterns of weaknesses need to be addressed by corrective action, as well as consequences for performance review and compensation.

I have not addressed here the issue of how business unit model responsibility should be divided between model builders and traders. This will have different solutions for different business units and different models. I will just note again my comments in Section 8.1 that any function seeking to shun responsibility for model error should accept that reduced responsibility and reduced compensation opportunities go hand in hand.

This emphasis on business unit accountability for model review is consistent with placing the main responsibility for model development with the business unit and allowing them as much freedom in structuring models as possible. Models need to take advantage of as much inside information, in the form of trader beliefs about the future, as possible. Firms must try to be open to as many trading ideas as possible and not dismiss ideas on the grounds that they do not line up with some approved theory (for example, rational expectations or marketplace efficiency). However, a culling process must also be available for measuring the success of trading ideas and eliminating those ideas that are not proving successful. Insiders should be given latitude in the theories used in deciding how to trade, but not in the theories used in deciding when to recognize P&L. Profits should not be booked and bonuses not be paid out until the forecasts of the trading models have proven correct.

For decisions on when to book P&L, it is better to rely on outsiders to avoid bias. You may lose accuracy by not having access to the insiders' market knowledge, but this will only result in delays in recognizing earnings, which is not as serious a problem as taking the wrong positions. Insiders may object that this delay in recognizing P&L will cause them to turn away good business, but they have two alternatives: find others in the market who share their opinions and sell off the risk recognizing the profits, or, if they are sufficiently confident, wait to recognize the P&L until after the risk position has matured.

The risk management units that are part of control functions and that constitute the independent "second set of eyes" in model review do not have the business units' incentive issues. It is very clear to them that unidentified model risks that lead to trading losses will cost them in compensation and may cost them their jobs and even their careers. The incentive problem here runs in the opposite direction: the negative consequences of approving a model that later proves defective are so clear that there is a danger of playing it safe by creating unreasonably high barriers to model approval. After all, a rejected model will never have a chance to show how it would

have performed, so there might appear to be little danger in being proved wrong by being overly cautious. Fortunately, in most firms, business units will press their case with sufficient passion and sound analysis to overcome such unreasonable barriers to new business. But managers of independent model reviews need to be alert to this temptation toward caution, and need to be constantly challenging model reviewers to make sure the right balance is being struck.

The problems for independent risk managers are more likely to rest with issues of expertise and access than they are with issues of incentive. As outsiders, they have less chance to build up the thorough knowledge of models and markets that business units possess. This can sometimes be addressed by employing former model builders as independent reviewers, but this is a career move that appeals to only a small subset of model builders. Other techniques for trying to overcome this gap in expertise will be addressed throughout the remainder of this chapter and this book. For issues of access, rules need to be put in place and enforced to see that business units are forced to share model code, documentation, and supporting data with independent reviewers. Claims of need for secrecy to protect proprietary model features must be viewed with suspicion—these are often just excuses to try to avoid independent scrutiny. When found legitimate, such claims need to be addressed by controls that restrict access to only those actually directly involved in the independent review; they must *never* be used as a reason to limit the scope of the independent review.

Independent model reviewers need to clearly identify steps that need to be taken when they find issues with models. As FRB (2011, Section VI) states, “Control staff should have the authority to restrict the use of models and monitor any limits on model usage. While they may grant exceptions to typical procedures of model validation on a temporary basis, that authority should be subject to other control mechanisms, such as timelines for completing validation work and limits on model use.” In all cases where follow-up action is called for, there should be definite dates for further review established and a well-organized procedure for making certain that a follow-up review is performed evaluating these follow-up actions.

The role just specified for independent model review is consistent with the guiding principle of FRB (2011, Section III) of “effective challenge” of models—“critical analysis by objective, informed parties who can identify model limitations and assumptions and produce appropriate changes.” Requirements for effective challenge outlined by the Federal Reserve are separation of the challenge from the model development process, knowledge and modeling skills adequate to conduct appropriate analysis and critique, and sufficient “influence to ensure that actions are taken to address model issues.”

To what extent should external resources (i.e., consultants) be used as part of the independent model review process? There are many reasons for wanting to minimize the use of external resources: the desire for confidentiality of proprietary models, the desire to build up in-house expertise through the experience gained by conducting model reviews, and the fears of discontinuity if an external resource becomes unavailable or proves unsatisfactory. There still may be times when use of external resources is desirable, either because of a lack of in-house expertise within the independent model review group or because manpower available is not sufficient to meet the demand of new models needing review. When external resources are utilized, care should be taken that a designated in-house reviewer becomes as familiar as possible with the work of the consultant. This serves the function of acquiring some in-house expertise that can be utilized in subsequent model reviews, as well as having someone in-house who can monitor and coordinate the work of the consultant, provide a point of contact for subsequent discussion of the consultant's work, and be able to ramp up involvement in case of discontinuity. (Compare the discussion in this paragraph with the segment headed "External Resources" in FRB [2011, Section VI]).

Finally, all model review activities, both those of business units and those of independent reviewers, must be properly documented. This covers both documentation of the model itself and of the model review process. Model developers are often anxious to get on to the next project, and business units are anxious to develop the next model; model documentation can get shortchanged in the process. Poorly documented models are likely to cost money in the long run, by making model revisions more difficult and time-consuming and by increasing the likelihood that model errors will be missed. If necessary, business units may need to establish separate model documentation teams to complete documentation based on interviews with model developers.

Standards for documentation of models and model reviews should include the following:

- A review of the adequacy of business unit documentation by the independent model reviewer should be included, with recommendations for gaps that need to be remedied.
- "Documentation and tracking of activities surrounding model development, implementation, use and validation are needed to provide a record that makes compliance with process transparent" (FRB 2011, Section VI).
- An inventory of all models that require review should be maintained, both by business unit and firmwide. This inventory can serve as a central control point for scheduling model reviews, keeping track of

documentation, providing information on contacts, and scheduling updates of model reviews (FRB 2011, Section VI, “Model Inventory”).

- Documentation is also required for the policies governing model review and the roles and responsibilities of business units and independent reviewers (FRB 2011, Section VI, “Policies and Procedures”).

What is the role for the senior management of the firm in model review? FRB (2011, Section VI) emphasizes senior management responsibility for assuring that a process is in place that meets the standards outlined in this section and in Section 8.2.1. This certainly includes providing adequate funding for these functions. Basel (2009b) has a similar statement in its Principle 1 and, in Principle 2, emphasizes that the review capacity has to be adequate to handle conditions of stress. This is obviously a response to the stressed conditions of the 2008 crisis.

In addition to these more formal requirements, senior management must be prepared to understand the aggregate level of model risk that the firm is exposed to and to set limits on this risk. Having a requirement that business units communicate model risk in an aggregated and comprehensible form to senior managers, as we have at the start of this section, entails a corresponding responsibility of senior managers to make use of this information. There will, of course, be cases of conflicting presentation to senior management—often a more sanguine view of risk from the business unit and a more cautious view from the risk management groups. Senior managers must insist that both sides get a fair hearing, preferably in the same room at the same time, and that arguments be presented in a comprehensible manner. At the end of the process, it is senior management that owns the risk and must reach a decision.

Boards of directors in principle exercise an oversight role over senior management in controlling model risk, as in all other critical aspects of the business. In practice, it is very difficult for directors, whose involvement is only for a small part of each month, to have much impact, particularly since their access to information is often tightly controlled by senior management. The one course of action I would recommend is that directors on the risk committee of the board insist on private meetings with senior risk management personnel. This will at least provide a forum for concerns to be expressed that could give directors enough information to pose questions to senior management.

8.2.3 Model Verification

Most model problems are related to the fit between the product or trading strategy being modeled and the model selected. This issue of model validation

we will address in Section 8.2.6. Here we deal with the simpler question of whether the model selected has actually been properly implemented—does the model actually do what it claims to do? This can be controlled by adequate model documentation and thorough checking by competent reviewers before the model is put into production. Here are a few rules that should be borne in mind to make sure this gets done properly (more detail on these points can be found in FRB [2011, Section V] and Morini [2011, Section 1.5.1, “Model Verification”]):

- Thorough documentation of what the model is trying to achieve, model assumptions, and derivation of formulas must be insisted on. All formula derivations should receive an independent check. Useful advice from Morini (2011, Section 1.5.1) is: “When a model is used for the first time the passages from dynamics to closed-form formulas or the other way around should be verified. This should be the case for any new model developed by a front office quant, also for any new model that simply appears on the Internet—or in a journal. There are errors even in published literature, never be too trusty.”
- Systems implementation of the model should be subject to rigorous standards of documentation, change control procedures, and systems testing.
- The best check on an implementation is to perform an independent implementation and see if the results agree. It is tempting to cut costs by confining checking to having an independent analyst read through the documentation, equations, and code of the model builder and confirm it is correct. But it is much easier to miss an error in reading through someone else’s equations or programming code; it is much more unlikely for two analysts working independently of one another to make the same error.
- Whenever possible, the independent implementation used as a check should employ a different solution methodology than the implementation being tested. For example, if the implementation being tested has used a Monte Carlo simulation, the test should be made solving backwards on a tree, where this is feasible. Using different implementation methodologies reduces further the chances that the two implementations will have the same flaw.
- Models should be tested on degenerate cases that have known solutions. For example, a down-and-out call with a barrier of zero is equivalent to a vanilla call, so setting the barrier to zero in a down-and-out call model should produce the standard Black-Scholes result. Other examples would be: (1) to always check that put-call parity for European-style options is preserved for any model used to price options, and (2) to

always check exotic option models against known analytic solutions for flat volatility surfaces (see the introduction to Section 12.1 and Section 12.3.1).

- Models should be tested for their impact on VaR and stress test calculations as well as on valuation and limit calculations.
- Models should be tested on extreme inputs to see that they handle these cases properly. For example, interest rates much lower and much higher than have recently been experienced should be input to see that the model can produce reasonable results. For more details on this model stress testing, see Morini (2011, Section 3.1).
- Produce graphs of model output plotted against model inputs and explore any instances where they do not make intuitive sense. This is another good check on the model's ability to handle extreme inputs, as per the previous bullet point. The impact of varying several inputs simultaneously should be compared to the sum of the individual impacts to check if the interactions of variable changes produce reasonable results.
- When a new model is replacing an existing model, a thorough benchmarking process should be used to compare results of the two models for an identical set of inputs. Model differences should be checked for reasonableness and unreasonable differences investigated. The same benchmarking standards should be utilized whenever one systems implementation of an existing model is being replaced or supplemented by another systems implementation of that model.
- Model error due to incorrect representation of transactions is just as worrisome as model error due to incorrect equations. This will be addressed in Section 8.2.4.
- A particular point of concern is approximation error introduced by the need for fast response time in a production environment. This will be addressed in Section 8.2.5.

Be careful about the degree of complexity introduced into models. Is there sufficient gain in accuracy to justify the reduction in intuitive understanding that results from added complexity? To illustrate with an example from my own experience:

I had recently taken a new job and found that my most pressing problem was widespread user dissatisfaction with a model upgrade that had recently been introduced. The old model had been easy for traders and risk managers to understand; the new one was supposed to be more accurate, but could be understood only by the model development group. My initial examination showed that, on theoretical grounds, the difference between answers from the two models should be too small to make an actual difference to decision making, so I tried to persuade the model builder to switch

back to the original, simpler model. Finding him adamant on the need for what he viewed as theoretical correctness, I examined the new model more closely and found a major implementation error—a factor of 2 had been dropped in the equation derivation. This is the sort of mechanical error that would certainly have been picked up as soon as a formal model review was performed. But a similar error in a less complex model would have been caught long before, by the people using it on a day-to-day basis.

8.2.4 Model Verification of Deal Representation

Verification of transaction details that serve as input to models can be just as important in avoiding valuation and risk measurement errors as verification of the model itself.

The quote from FRB (2011) in Section 8.2.1 that “Models are of necessity simplified representations of real-world relationships and so can never be perfect” applies just as much to the representation of transactions in models as it does to the models themselves. A single transaction confirmation document often runs to tens of pages and its representation in the model is just a few numbers, so inevitably some simplification and approximation are being utilized.

In some ways, this is a more difficult issue to deal with than verification of the model itself, because of the large number of transactions that are often input to a single model. Reconciling transaction details between confirmations and position entries to models is an important middle-office control function, as emphasized in Sections 3.1.1 and 3.1.2, and will certainly be expected to catch numerical errors in data entry and details such as correct day count convention. But middle-office personnel lack the intimate knowledge of the model that might allow them to identify an important contract detail that is not being captured in the way the transaction is being represented in the model. Some steps that should be taken to control this risk are:

- While model builders and independent model reviewers who do have intimate model knowledge won't have the time to review every transaction confirmation, they should review all of the very largest transactions and a sample of the remainder, to look for both individual errors and patterns of errors. Samples should be selected at random, but with some weighting scheme that makes it more probable that large-impact transactions have more of a chance of being part of the reviewed sample than those of lesser impact. Review should consist of a thorough reading of the confirmation, comparison with how the confirmation has been represented in the model input, and consideration of any possible gaps in the representation.

- Middle-office personnel should be strongly encouraged to immediately raise any question they have about adequate representation with an independent model reviewer.
- In the daily P&L verification process, discussed in Section 8.2.7.1, any transaction that makes a payment significantly out of line with the payment projected by the model should be investigated. This may uncover an outright error in data entry, but may also identify a facet of the contract that has not been adequately represented in the model input.
- Some wording differences between different variants of contracts may be very subtle (for example, see the discussion of legal basis risk on credit default swaps in Section 3.2 and Section 13.1.1.2). The best approach in this type of case may not be to try to capture all these variants in the model input. It may be better to have a separate offline calculation of the risks arising from wording differences and to establish limits and reserves against this risk on the basis of this offline calculation. This can be regarded as a type of liquid proxy, per our discussion in Section 8.4, with the most common contract type serving as the liquid proxy in all standard risk calculations, such as VaR and stress tests, but with the separate offline calculation capturing the nonliquid risk.

8.2.5 Model Verification of Approximations

The quote from FRB (2011) in Section 8.2.1 that “Models are of necessity simplified representations of real-world relationships and so can never be perfect” could be paraphrased as saying that “All models are approximations.” But for model review it is very important to distinguish between two different types of approximations:

1. Approximations in which some source of risk or driver of value has been omitted to reduce model complexity.
2. Approximations in which a computational approximation is being used in order to speed calculations and reduce cost.

Approximations involving the omission of a risk factor pose greater challenges for model review, since it can be very difficult to estimate the potential impact on earnings and risk. This issue dominates our discussion of model validation in Section 8.2.6 and in Sections 8.3, 8.4, and 8.5. Approximations involving computational approximation are much easier to control, since the model review process can create a detailed comparison between the production model and a more thorough model that is run less often or only on a selected sample of transactions. This section focuses on techniques for dealing with computational approximations.

We begin by looking at a set of suggested controls for computational approximation and then illustrate with a detailed example. Suggested controls are:

- Model reviews should explicitly recognize the trade-offs between model accuracy and investment of resources. Models used in production must be sufficiently fast to produce answers within the time frame required for providing quotes to customers and providing risk analysis to the trading desk and senior management or they will prove useless. Their development cost must be reasonably related to the revenue that can be realized on the products they support. The time required for development must be consistent with overall business plans.
- Evaluations of the inaccuracy of a production model need to be made by comparison to a more thorough model. Since model testing can be performed over a period of days or weeks, as compared to the minutes or seconds required of a production model, there is ample room to develop much more thorough models in testing environments. Comparison of results to the production model will show just how much accuracy is being lost.
- Comparisons of the production model to a more thorough model need to be performed not just for current market conditions. Tests should be performed to anticipate the impact on approximation of potential future market conditions.
- Where this test shows significant loss of accuracy, this identifies a good target for improved approximations. Until such improvements can be implemented, remedies can include valuation reserves against inaccuracy along with periodic revaluations with a slower but more accurate model, as well as traders exercising a degree of conservatism in pricing and hedging.
- Improved approximations can be achieved by “throwing more money” at the problem—buying more hardware to increase the number of calculations that can be performed in a given period of time. But more can usually be accomplished by the design of clever approximation algorithms. Indeed, one of the dirty little secrets of industry quants is just how much of the effort of people with PhDs goes into applying advanced mathematics to creating better approximation algorithms, rather than to the creation of new ideas for financial modeling. For example, see Section 13.3.3.
- In some cases, the thorough model can be so computationally intensive that it can be evaluated on only a sample of transactions. This is clearly a less desirable test of accuracy than one that looks at the full portfolio, but when this is necessary the reviewed sample should include all of the

very largest transactions and a random selection of the remainder. The random selection should be chosen with some weighting scheme that makes it more probable that large-impact transactions have more of a chance of being part of the reviewed sample than those of lesser impact.

- Approximations need to be reevaluated periodically, since approximation inaccuracy can be strongly related to portfolio size and composition. This point will be discussed in more detail in Section 8.2.8.1.
- A clear distinction needs to be made between the degree of accuracy needed to specify initial market conditions and the degree of accuracy needed to specify the evolution of market conditions. For example, as we will see in Section 12.5.2, a multifactor model for the evolution of interest rates does not offer much added accuracy over a single-factor model for the valuation of Bermudan swaptions. But a specification of the shape of the initial yield curve that does not utilize a full set of liquid points on the yield curve can have a very significant impact on this valuation.

As an illustrative example, I want to consider a situation I was involved with as a model reviewer. It involved a large portfolio of illiquid interest rate derivatives, all of which required Monte Carlo simulation for valuation calculations and for calculations of risk parameters. The number of simulation paths being run had been very clearly selected by the traders as the number that would allow all of the needed calculations to be performed between the close of business and the opening of trading the next day.

To test for the impact of this choice on accuracy, I first set up a simulation for the whole portfolio that would run continuously for about a week in an offline environment. The results from this very much larger number of runs allowed me to estimate the number of runs needed to determine accurate valuation to within the tolerance required for financial significance (see, for example, Hull 2012, Section 20.6, “Number of Trials”). I was also able to see how much variance from this accurate valuation resulted from the smaller sample being used in production runs.

I next used the larger sample to estimate the impact of the selection of different sets of Monte Carlo simulation paths on the production run. I determined that there was reasonable stability over a sufficiently small time frame; a set of paths that produced accurate values when comparing results from the smaller number of paths in the production run to the larger number of paths used in the offline run would also be fairly accurate over the next few days. But as the time period lengthened from days to weeks, a set of paths that had previously produced accurate values lost accuracy, both as a result of shifts in composition of the portfolio as new transactions were added and older transactions had less time left to expiry, and as a result of changes in market parameters.

I therefore set up the following process. Once a month, a new offline run would take place and would be used to determine the set of paths that was going to be used in the production runs for the next month. These production runs determined valuations and sensitivities reported for P&L and risk management purposes. Each month, the shift from one set of paths used in the production runs to a new set of paths would cause a change in P&L. These changes were randomly distributed, as likely to be increases in P&L as decreases, with a standard deviation that could be estimated from comparisons between valuations of different subsets of paths in the offline run.

I argued that while these changes were randomly distributed, the firm should have a reserve against negative changes so that we were reporting to shareholders only valuations we could be fairly sure were actually achievable in the long run (since the portfolio consisted of illiquid positions, it was not possible to realize a current value just by liquidation—we were committed to holding the portfolio for a longer time period). Therefore, each time new transactions were added, the reserve would be increased to reflect added uncertainty, while as older transactions got closer to expiry, the reserve would be reduced to reflect less uncertainty. Each month, when the change in paths took place, the resulting gain or loss would impact the reserve and would not impact P&L (the reserve would have had to be exhausted to impact P&L, but this never occurred in practice). I was able to get agreement from the firm's finance function and accounting firm to make this process part of the official books and records of the firm, and not just the risk management reports.

Note that this process had a built-in set of controls against changes in the market environment or portfolio composition, since the monthly offline run would automatically pick up any such impacts; the size of required reserve was recomputed each month by recomputation of the standard deviation between valuations of different sets of paths.

The issue of computational approximations is particularly important for credit portfolio models and the closely related collateralized debt obligation (CDO) models. These will be covered in some detail in Sections 13.3.3 and 13.4.2. Another area in which computational approximation plays a major role is in multifactor interest rate models. This is discussed in great depth in Morini (2011, Section 6.2), which includes extensive examination of the accuracy of these approximations.

8.2.6 Model Validation

We must now move beyond the tests of internal model consistency we have focused on in Sections 8.2.3, 8.2.4, and 8.2.5 to look at the fit between the model and the product or trading strategy being modeled, what Morini

(2011) calls *model validation*. It is not surprising that this more challenging task does not lend itself to the easy consensus and process-oriented approach of these last three sections. We will distinguish between three basic approaches to model validation: one focused on interpolation, one focused on the long-term cost of hedging, and one focused on discovering the prevailing market model. These are not completely competitive approaches—there is some overlap among them—but they do have distinctly different emphases. As we will see, the appropriate approach has much to do with the purpose the model will be used for and the liquidity of the product being modeled. We will also look at the issue of how to deal with risks that may not be evident to model reviewers. I would strongly recommend comparing my approach in this section to Morini (2011, Sections 1.5.1 and 1.5.2).

8.2.6.1 Interpolation Approach In the interpolation approach, models are viewed as primarily serving as interpolation tools from observable to unobservable prices. This is closely related to the view discussed in Section 8.1 that downplays the importance of models. Viewing models as interpolation tools provides valuable insight into why certain models have been able to achieve a high degree of acceptance in financial management. It is much easier to agree on an interpolation methodology than it is to agree on a fundamental method for pricing an instrument. The danger is that this view leads to unwarranted complacency, since model builders often regard interpolation as being a mathematically trivial or uninteresting task. The result can be uncritical acceptance of what seems a plausible interpolation method or a view that the choice of interpolation methods is somehow a matter of taste.

A closer examination will show that every choice of interpolation method entails significant financial assumptions. The interpolation of an unobservable price based on a set of observable prices amounts to the theory that the instrument with the unobservable price can be well hedged by the set of instruments with observable prices. As with any theory, this should be subjected to empirical testing and competition with alternative hedging proposals. Even the simplest-sounding interpolation proposal (for example, calculating the two-and-a-half-year rate as a 50–50 average of the two- and three-year rates) should be regarded as a model subject to the same tests as more mathematically complex models. We examine this in more detail in Sections 8.3 and 10.2.1. Models rarely cost firms money because modelers have made an error in complex mathematics; they frequently cost firms money because they embody financial assumptions that are not borne out by future events.

8.2.6.2 Cost of Hedging Approach Basing model validation on an examination of the possible costs of hedging a transaction over the long term is closely

related to the approach advocated in Section 6.1.2 of establishing a liquid proxy for an illiquid instrument and then simulating the difference between the liquid proxy and the actual trade. This viewpoint has been laid out very eloquently in Derman (2001):

It's never clear what profit and loss will result from hedging a derivative security to its expiration. Markets will move in unexpected ways, sometimes intensifying transactions costs and often dismantling what seemed a reasonable hedging strategy. These effects are rarely captured by the conventional models used in front-office valuation systems. . . .

Therefore, for illiquid positions, it is important to estimate the adjustments to conventional marked values that can occur as a result of long-term hedging. One should build Monte Carlo models that simulate both underlying behavior and a trader's hedging strategy to create distributions of the resultant profit or loss of the whole portfolio. These distributions can be used to determine a realistic adjustment to the trading desk's conventional marks that can be withheld until the trade is unwound and their realized profit or loss determined. . . . Monte Carlo analysis provides a good sense of the variation in portfolio value that will be exhibited over the life of the trade due to transactions costs, hedging error and model risk. Ultimately, such analyses should be part of the desk's own front-office valuation system.

Note: There is a rough equivalence between Derman's use of "underlyer" liquid instruments being used to hedge the illiquid instrument and my emphasis on a representative liquid hedge. As Derman says: "Derivative models work best when they use as their constituents underlying securities that are one level simpler and one level more liquid than the derivative itself."

8.2.6.3 Prevailing Market Model Approach Rebonato (2003) emphasizes model validation based on the model that prevails in the marketplace and anticipation of directions in which the prevailing market model might evolve:

"Model risk is the risk of occurrence of a significant difference between the mark-to-model value of a complex and/or illiquid instrument, and the price at which the same instrument is revealed to have traded in the market."

- "From the perspective of the risk manager the first and foremost task in model risk management is the identification of the model ('right' or 'wrong' as it may be) currently used by the market to arrive at traded prices."

- “[M]arket intelligence and contacts with the trader community at other institutions are invaluable.”
- Requires a variety of models to reverse-engineer observed prices.
- Requires information about as many observed prices as possible.
- “No matter how good or convincing a theoretical model might be, few states of affairs should worry a risk manager more than the trader who, using this model, consistently beats all competing banks in a competitive-tender situation.”
- “The next important task of the risk manager is to surmise how today’s accepted pricing methodology might change in the future” (including changes to model, changes to calibration, and changes to numerical implementation). “Being aware of the latest market developments, and of academic papers can be very useful in guessing which direction the market might evolve tomorrow.”
- “To a large extent, the model risk management task can be described as an interpolation and extrapolation exercise that simply cannot be carried out in an informational vacuum . . . without at least some anchor points of solid knowledge about the levels and nature of actual market transactions.”

Hull and Suo (2001) present an approach to model validation closely related to Rebonato’s. They quantify the risk of a model being used by a trading desk by estimating how much of a loss the trading desk would suffer if a different model turned out to be correct.

8.2.6.4 Matching Model Validation to Model Purpose There are two dimensions to matching model validation to model purpose. The first relates to the degree of liquidity of the instrument being modeled. The second relates to differences between models being used for managing the firm’s overall risk, the models of valuation and position measurement described in Section 6.1, and models being used to make trading decisions.

The key to designing a proper model valuation procedure for models being used to manage the firm’s risk is to fit the model review to the degree of liquidity of the instrument for which the model is being used. Essentially, we need to work out the model review implications of the liquidity differences discussed in Section 6.1.1.

Please note that the distinction here is between liquid and illiquid *instruments*, not liquid and illiquid *positions*. If a position in a liquid instrument is so large as to create an illiquid position, this needs to be dealt with by modifying VaR and stress test calculations, as discussed in Section 6.1.4, but does not require a different model or model review than would be needed for a smaller position in the same instrument.

The interpolation approach to model validation is usually very reasonable for liquid instruments. We will look at the details of applying the interpolation approach to liquid instruments in Section 8.3. For illiquid instruments, the interpolation approach has little relevance; the kind of frequent checks of interpolation methodology with the market recommended in Section 8.3 are not possible because of illiquidity. Given that interpolation will be of little use, model verification for illiquid instruments must therefore rely on either the cost of hedging approach or the prevailing market model approach or some combination of the two. We explore this issue in detail in Section 8.4. For models used in making trading decisions, discussed further in Section 8.5, the prevailing market model approach is the most salient, as the next example illustrates.

To better understand the implications of different types of model use, consider the case of the 1998 breakdown of the historical relationship between the pricing of interest rate caps and interest rate swaptions, discussed in detail in Morini (2011, Sections 11.3 and 11.4). For models utilized for managing the firm's overall risk on liquid instruments, this breakdown was probably a nonevent. Since both caps and swaptions had adequate liquidity in external price quotes, most firms would be using some form of interpolation model to value and compute risk statistics for their caps just using market cap prices, and a separate interpolation model to value and compute risk statistics for their swaptions just using market swaption prices. There would have been no interaction between the two models. (Some possible exceptions where liquid prices in one market would be allowed to override liquid prices in another market will be examined in Section 8.3.)

There may have been a different story for models used in managing the firm's overall risk on illiquid products. Some products, such as Bermudan swaptions, knock-out caps, and forward-start interest rate options (see Section 12.5.2 for details), may have been priced using models that incorporated both market cap prices and market swaption prices as inputs. Would the breakdown in the historical relationship have caused problems for these models? Using the approach we discuss in Section 8.4, there should be a very significant degree of conservatism relative to historical relationships built into the reserves kept against model risk and, in any case, the relationship that is important is what will happen over the lives of the deals. Could a temporary period of breakdown in historical relationships be enough to call into question the adequacy of these reserves? We discuss this further in Section 8.4.

But models being used by the trading desk to determine trading strategies involving caps and swaptions would definitely have been impacted. Here's the description by Morini (2011): "No matter whether or not the long term equilibrium was going to come back, the market had gone too

far from it for too long for a bank or a fund with a risk management unit to stand it.” Here it is clearly the case that the prevailing market model approach must govern.

8.2.6.5 Capturing Risks That Are Difficult to Identify I once heard a senior risk manager for an investment bank say, “I don’t stay awake at night worrying about the risks I know, but about the risks I don’t know.” This is a sentiment with which I could readily identify. In performing model validation, the great fear is that there will be some exposure that is not being captured by the model and that you, as the independent model reviewer, don’t even know about. For example, there might be some potential piece of legislation or judicial decision that would have a big impact of the transaction being modeled, but it has never been discussed in the literature you have access to.

This type of potential exposure is probably known to the front-office personnel who specialize in the product. Ideally, they should consider it in their internal model review and share their concerns with the independent model reviewer. But the usual moral hazard concerns come into play, with the incentives discussed in Section 2.1, motivating front-office personnel to be reluctant to share information that might lead to tightened controls.

JPMorgan in the late 1990s instituted an internal system called Risk Identification for Large Exposures (RIFLE) to try to address this issue. It is still in operation, as can be seen from the following quote from JPMorgan’s 2010 annual report:

Individuals who manage risk positions in the Investment Bank are responsible for identifying potential losses that could result from specific, unusual events, such as a potential change in tax legislation, or a particular combination of unusual market events. This information is aggregated centrally for the Investment Bank. Trading businesses are responsible for RIFLEs, thereby permitting the Firm to monitor further earnings vulnerability not adequately covered by standard risk measures. (p. 145)

But even with a mechanism like this in place, the incentive issue remains. Traders who do an honest job of reporting these risks may thereby lower their return on risk measures and attract added scrutiny of position sizes. To attempt to overcome this, a firm needs to make clear distinctions in performance evaluation between losses that occurred due to an event that the traders had made certain received adequate firmwide attention in advance of the trade being approved and losses that occurred due to an event where this type of advance notice was not provided.

8.2.7 Continuous Review

FRB (2011, Section V) calls for “validation activities [that] continue on an ongoing basis after a model goes into use.” Three major components of ongoing validation activities are (1) daily P&L reconciliation for models being used for valuation and risk-reporting purposes, (2) back-testing for statistical forecasting models, and (3) analysis of overrides for cases where model output needs to be altered based on the expert judgment of model users. We consider each in turn.

8.2.7.1 Daily P&L Reconciliation In Sections 3.1.1 and 3.1.2 we stressed the importance of a daily explanation of P&L produced by independent support staff as a control measure against both fraud and nondeliberate incorrect information. Here, we want to stress its equal importance as a tool for identifying model weaknesses.

The basic approach of P&L reconciliation is to take position reports from the close of business (COB) of the prior day and combine them with actual market movements from the previous day’s COB to the current day’s COB to estimate P&L for the day. The idea is that since position reports show sensitivities to changes in market variables (e.g., option “greeks”), multiplication of these sensitivities by actual price changes should produce a reasonable estimate of P&L. Of course, this applies only to those positions that were in place as of the COB of the previous day, so it is first necessary to identify and segregate the P&L due to trades booked during the day (this includes hedges that may have been put on during the day in response to market moves). This segregation of P&L between that due to previous COB positions and that due to trades booked during the day is already valuable as a tool in avoiding inadvertent Ponzi schemes in which profits on newly booked trades cover up hedge slippage on existing trades (compare with the discussion in Section 2.2).

If the initial estimate of P&L is significantly different than actual P&L, what can be the possible causes? Incorrectly recorded positions are certainly a possibility; this is why P&L reconciliation is a valuable tool in uncovering fraud and incorrect information. It’s true that an incorrectly reported position might impact both the COB position report and the daily P&L record, but at some point there will be an actual payment on the position, and at this point, when the payment becomes part of the daily P&L, a discrepancy between projected P&L and actual P&L will show up.

Another possibility is that the COB position has not been reported with sufficient detail. For example, an option position might be reported using just the first-order greeks, such as delta and vega, and not the second-order greeks, such as DdelV and the price-vol matrix (see Section 11.4 for a

detailed discussion of option sensitivity measures and how they can be used in P&L reconciliation). In this case, the P&L reconciliation will identify the need for more detail in the position report, which will enhance management's ability to accurately measure exposure.

A third possibility is a deficiency in the model. For example, it could be that even though the contract details of a position are correctly entered, the model is not using these details correctly in computing P&L or in computing the position. Again, we might worry that the model flaw will impact both P&L and position reporting and so will not be spotted in reconciliation. But when a payment date is reached and actual payment becomes part of the daily P&L, a discrepancy should appear. More likely, the model is missing or mishandling some key factor of risk, and a position that appears hedged is actually suffering some hedge slippage. An example might be a Bermudan swaption model that fails to identify some circumstances in which early exercise becomes more profitable for the counterparty. So daily P&L reconciliation should form an important part of identifying model problems that may have been missed in initial model validation.

8.2.7.2 Back-Testing Any statistical forecasting model needs to have continuous monitoring of actual performance. The best-known example in risk management is back-testing of VaR models, discussed in detail in Section 7.1.2. Since VaR models produce statistical distributions of the size of losses that can be expected to occur at different percentiles (e.g., 1 percent of the time, 2 percent of the time), these projected distributions need to be continuously compared to actual experience. Statistical analysis should be applied to results, and when this analysis indicates a strong probability that the actual distribution differs from the projected distribution, corrections to the model need to be considered. Until corrections can be made, an extra layer of conservatism may be necessary in utilizing limits and reports based on the existing model.

Similar back-testing is called for in any statistical model being used to suggest a trading strategy. In hedge funds and on trading desks, one frequently finds trading strategies employed based on statistical studies of how the strategy would have performed historically (see Fabozzi, Focardi, and Kolm 2010, Chapter 7). Continuous back-testing is needed to update evaluation of the model's performance and identify changes in market environment that require alteration or even abandonment of the model.

8.2.7.3 Analysis of Overrides When hedge funds and trading desks employ statistical models for trading, there will occasionally be the need for a trader to override the trading strategy recommended by the model because of economic insights that cause the trader to doubt the advisability of the model's

recommendation. It is important that all such overrides be recorded and analyzed, with performance of the model strategy not pursued compared to success of the override, to spot possible needs for model modification.

It is less common for models used for valuation and risk reporting to be overridden, but there are examples. One typical case is the representation of binary options in option greeks. When trading desks do not use the liquid proxy representation of binary options by a call spread recommended in Section 12.1.4, it often happens that the trading desk must come to the risk managers and ask for an override on the large delta and gamma positions produced by a binary option nearing expiry at a price close to the strike. Risk managers will be willing to grant these requests, since the trading action that would be required to get back within the limit would be a foolish position to take, as discussed in Section 12.1.4. But any such overrides should be recorded and analyzed. It is just such analysis that has persuaded some firms to move to the use of the liquid proxy representation, since a model that is producing position reports that would be foolish to act on is clearly flawed.

8.2.8 Periodic Review

Once a model has been approved and is in production, ongoing validation is still required. FRB (2011, Section V) reiterates what has been a long-standing regulatory requirement that existing models should be reviewed at least once a year, but also calls for continuous monitoring of model performance. We'll examine periodic review in this section and ongoing monitoring in the next section.

To be productive, the periodic (generally, annual) review of existing models must be carefully designed. Merely replicating previous tests is likely to be both unlikely to produce new insights and wasteful of resources. Reviews need to be focused on changes to the environment in which the model is being used that should trigger new testing and possibly new conclusions. We'll focus on four types of environment changes that should be investigated: (1) changes in the population of transactions the model is being applied to, (2) changes in the market environment, (3) changes in the academic literature or in market practices, and (4) changes in technology. In addition, the periodic review should examine any patterns that have been revealed by the ongoing monitoring we describe in the next section.

8.2.8.1 Changes in the Population of Transactions Consider the Kidder Peabody disaster, discussed in Section 4.1.2, as an illustration. Whatever your opinion about whether Joe Jett was deliberately gaming the system, there is no doubt that the firm was ill-served by having a model that computed the value of

forward transactions without proper discounting. But Kidder Peabody was hardly alone in the industry in using a model that omitted discounting of forwards. This is not due to a widespread ignorance of this fundamental principle of finance. What does often happen—and this is a pattern I have seen over and over again—is that a sensible decision is made at the time a model is built but is not subjected to adequate review as circumstances change.

So a model might be set up for valuing forwards that at the time of implementation is being used to evaluate trades that are of moderate size and no more than a few days forward. The added accuracy that comes from correct forward discounting might be quite small and thus easily justify a decision not to devote the added programming time and computational resources to include this factor. The situation changes as larger transactions with longer forward periods are added. As the situation changes through time, there comes a point at which the proper decision would be to change to a more accurate model. But the decision to invest the resources needed to improve accuracy can be a difficult one, involving considerable expense, diversion of resources from important new ventures, and perhaps a limitation on trading volume until the change is made. The environment may be changing gradually, so that no single point in time stands out as the time at which to switch.

This is the kind of situation in which a periodic review of the impact of changes in the population of transactions being valued by a model can be of tremendous value. Note that the change in population of transactions could be due to changes in number of transactions (an approximation that had little impact when the model was being used to value just a few deals causes more concern when the model is being used to value many deals); size of transactions (maybe the model is being used to value just a few deals, but the average size of deals has grown to the point that approximations have become worrisome); terms of transactions (for example, the large increase in the length of the forward period in the case discussed previously, or an increase in time to maturity, or the more frequent use of features that are difficult to evaluate); or a combination of all three.

Morini (2011, Section 1.4.1) explains the reasoning behind this policy very well:

A bank cannot expend big resources for a small exposure; and additionally banks and traders learn by trial and error, a new model needs to be tested for a while to really know its risks. When the exposure starts growing, a previous model validation must not automatically be considered valid: a surplus of effort can be spent on the model used, an effort that was not economically meaningful in the past but is crucial in the face of the increased exposure.

When faced with a large change in deal population, an independent model review group must think very carefully about what is behind the change. Is it just a new market taking off to meet a customer need, or is it a structuring group looking to arbitrage a deficiency in the way a transaction is being valued or a risk is being measured? If traders and structurers are hired because of their skills in uncovering complex arbitrage opportunities in markets, one shouldn't be too shocked if they sometimes use the same skill set to try to find arbitrage opportunities in regulations, whether external (government) regulations or internal (risk management) regulations. When independent reviewers see signs that such an opportunity is possibly being exploited, they must expend extra effort on trying to uncover the motivation and possible consequences.

If a reviewer does spot a loophole being exploited, there should be no hesitancy in quickly improving the valuation procedure or risk measure. There will inevitably be cries of foul play from structurers who can no longer take advantage of the old system and complaints that "The rules of the game have been changed without warning." Those complaining need to be reminded that risk management is not a game but a serious endeavor to protect shareholders, depositors, and taxpayers. Keeping a fixed set of rules and allowing structurers to experiment and see where the weaknesses are is a recipe for disaster, as the rating agencies amply demonstrated by publishing fixed models for evaluating the risk of CDO tranches and letting bank structurers play with the models until they had designed trades that optimized the degree to which the models underreported deal risk (see Section 5.2.3 for further discussion of this example).

8.2.8.2 Changes in the Market Environment An example of a change in market environment would be a risk factor that previously could not be priced based on market observation but now has liquid prices available. This could change previous conclusions about which model inputs need to be derived from market prices. In the other direction, deterioration in the liquidity of a pricing source might prompt the need for new reserves or limits.

Another example would be changes in levels of prevailing market prices that might prompt reruns of sensitivity analyses and model stress tests. FRB (2011, Section V) states, "Sensitivity analysis and other checks for robustness and stability should likewise be repeated periodically. . . . If models only work well for certain ranges of input values, market conditions, or other factors, they should be monitored to identify situations where these constraints are approached or exceeded."

Another instance of change in market environment would be rapid growth in the size of a market. This should prompt reexamination of the relevance of historical data, since rapid growth may be the result of major

changes in the nature of the market. A recent example was the explosive growth in U.S. subprime mortgages. The use of historical data on default rates for subprime mortgages should have then been treated with extreme caution. As it soon became clear, underwriting standards for approving these mortgages had become drastically more lax than in previous eras, which contributed to steeply rising default rates (see Section 5.2.1). A prior example was the precipitous growth in non-investment-grade bonds in the late 1970s and early 1980s. This growth was largely due to the efforts of Michael Milken at Drexel Burnham Lambert, and a major contributor was studies by Milken and others showing the very favorable historical returns of these bonds after adjusting for default losses. But as soon as there was a large increase in the issuance of these bonds, it should have been suspected (as turned out to be the case) that the growth was largely being fueled by types of transactions that had rarely been done previously and for which the historical data was of dubious relevance. Bruck (1988) is a good account of the Milken story; see particularly page 28–29 on historical return studies, and pages 266–270 on skepticism about their continued relevance.

8.2.8.3 Changes in the Academic Literature or in Market Practices Periodic reviews offer a good opportunity to consider any new approaches to modeling a particular type of transaction that have appeared in the academic literature, have been discussed at conferences, or have begun to be used by other market participants. This is where the emphasis in Rebonato (2003) on “market intelligence and contacts with the trader community at other institutions” and the reverse engineering of observed prices from other firms can be of particular value.

8.2.8.4 Changes in Technology Increased computational capacity may change the conditions on which previous decisions about approximation techniques have been made. Increased computational capacity could be due to newly purchased or upgraded hardware or to advances in computational theory. New conditions should lead to a reassessment of prior decisions—replacing existing approximation techniques either with more accurate ones or with full computations.

8.3 LIQUID INSTRUMENTS

Models for liquid instruments are robust and easy to test, since they can constantly be checked against actual liquid market quotes. This is why they lend themselves so readily to the interpolation approach to model validation outlined in Section 8.2.6.1. Risk reports only need to look at exposures,

such as delta and vega, measured against current market prices. If changes in price levels lead to new exposure levels that concern senior managers, the liquidity of the instrument will allow for reduction in positions at the time the exposure exceeds desired levels.

We illustrate with an example of a model review for a very liquid instrument. Consider a portfolio of U.S. dollar interest rate instruments (e.g., interest rate futures, forward rate agreements, interest rate swaps, government bonds) with no option component (see Chapter 10 for a detailed discussion of how the risk on such a portfolio is managed). There will be liquid market quotes available throughout the day for trades on a large subsection of these positions. But many instruments will need some form of modeling for valuation, since even a very liquid (“on-the-run”) instrument at the time of original transaction may soon become less liquid (“off-the-run”) through the passage of time (e.g., a five-year swap, for which direct market quotes are readily available, soon becomes a four-year 11-month swap, whose liquidation price needs to be inferred from market quotes for on-the-run instruments). The models used for this off-the-run valuation will also be needed for computing the change of the portfolio’s value in VaR and stress test simulations.

The models needed for these computations are quite standard throughout the financial industry by now, but there are still choices in interpolation methodology that need to be made that constitute forecasts of relative movements between instruments (Section 10.2.1 provides details). These modeling choices are best made by the front-office personnel who have the product expertise and superior data access. In addition, it is the front office to whom the profits from correct forecasting decisions (and losses from poor forecasting decisions) properly belong.

Model validation by outside reviewers only requires periodic checking of model valuations against actual market prices. Close agreement shows model adequacy; significant differences indicate the need to establish limits and/or valuation reserves and may serve as clues for model revision. The most robust price checks come when there is an actual transaction in an off-the-run instrument, but price checks can also be performed by polling brokers, dealers, and other independent sources of pricing information (issues involved in obtaining such quotes are addressed in Section 6.1.3). While actual conduct of the price checks may be performed by support staff, model reviewers and other senior control personnel should be involved in the design of the price check procedures, with regard to frequency and standards for confirmation of model adequacy.

The type of price check just discussed should be complemented by the daily P&L explanation exercises discussed in Section 8.2.7.1. As observed there, the P&L explanation process often identifies model deficiencies when

there are unexplained P&L changes, particularly around transaction dates and dates for scheduled payments and resets. But even a thorough daily P&L explanation process should not be regarded as a full substitute for price checking; it may be that a model performs very well in handling on-the-run transactions from inception through maturity and is rarely tested on off-the-run transactions because the trading desk almost always transacts on on-the-run dates. But what happens if the desk goes through a stop-loss limit or if the firm's appetite for risk decreases? A reduction in positions may need to take place by reversing previously booked transactions that are now off-the-run. Risk managers will want to have prepared for this eventuality by testing model pricing of off-the-run positions.

If the disagreement between an observed market price and a model value represents a clear difference between where a risk can be sold at the current time and a theory as to the value of the asset over a longer period of time, then no matter how sound the reasoning behind the theory, I would recommend holding to the mark-to-market principle. If a firm deviates from this principle and values based on longer-term values that it believes can be realized, rather than on prices at which risks can currently be exited, it is turning short-term risks into much harder-to-evaluate long-term risks. Morini (2011, Section 1.2.1) supports this view colorfully: "on intuitive grounds, anyone who claims that arbitrage opportunities are abundant in the market should always be asked if he is fabulously rich. If he is not, it is not clear why he knows of so many free lunches and yet, rather than exploiting them, goes around passing preposterous judgments about market functioning."

However, sometimes the difference between an observed market price and a model value represents two different ways in which a risk can be sold at the current time. Although this would seem to violate several important axioms of finance theory—the *no-arbitrage principle* and the *law of one price*—these are just models and cannot expect any absolute deference in the face of empirical exceptions. However, there needs to be careful evaluation of what lies behind an observed difference between a market price and a model price before an intelligent decision can be made as to which is the best of two different ways to represent the risk.

Let's focus on a concrete illustration. You have observable market prices for a European call option, a European put option, and a forward to the same expiry date, with the same underlying and, in the case of the put and the call, the same strike price. The combined prices, however, do not agree with put-call parity. This would imply, for example, that a position in the put that you have sold can be offset in two different ways—you could buy a put, or you could synthetically create a put by buying a call and entering into a forward. It also implies that the call-forward combination will offset the position at a cheaper price than the direct purchase of the put.

What should a risk manager recommend in such circumstances? Since the main argument behind a no-arbitrage principle such as a put-call parity is that the lack of parity will be quickly eliminated by profit-seekers taking advantage of a riskless opportunity to make money, any persistence of parity violation is suggestive of some liquidity difficulties preventing the opportunity from being exploited. We'll consider some possibilities:

- This is an arbitrage of which very few market participants can take advantage, but your firm is one that can. This could be because the market for the put is in some way restricted to only a few firms. It could be an arbitrage that is difficult to identify computationally and your firm has a computational advantage. It could be a diversified basket of assets that is difficult to accumulate and your firm has an advantage in its market access (see the discussion in Section 12.4.1.1). In such cases, it is right to base valuation on the model-derived price (in this instance, the call-forward combination), since this represents a liquid external price at which risk can actually be extinguished in the short term.
- One of the prices is less liquid than the others. For example, the amount of trading for that strike and date could be much more active in calls than in puts. This would be a strong indication of the desirability of using a model (put-call parity) to supply a price based on more liquid quotations rather than utilizing a less liquid price. The same reasoning would apply if the call and put markets are significantly more active than the forward market, in which case I would recommend replacing an illiquid forward price with a put-call parity–derived price based on liquid put and call prices.
- A timing difference exists in price quotations. Perhaps the options market posts closing prices at an earlier time of day than the forwards market. It is certainly legitimate to use a model to update both call and put quotes to adjust for changes in the forward since the time the options market closed.
- Some contract features make the model not completely applicable. Sometimes, on closer examination, contract provisions call into question the applicability of a model. In this case, it might be an allowance for early option exercise in certain circumstances, whereas put-call parity applies only to options without early exercise provisions.

This last type of case has led to a considerable number of disputes between risk managers and trading desks. One example that has arisen at several firms is traders' desire to unlock stock option values contained in convertible bonds. Option models applied to convertible bond prices frequently indicate implied volatilities that are quite low compared with the

implied volatilities that can be derived from plain equity options on the same stock, leading traders to conclude that buying the convertible is a good value trade. Trading desks hungry to book immediate profits have pressed for overriding reasonably liquid convertible price quotes with a model-driven quote based on the implied volatility from the equity options market. But a convertible bond contains the option to exchange a bond obligation for a stock obligation rather than to exchange cash for a stock obligation, so it cannot be completely reduced to the value of an equity option (see the discussion in Section 12.4.4). When turned down on their first attempt, some trading desks have shown good enterprise in marketing total return swaps on the bond portion of the convertible in an attempt to isolate the equity option portion. So long as the swap has been properly engineered to cover all contingencies, such as canceling the swap without penalty in the event that the bond is converted for equity, a complete decomposition can be achieved and it is legitimate to value the resulting position as an equity option. Risk managers have, however, been very careful to check that no uncovered contingencies are present before allowing this valuation change.

8.4 ILLIQUID INSTRUMENTS

8.4.1 Choice of Model Validation Approach

Model use for illiquid instruments is much more critical than it is for liquid instruments and, unfortunately, model validation is also much more challenging. There may be a complete absence of actual market prices at which positions can be unwound, so modeling assumptions and inputs for unverifiable model parameters now become a key driver of model valuation.

Both Derman (2001) and Rebonato (2003) have strong statements as to the difficulty these risks can entail.

- Derman: “Because of their illiquidity, many of these positions [in long-term or exotic over-the-counter derivative securities that have been designed to satisfy the risk preferences of their customers] will be held for years. Despite their long-term nature, their daily values affect the short-term profit and loss of the banks that trade them.”
- Rebonato: “What differentiates trading in opaque instruments from other trading activities is the possibility that the bank might accumulate a large volume of aggressively-marked opaque instruments. When, eventually, the true market prices are discovered, the book-value re-adjustment is sudden, and can be very large. Stop-loss limits are ineffective to deal with this situation, since the gates can only be shut once the horse has well and truly bolted.”

Let's consider which of the model validation methodologies of Section 8.2.6—cost of hedging or prevailing market model—is more appropriate for illiquid instruments. My own view is that the cost of hedging approach is the more relevant for independent model reviewers for two reasons:

1. Even if a given model prevails in the market place, so long as the trading desk can't actually extinguish positions at the prices implied by the model, owing to illiquidity, it is actually the hedging costs that will determine the firm's P&L on the product. The model might continue to prevail in the marketplace for many years, and all the while the firm loses money on its hedging strategy. An advocate of the prevailing market model approach might respond that if, in fact, the model leads to hedging losses, then firms will eventually replace the model, so this is just a case of anticipating the direction in which the prevailing market model may evolve, in line with Rebonato's proposed criteria. But then I would still want to utilize the cost of hedging approach as a key tool in anticipating prevailing market model evolution.
2. I am wary of the ability of risk managers to anticipate prevailing market model evolution using any other tool besides cost of hedging simulation. Some of the tools Rebonato recommends—market intelligence and contacts with the trader community at other institutions—seem much easier for traders to utilize than independent reviewers.

If independent reviewers do rely on the cost of hedging approach, it would still be valuable for the front-office reviewers to utilize the prevailing market model approach as a supplement. This is particularly true when mark-to-market policies require marking to the prevailing market model, so that even if an instrument is being priced and hedged in a way that will virtually guarantee long-term profit, accounting losses may need to be booked in the shorter term. Rebonato (2003) makes this clear in Section 2.1, saying that "model risk arises . . . because of a discrepancy between the model value and the value that must be recorded for accounting purposes." This would not be the case for the mark-to-market policy I advocate in Section 8.4.4.

In a thorough review of cost of hedging, Monte Carlo simulation allows systematic consideration of many possible future paths of relevant liquid market prices and other economic variables. The soundness of the model can be judged only over longer time periods, when longer-term unobservable prices transform into shorter-term observable prices, when there is enough time to observe the impact of required rehedges, or when trades reach maturity and require contractual payments. Over a short time period, almost any model chosen will appear to perform well by a type of circular

reasoning: The instrument with unobservable prices will be valued using the model and the observable price inputs. Therefore, the movement of the unobservable prices relative to the observable prices will seem stable since the same model is being used for valuation throughout the time period (see the example of this discussed in Section 10.2.1).

Proper design of a model review of an illiquid instrument utilizing Monte Carlo simulation has two parts: (1) choice of the liquid proxy, which will be analyzed in Section 8.4.2, and (2) design of the simulation, discussed in Section 8.4.3. I will then look at issues this approach raises for mark-to-market policies in 8.4.4 and for risk measurement in 8.4.5.

Another application of this approach to creation of liquid proxies and simulations of hedge slippage involves investments in hedge funds for which you lack data on current holdings of the hedge funds. Trying to draw conclusions from the historical pattern of the returns for the hedge fund and historical correlation of these returns with other positions is dubious, given the possibility that current holdings of the hedge fund may not resemble historical holdings. The case for treating these investments as illiquid rests both on limitations on hedge fund withdrawals and on lack of information on true exposure. A liquid proxy can be built by making reasonable inferences about the current style of the hedge fund, based on whatever public and private information you have available from the fund. Ineichen (2003) is an excellent starting point for explanations of the differing hedge fund styles (Chapter 5) and detailed examination of each style in relation to indexes of liquid investments (Chapters 6 through 8). A statistical approach to creating liquid proxies for each hedge fund style can be gleaned from Hasanhodzic and Lo (2007). The liquid proxy can be used for representing hedge fund investments in VaR and stress test calculations. Statistical analysis can then be performed on deviations between the liquid proxy and historical returns on hedge funds.

8.4.2 Choice of Liquid Proxy

A choice of liquid proxy is equivalent to a choice of what liquid market prices are utilized in modeling the illiquid instrument. Every model choice implies a liquid proxy, and every liquid proxy choice implies a model.

In evaluating whether a liquid proxy choice is correct, it is necessary to ask whether the implied model makes adequate use of available liquid market prices. This is closely related to one of the key questions in Derman (2001): “Has the model been appropriately calibrated to the observed behavior, parameters and prices of the simpler, liquid constituents that comprise the derivative?” This point can be most clearly made using a concrete example, which is discussed more fully in Section 12.4.2.

Consider an option written on a basket consisting of two stocks. You could choose two different ways to model this: (1) have a complete model of the price evolution of each of the two stocks individually and assume a correlation between them, or (2) directly model the price evolution of the basket. We'll call these the correlation model and the direct model, respectively. Assume that there are liquid market prices for options on the individual stocks but no liquid market prices for options on the basket or on the correlation, which is a fairly standard situation.

It can be argued that either model is a reasonable choice. In either case you will need input for a variable that cannot be observed in the market. In both cases, you have included all the sources of risk in your model.

But, as will be shown in our more detailed discussion in Section 12.4.2, where correlation is not expected to be too negative, the first model offers definite advantages in terms of making better use of liquid market prices. Options on the individual stocks will serve as effective partial hedges for the basket option, so utilizing the first model, which can be calibrated to current market quotes for these options, offers the following advantages over the second model:

- The correlation model implies a liquid proxy that represents the basket trade in the exposure reports for options positions on the two individual stocks. This encourages the use of liquid hedges.
- The correlation model will require valuation changes in the basket when there are changes in the implied volatility of the two individual stock options. The direct model does not require such valuation changes and so can result in stale valuations not fully reflecting the cost of unwinding some of the risk in the trade.
- The correlation model exhibits significantly lower statistical uncertainty of results compared with the direct model. This should permit lower required reserve levels and larger limits than could be allowed if the direct model was used.

Note that these advantages of the correlation model over the direct model are based on empirical, not theoretical, findings. As can be seen in the fuller discussion, if correlation levels are expected to be very negative or if the product were structured differently (for example, an option on the difference between the stock prices rather than on the basket), the advantages of the first model over the second would diminish to the point of indifference between the models.

In some cases, the liquid proxy used could consist of an instrument that is not itself liquid, but for which modeling in terms of a liquid proxy and simulation of the remaining risk have already been incorporated into

the firm's risk management system. This is reminiscent of the quote from Derman in Section 8.2.6.2: "Derivative models work best when they use as their constituents underlying securities that are one level simpler and one level more liquid than the derivative itself." For example, in Section 12.3.3, in examining a liquid proxy for barrier options based on Peter Carr's approach, I advocate the use of illiquid binary options as part of the liquid proxy, noting that "techniques we have already developed for managing pin risk on binaries" in Section 12.1.4 "can now easily be brought into play."

8.4.3 Design of Monte Carlo Simulation

Modeling the differences between the actual trade and its liquid proxy must go all the way to final payout or to when the trade becomes liquid. Modeling must reflect the possibility that the model used for pricing and trading the product may be wrong. Modeling should be by Monte Carlo simulation to reflect a full range of possible outcomes and to generate a statistical distribution that can be used in assessing issues such as capital adequacy. Let us take these points in more detail:

- Don't assume that an illiquid instrument will become liquid—it may happen but it shouldn't be assumed. Another way of saying this: It is important that statistical analysis of the distribution of parameters be based on actual market observations and not on derived values, since the derived values often themselves contain modeling assumptions subject to error. For example, if a given market is currently liquid only out to seven years, use only quotations out to seven years in your hedging simulations; 10-year quotations derived by extrapolation should not be used. This is analogous to the point made earlier in Section 8.4.1 about avoiding circular reasoning in model validation.
- Statistical assumptions used in determining distributions should not be constrained by any assumptions made within the valuation model. For example, the valuation model may assume a normal distribution of a factor because it is computationally simple and the increase in accuracy from using a different distribution can be shown not to be worth the added investment. This would not in any way justify assuming that the corresponding input variable is normally distributed in a model-testing simulation, since the computational trade-offs motivating the model-building decision do not apply to the model-testing calculation.
- Independent reviewers must be careful not to rely on statistical analysis prepared by traders. It is notoriously easy to employ data-mining techniques to find statistical proofs of nearly any relationship by selecting the right historical data set. Statistical controls, such as careful

discipline about segregating historical data into sample periods to fit parameters and out-of-sample periods to test results, are useful, but can still be defeated by sufficiently industrious data mining. It is better to have truly independent analysis, even at the risk of inaccuracy (on the side of conservatism) from lack of insider information.

- Use of Monte Carlo simulation allows for generation of a full statistical distribution of results, which can be very useful for issues such as determining capital adequacy on illiquid positions. This is a necessity if the capital adequacy proposal of Section 8.4.4 is to be followed.
- It must be emphasized that any statistical distribution involving tail risks requires subjective probability judgments (as discussed in Section 1.3). Still, the basic approach of insisting that simulation be of hedge trades involving liquid instruments, and that simulation go all the way to the point at which the original position becomes liquid, means that there will be a lot of historical liquid pricing data that can be utilized in forming these probability judgments. In essence, while illiquid instruments cannot be fully evaluated based on current liquid prices, they can be evaluated based on the future evolution of liquid prices. For illustrative examples, see Sections 10.2.2, 12.1.4, 12.3.3, 12.4.2, 12.5.2, and 13.4.3.
- Use of Monte Carlo simulation avoids the overstatement of risk that can result from more formulaic risk calculations. For example, if the desire is to reserve to a 90th percentile degree of certainty, using 90th percentile values of the distribution of two or more input parameters will likely result in a far greater than 90th percentile degree of certainty in the reserve. In a Monte Carlo simulation, many reruns of the valuation model are made based on sample points chosen randomly from the assumed distribution of each nonliquid variable, and with explicitly assumed correlations between variables. The 90th percentile of model outputs can then be estimated.

Derman recommends a full simulation that includes both underlying behavior and trader hedging strategy. Section 11.3 contains an example that comes close to Derman's proposed full simulation: a Monte Carlo simulation of dynamic hedging of a less liquid option (less liquid because of a nonstandard strike). Sampling over the simulation paths yields a statistical distribution of the differences between the payout on the option and the costs of the hedge.

Derman's recommendation of a full simulation including trader hedging strategy represents an ideal that may sometimes be difficult to achieve in practice. In the simulation in Section 11.3, a full simulation is possible because the assumed trader strategy is very simple, just varying the delta

hedge of the underlying forward. Trader strategies that involve changes in options positions are much more difficult to simulate, because a full specification of the volatility surface is required at each node of the simulation. An illustration of this point can be found in the discussion of barrier options in Sections 12.3.2 and 12.3.3.

When a full simulation is not practical, then I still believe that a simulation should be done, but computation can be simplified by restricting hedging strategies. Easier implementation comes at a cost of greater conservatism, since the full range of possible trader hedging strategies will not be captured. The simulations that I refer to in the next-to-last bullet point of the preceding list can serve as helpful paradigms.

8.4.4 Implications for Marking to Market

Choosing a good liquid proxy, following the guidelines of Section 8.4.2, should assure that illiquid positions are marked to market to reflect changes in liquid market prices. To illustrate with the example used in that section, when there is a change in the implied volatility of one of the two stocks in the basket, it will be immediately reflected in the marking to market (MTM) of the liquid proxy and hence in the MTM of the basket option, which consists of the sum of the MTM of the liquid proxy and the reserve for the difference between the basket option and the liquid proxy.

But should there also be an adjustment to the MTM of the illiquid position based on new information about parameters that cannot be sourced from a liquid market? Continuing with the same example, the question would be whether to change the MTM of the basket option based on new information about correlation between the two stocks. My answer would be that this should be done only very rarely. We have classified the correlation parameter as one that has no liquid market pricing source, so where would frequent updates be coming from? There are two possible sources:

1. Analysis of historical price data has led to a change in estimates of the correlation to be used. But this will only occur infrequently—if the correlation has been estimated from a long data history, then it will usually take months of new data before conclusions will change significantly.
2. There is evidence that the price being charged customers for correlation has changed. But since this is not a liquid market at which risk can be exited, the argument for making immediate use of such new data is not nearly as strong as it is for liquid instruments.

In both cases, new information on correlation might ultimately impact the reserve for the difference between the basket option and the liquid

proxy, and thereby impact the total MTM of the basket option. But in both cases, you would expect to see this impact take place infrequently. In fact, I would argue for designing reserve calculations in a way that would make such changes extremely infrequent. For example, in this case, calculate the reserve based on an extremely unlikely *level* of correlation as opposed to an extremely unlikely *change* in correlation from the long-term average. That makes it less likely that new information about a shift in the long-term average will require a change in reserve level.

The reason I want to make reserve changes infrequently is that I don't think reserve level changes provide good incentives to traders and marketers. Changes in MTM of liquid instruments provide good incentives for exiting positions—either because stop-loss limits are being breached or because accumulating losses cause traders to rethink the desirability of positions (this includes changes in MTM of liquid proxies, which can trigger hedging actions in liquid markets). But changes in reserve levels won't provide much incentive to exit existing positions, since the illiquidity of the instrument makes such exits very difficult. It is true that raising reserve levels may send a signal to marketers to be more reluctant to book new trades, which might argue for raising reserve levels on new trades but not on existing ones.

Even if you are convinced this policy makes good risk management sense, you still might be reluctant to have it guide the MTM reporting of the firm. Financial controllers, independent accounting firms, and regulators all tend to be suspicious of policies that involve high reserve levels that shield reported earnings from fluctuation; it looks like an attempt to smooth reported earnings. Let me make the following points concerning this:

- I believe that the policies I am advocating here represent an accurate picture of what is known about earnings. The true earnings on illiquid positions are often not known until the trade matures. A highly conservative reserve level is therefore justified, and it is unreasonable to expect much new information to arrive from outside sources; the real information will come over time as the trade matures. There are exceptions—new information that would change your outlook for the whole distribution of an illiquid input. An example would be long-term default rates on home mortgages in 2007 when new information on deteriorating underwriting standards would have impacted reserve levels that were previously viewed as prudently conservative.
- Reserving policies can be designed to assure independence and shielding from manipulation that attempts to use reserve levels to smooth earnings. See Section 6.1.4.
- These policies could help to deal with some of the concerns being expressed about the harmful impact MTM policies are having on bank

management (see the reducing procyclicality discussion in Section 5.5.8.1). MTM losses for liquid instruments encourage banks to shed volatile assets, but MTM losses for illiquid instruments, since the banks can't shed the assets, result in a need to raise new capital, often in economic environments that are the most challenging for raising capital, leading to paralysis of the banking system. (This is discussed in greater detail, in the context of the 2007–2008 crisis, in Section 5.3.2.) My proposal causes large reserves to be taken up front, when the environment is still favorable for raising capital, and then releases the reserves, and hence frees capital for new investments, as the existing investments unwind.

- I would definitely advocate strong controls on the use of this accounting policy, only permitting it for positions the firm designates at the time of creation as illiquid.
- I have experience with a policy close to the one described working in practice over a several-year period, from 1996 to 2003, at Chase and JPMorgan Chase, with the full knowledge of risk managers, financial controllers, independent accountants, and regulators. Reserve levels established were sufficiently conservative that they almost always proved adequate at an individual product level, and always proved more than adequate at an aggregated firm level.
- In the current environment, following the debacle of 2007–2008, it may no longer be possible to get independent accountants and regulators to go along with a policy like this; it requires more trust of the motivations of firm risk management than may now be achievable. In that case, I think risk managers should argue for keeping an internal set of accounts that most accurately reflects the economics of a business, even where this diverges from external reporting.

8.4.5 Implications for Risk Reporting

In Section 8.3 we noted that for liquid instruments risk reports only need to look at exposures measured against current market prices, since future exposures due to changes in price levels can always be reduced utilizing the liquidity of the instrument. This approach will not work for illiquid instruments. To take an example, discussed at greater length in Section 12.1.4, a binary option might currently show very little gamma exposure but might have an unacceptably large gamma in the future if prices are close to the strike level when little time is left until option expiry. You can't just wait to see if this will happen, since if it does you can't count on being able to extinguish the risk by selling the digital option. You need to deal with this contingency at the time you are considering creating the option.

One way of handling this is to run risk reports at the time you are considering creating the position that look at a range of future possible price levels for future dates. Acceptability of possible future risk exposures are evaluated as part of the decision-making process for taking on the position.

Another way of handling this is to make sure that the liquid proxy and simulation methodology of Sections 8.4.2 and 8.4.3 adequately control possible future exposures. Continuing with the binary option example, you would make sure that the call spread liquid proxy chosen can only give rise to reasonable future gammas, by making sure that there is a sufficiently wide gap between the strikes of the call spread. As you will see in Section 12.1.4, widening the gap between the strikes will lead to more uncertainty in the simulation and hence higher reserve levels and tighter limits, but this should be viewed as a necessity for controlling future gamma exposure.

8.5 TRADING MODELS

When a model is being used as part of a trading desk's decision-making process, it clearly requires internal model review by the model creators and users. For the model validation part of this process, it is particularly important to review how the model relates to the prevailing model being used in the market and to try to anticipate evolution of the prevailing market model, as argued in Section 8.2.6.3. The question I want to examine here is whether such models also require an external review by an independent group if the model is to be used only for trading decisions and not for the firm's official valuations and measurement of risk.

Major trading losses are frequently ascribed to the firm having the wrong model. What is often unclear in these claims is whether "having the wrong model" just means making incorrect forecasts about the future direction of market prices or if it means misleading the firm's traders and managers about the nature of positions being taken. A good illustration is the discussion in Section 4.2.1 of whether the reliance by Long-Term Capital Management (LTCM) on models should be viewed as a primary cause of the collapse of the fund.

Any firm engaged in making markets or investing funds must take positions whose profit or loss will depend on the correctness of forecasts of moves in market prices. Different strategies will be tied to different price relationships. Some depend on overall market direction, whereas others depend on the relative price of related assets; some depend on getting a long-term trend right, whereas others depend on correctly anticipating short-term moves. However, traders will always need to make judgments about an uncertain future, and firm managers in turn will always need to make

judgments about how much of a risk of loss they will allow a trader to take in exchange for a possible gain. When making this assessment, management will be guided by evidence of prior accuracy of the trader's forecasts.

Nothing in the last paragraph will be altered by whether a trader uses a model as a computational aid in forecasting, unless perhaps management is lulled into a false sense of security by believing that the use of a model lessens the chance of errors in trading judgment. However, if a model results, either purposely or inadvertently, in misleading traders and managers about the relationship between positions being taken and the size of possible losses, then the accusation that model error resulted in the loss is far more plausible.

For example, a spot foreign exchange (FX) trader could be using a very complex model when deciding which positions to take. This could even extend as far as program trading, in which a computer actually issues the buy and sell instructions based on model output. However, spot FX positions can easily be valued based on external quotes, and position size is extremely easy to understand without the aid of models (see the discussion in Sections 9.1 and 9.2). So it is easy for management to see what the profit and loss (P&L) is every day and to cut the risk if P&L performance has been poor. Thus, the modeling does not have any of the dangers of hidden risk, such as Ponzi schemes (see Section 2.2). No FX trader would dream of asking to report more profits this year because he can "prove" that his model (or trading style) will work better next year than this year.

When I was in the position of managing the independent model reviews for a firm, I argued strongly against my group reviewing the validity of models that were being used only for trading decisions. Partly, this was an attempt to conserve resources for what I viewed as the more important task of validating models used for valuation and risk measurement. But even more, I was concerned that traders would use model validation by my independent reviewers as a stamp of approval that would discourage critical review of trading strategies by senior management. I argued that since we weren't being asked to review trading strategies that didn't involve models, the use of a model did not transform us into experts on trading strategy. In particular, how could we obtain the insider knowledge that could allow us to anticipate evolution of the prevailing market model? This is the position I advocated in the first edition of this book, but on reflection, I would reconsider my previous stance.

When position limits are being set and when actions following a stop-loss limit overage are being reviewed, there is no question that traders will utilize results from their trading models to make their case to senior managers. Since senior managers will not have the time or, usually, the skill set to form their own judgment of these models, it is only by having independent

reviewers look at the models that an effective challenge to trader claims can be prepared. Independent reviewers must make clear the limited scope of their review, but can certainly raise issues concerning possible cherry-picking of historical data or reasons why shifts in the economic environment might bring conclusions based on historical data into question. These challenges may prove of value to traders as well as senior managers. And certainly, independent review of model mechanics—the model verification of Sections 8.2.3, 8.2.4, and 8.2.5—can add value.

FRB (2011) seems quite clearly to endorse independent review of trading models. Its Section III, which examines the criteria for which models need to be subject to the review standards of the document, states, “Models meeting this definition might be used for analyzing business strategies, [and] informing business decisions” and “The definition of *model* also covers quantitative approaches whose inputs are partially or wholly qualitative or based on expert judgment, provided the output is quantitative in nature.”

Managing Spot Risk

S*pot trades* are trades that involve an immediate exchange. This includes trades such as purchases of stock, purchases of gold, and exchanges of one currency for another. It excludes trades that involve a promise to deliver at some future time. Most of our study of risk involves future promises to deliver—unconditional promises constitute *forward transactions*, and promises whose payments are predicated on some future condition constitute *options transactions*.

The mathematical modeling and risk management of forwards and options are far more complex than the corresponding elements of spot transactions, and far more space in this book is devoted to forwards and options than to spot positions. However, positions in spot trades often constitute the largest portion of a firm's risk. Spot transactions are also the fundamental building blocks for valuing and risk managing forward and option positions. We can find the present value equivalent of a set of forward cash flows or the delta equivalent of an options position, but we then need to be able to value and risk manage these resulting spot positions. So a brief survey of the management of spot risk is in order.

9.1 OVERVIEW

All instruments traded by financial firms are *commodities* in the sense of not being individually identifiable. (If I borrow—that is, rent—a house from you, you expect me to return that exact same house, so houses are not a commodity; this is not true for dollar bills, bars of gold, barrels of oil, shares of IBM stock, specified amounts of a given bond, and so on.) This commodity feature means that traders are free to sell before they buy, since they can always borrow the instrument in order to make delivery. In this way, financial markets are more symmetrical than noncommodity markets such as houses, where you must build up an inventory by buying before you can sell.

Commodities can be divided into *physical commodities*, such as gold and oil, and *financial commodities*, such as stocks, bonds, and currencies. We do not study any trading in bonds in this chapter. Since bonds represent a fixed obligation to deliver an amount of currency, they are studied in Chapter 10 on managing forward risk. A general convention in the market is to use the term *commodities* to mean physical commodities only. Financial commodities are now almost universally transferable from one location to another in electronic form, so they have negligible transportation and storage costs per unit. Physical commodities have nonnegligible transportation and storage costs, which will have consequences we will study shortly.

Let us begin by looking at the hedging activities of a market maker in the dollar versus yen spot foreign exchange (or to adopt the terminology of that market, USD–JPY FX). In terms of instruments used, this represents the simplest type of trading possible—it is completely one-dimensional. The trader's position at any point in time can be represented as either long or short a certain quantity of JPY (or, completely equivalently, short or long a certain quantity of USD). In a more complex spot market, such as the commodities market for wheat, a trader's position would need to reflect being long or short different grades of wheat. However, currencies do not have grades—\$1 million is \$1 million, whether it is made up of 10,000 \$100 bills, 100,000 \$10 bills, 1,000,000 \$1 bills, or 100,000,000 pennies.

Our market maker will receive orders throughout the day from customers who are either looking to sell JPY and buy USD or looking to sell USD and buy JPY. Each customer will state the quantity of USD she wishes to sell and ask for a bid of the quantity of JPY that the market maker will exchange for it, or state the quantity of JPY she wishes to sell and ask for a bid of the quantity of USD the market maker will exchange for it. Trading screens are available at all times that show the best bids currently available from other market makers for selling JPY in exchange for USD and for selling USD in exchange for JPY. Market makers are constantly submitting their own bids for these two trades for the consideration of other market makers. When a customer's inquiry is for a small enough quantity, the market maker can guarantee a profit by quoting a bid just slightly higher than the best bid currently quoted on the trading screen, and if the customer accepts the bid, the market maker will immediately be able to close out the position created by hitting the bid quoted on the trading screen and making the small differences between the two as profit.

The market maker is only required to decide how much of a margin to build into the quote to the customer. The higher the margin, the higher the profit, but the greater the chance that the customer will turn down the quote and seek a quote from another market maker. The size of margin quoted must depend on the market maker's knowledge of the customer—how likely

is this customer to be polling a large number of market makers simultaneously rather than just coming to a single firm seeking a quote? In practice, the decision making at a firm will probably be divided up between a trader and a salesperson. The salesperson, who has a close knowledge of and continuing relationship with the customer, will bear the primary responsibility for determining the size of margin quoted. The trader will be credited, in the internal record keeping of the firm, with only a small portion of this margin.

A trader who followed this risk-averse a strategy would be unlikely to retain a job for long. The firm would probably judge that the profit the trader was making for the firm was not worth the opportunity cost of the trading seat. Higher profits would likely come from giving the seat to a more aggressive trader who would choose to take some risk by not closing positions out immediately. It is true that more aggressive traders are running the risk that prices will move against them, but, assuming that the firm sees a decent flow of customer orders, it is likely that a customer order will soon come in on the other side, and, on average, over time, the spread between the bid on each side of the market will be greater than losses from price movement through time.

When a large customer order comes in, then the market maker has no choice but to take some risk—the only choice is how to divide the risk between the liquidity risk of trying to offset the position immediately and the basis risk of offsetting the position over time. With a large order, the trader can no longer count on being able to close the position out at the price posted on the trading screen since this quote will only be for a reasonably small transaction. Of course, the customer will be charged a premium for the liquidity risk posed by the size of the order, which will provide some cushion to the trader against the risk that must be taken. The trader needs to make a judgment as to the relationship of this large customer order to overall market conditions. Is it an order that simply reflects the idiosyncratic circumstances of this customer, perhaps a payment that needs to be made in the customer's business? In this case, it is unlikely that a relationship exists between the order and any price trend in the market. Unless the trader has some other reason to believe that the market will be trending in a direction that will cause losses to this position, it will be better to close the position slowly, relying on customer orders and small trades with other market makers, minimizing liquidity risk. However, if the large customer order is likely to be part of a large movement, such as a customer wanting protection against the announcement of economic data that may impact the market, it may be better to close the position more quickly, bearing some liquidity cost in order to reduce the exposure to market trend.

Almgren and Chriss (2001) show how to calculate the efficient frontier of strategies that have the optimal trade-off between the liquidity costs of

offsetting the position in large blocks and the volatility risk (which we call *basis risk*) that the price at which the offset occurs differs from the price at which the position was put on. In the absence of price drift, the strategy that minimizes liquidity cost is one in which position covering is spread out over as long a period of time as possible, minimizing transaction size, and the strategy that minimizes volatility risk is one in which the entire position is offset at once, with as little chance for prices to change as possible.

Thus far, we have pointed out two advantages of seeing good customer order flow to a market-making firm: the increased likelihood of closing out positions at the favorable side of the bid spread and knowledge about the motives behind large orders. There are other advantages as well. Working with customers closely enables a firm to anticipate a large order and allows positions to accumulate through customer flow to meet part of the order in advance, thereby further lowering liquidity risk. When a firm's traders have a market view and want to put on a position, customer order flow enables them to put positions on and close out the positions more cheaply than if all positioning had to be done by aggressively seeking bids from other market makers. All of these advantages of customer order flow and the trade-offs of liquidity versus basis risk are present in all market-making activities, but can be observed in their purest form in spot risk market making, where other complicating factors do not intrude.

Even for the simplest spot product, FX spot, positions can be closed over time in other possible ways. For example, another source of liquidity is to spread out the closing of the position between the spot FX market and forward FX markets. This introduces a new basis risk in the form of the risk of unfavorable interest rate movements between the time the forward position is put on and the time it is closed out, but lowers the time basis risk. The trader must judge which is the most favorable risk mix. A trader in the currency of a smaller economy, let us say one trading the Danish krone against the dollar, might choose to temporarily hedge some of a position by a euro-USD trade that will eventually be closed out by a krone-euro trade. Adding a leg to the trade adds transaction costs, but euro-USD has more liquidity than krone-dollar and the trader's judgment may be that the basis risk of a krone-euro position is considerably smaller than that of a krone-USD position, given the closer tie of the Danish economy to the economy of the euro bloc countries than to the U.S. economy. When we move to more complex spot products such as commodities or equities, the potential avenues for redirecting basis risk multiply enormously. A position in IBM stock could be temporarily hedged by a Standard & Poor's (S&P) index future, judging this basis risk to be smaller than an outright IBM stock position. A position in one grade of wheat could be temporarily hedged with a position in another grade of wheat that trades with greater liquidity.

Firm-level risk management for spot risk is relatively straightforward. The more liquid spot positions can be valued by directly obtaining market prices. As a result, it is not necessary to utilize models for valuation and to establish reserves against possible model errors. Most spot markets are liquid enough that prices can be obtained from trading screens or closing prices on public exchanges, so it is not even necessary to arrange for a price collection from brokers. For market-making trading desks with reasonable customer order flow, positions should be marked to midmarket, since the presumption is that, on average, most positions can be unwound without needing to aggressively seek bids from other market makers. The only adjustment that might arise with any frequency is a reserve against liquidity risk if a spot position grows sufficiently large relative to the size of customer order flow that significant liquidity costs may arise in closing the position. For proprietary trading desks, positions should generally be marked to the side of the bid-ask spread that is least favorable for the position, since, in the absence of customer order flow, it should be presumed that closing out the position will require aggressively seeking bids from market makers.

Less liquid spot markets may require some form of modeling for valuation purposes. For example, an over-the-counter stock that does not trade very often or a commodity grade that is thinly traded may not have readily available price quotes. A model may need to be established that relates this price to the price of a more liquid instrument. For example, the over-the-counter stock price could be priced in relationship to a stock index, or a less liquid commodity grade could be priced as a spread to a more liquid commodity grade. In this way, the valuation can be updated daily based on quotes for the more liquid instruments. The relationship can be reestimated less frequently as reliable trading prices for the less liquid instrument are obtained. When models of this type are used, a reserve is needed against the statistical uncertainty of the relationship between liquid and less liquid prices being utilized.

The issues of nonstatistical limits and risk reporting to senior management for spot positions center completely on issues of which positions should be grouped together, since the position in any particular spot instrument is a single number. We'll discuss this issue for each of the spot markets: first FX, then equity, and finally physical commodities.

9.2 FOREIGN EXCHANGE SPOT RISK

To consider a concrete example, a USD-based firm will want to limit and report to senior management its net FX spot exposure to USD. This firm

will also want to have individual currency limits for FX spot exposure for every currency it trades. It will set limit sizes relative to the overall liquidity of the market for that currency and the firm's degree of customer order flow in that currency to ensure that traders have explicit management approval to build up positions that will require large time periods to reverse. However, senior management would probably need to be informed only of the largest individual currency positions. The remaining decision is determining which currency groupings are the best to use in setting net FX spot exposure limits and reporting to senior management. For example, does a grouping of all-Asian currencies make sense? A grouping of all-Asian currencies excluding the yen, Australian dollar, and New Zealand dollar? Should Asian currencies be divided into groupings based on national gross domestic product (GDP) per person? Should all currencies of countries with lower GDP per person be grouped together as emerging market currencies? Each firm will reach its own conclusions based on economic theory, trading experience, and, perhaps, statistical analysis of which currency movements tend to occur together.

9.3 EQUITY SPOT RISK

Equity reporting and limits can begin from a similar starting point as for FX. There should be reporting and limits for positions in individual stocks, for an overall long (or short) net position in all stocks, and for groupings by geographic region. Decisions on whether to group together stocks in all companies based in Europe or based in emerging markets is subject to the same type of analysis as the decisions for FX.

But geography is just a starting point for stocks. There are several other considerations: industry and industry sectors, and *style*. Much research has been devoted to which factors play the largest role in explaining the performance of equity managers, an issue that is known as *performance attribution*; the classic article in this area is Sharpe (1992), which was highly influential in the recognition of the importance of stocks of smaller-capitalization firms versus larger-capitalization firms and growth stocks versus value stocks as important style attributes in explaining performance. Much of this analysis translates very directly to how to group stocks together for purposes of risk reporting and limits.

Here are examples of some popular classifications for performance attribution:

- Morningstar, in its evaluations of mutual funds that invest in equities, has created a very influential style box based on smaller-capitalization

firms (less than \$2 billion) versus larger-capitalization firms (more than \$10 billion) and growth stocks versus value stocks. Morningstar follows Sharpe (1992) in defining growth stocks as those with little or no dividend payout, high price-to-book and price-to-earnings ratios, and promising capital appreciation, and value stocks as those likely to pay high dividends but with low price-to-book and price-to-earnings ratios.

- The Global Industry Classification Standard (GICS) developed by Morgan Stanley Capital International (MSCI) and Standard & Poor's (S&P) consists of 10 sectors, 24 industry groups, 68 industries, and 154 subindustries. The 10 sectors are energy, materials, industrials, consumer discretionary, consumer staples, health care, financials, information technology, telecommunication services, and utilities.

9.4 PHYSICAL COMMODITIES SPOT RISK

Physical commodities are further complicated by the presence of transportation costs, which leads to different markets for the same commodity in different locations (for example, oil for delivery in Seattle is a different product from oil for delivery in El Paso). This plays a role in valuation, since delivery at a location where liquid prices are not available could be priced using a model based on a more liquid price for delivery at another location and estimated transportation cost between the two locations. It also plays a role in the design of limits and reporting. Locations that are reasonably closely related in price, by having low transportation costs between them, should have their positions summed into a net position for reporting and perhaps limits.

An interesting analogy can be made between location relationships based on transportation costs and relationships between forward prices for different time periods. In Section 10.3.2, we will see that some commodities have forward prices for different times tightly linked by the possibility of cash-and-carry arbitrage. It is instructive to think of this as a form of location relationship, with the storage and financing costs as the cost of “transporting” the commodity from one time period to a later one. Just as transportation can be so expensive between some locations that they virtually form independent markets, storage can be so expensive for some commodities, such as electricity, as to virtually eliminate the possibility of cash-and-carry arbitrage. However, although transportation costs are almost always symmetrical (it costs just as much to ship from A to B as from B to A), a commodity cannot be transported from a later period to a former period, so cash-and-carry arbitrage works only in the forward direction.

Other types of potential transformations besides location play a role in physical commodities. To take two examples from McDonald (2006, Chapter 6):

- Soybeans can be crushed to produce soybean meal and soybean oil. A trader with a position in soybean futures and an opposite position in equivalent quantities of soybean meal and soybean oil futures is trading the *crush spread*. The trader is taking a position not on what will happen to the cost of soybeans but on what will happen to the cost of processing soybeans into soybean meal and soybean oil. To the extent that positions in soybeans and soybean meal and soybean oil offset, the resulting position should be reported and limits set on the crush spread and not on the individual legs.
- Crude oil can be separated into different petroleum products such as heating oil and gasoline by a refining process known as cracking. A trader with a position in crude oil futures and an opposite position in equivalent quantities in heating oil and gasoline futures is trading the *crack spread*. The trader is taking a position not on what will happen to the cost of crude oil but on what will happen to the cost of processing crude oil into heating oil and gasoline. To the extent that positions in crude oil and heating oil and gasoline offset, the resulting position should be reported and limits set on the crack spread and not on the individual legs.

Reports should also be designed and limits set on aggregated positions across physical commodities whose prices tend to be highly correlated. So there might be an overall limit on total net long (or short) exposure to all energy products, summed over crude oil, heating oil, gasoline, natural gas, and electricity. Which products get grouped together may differ by firm, based on economic theory, trading experience, and statistical analysis.

EXERCISE

9.1 Simulation of the Impact of Trading Rules on Expected Return and Risk

A market maker in a spot market is operating under the constraint that she must close out her position by the end of each trading day. We want to see the impact of different possible trading limits on the size of the position that can be built up.

Divide the trading day into 100 time segments. In each time segment except the last, there is a 50 percent chance of receiving a customer order for one unit. A customer order has a 50 percent chance of being a buy and a 50 percent chance of being a sell.

Customers pay \$0.10 per trade in transaction costs. So if the midmarket price is \$100.00, a customer will purchase at \$100.10 and sell at \$99.90.

The market maker cannot close out a trade without waiting at least one period. Midmarket price changes from one period to the next are normally distributed with a standard deviation of \$0.10 (assume a starting midmarket price of \$100.00). The market maker must close out her open position by the last trading period. She pays \$0.05 per trade in transaction costs to close positions with another market maker. So if the midmarket price is \$100.00, she sells positions at \$99.95 and purchases at \$100.05.

It is to the market maker's advantage if she can wait until a customer order comes in to close out her position, since she will make a \$0.10 transaction spread on each side of the trade, for a total of \$0.20, rather than making only \$0.10 minus \$0.05, for a total of \$0.05 in transaction spread by closing out with another market maker. However, the longer she waits for a customer order, the greater her risk of prices moving against her.

Simulate a set of trading rules to see the trade-off between expected return and risk. Use 1,000 paths for each simulation. The measure of expected return should be simply the average over these paths. You can choose any reasonable measure of risk, such as the 95th percentile loss or the standard deviation. One trading rule should be to never close out until the last period. Another should be to always close out in the period immediately after the customer trade. Intermediate rules can be based on a limit of how large the absolute size of a position is allowed to grow—when the position gets larger than this limit, the excess must be closed out.

1. Determine the impact on the risk/return trade-off of a lower standard deviation of the midmarket price of \$0.05 per period.
2. For a more extended exercise, you could experiment with more complex trading rules, such as having the transaction cost for closing a position be an increasing function of the absolute size of position to be closed, or allowing the market maker to influence the probability of customer trades being buys or sells by shifting her quoted price away from the midmarket price.

Managing Forward Risk

Managing forward risk is considerably more complex than managing spot risk due to the large number of dates on which forward payments can take place. With some forward markets going out to 30 years and even beyond, even if we restrict deliveries to take place on the 250 business days of a year, it still leaves $30 \times 250 = 7,500$ days on which future flows can occur, each of which requires a mark-to-market valuation and risk measurement. It is clearly impractical to have liquid market quotations for each possible forward, so modeling needs to be heavily relied upon.

Having a spot versus forward position is an interest rate differential position, not a price view. If I believe the market will get a surprise announcement that will raise the stock price, even if I think it will not come for three months, I don't want to be long the forward and short the spot. When the announcement comes, both will be roughly equally impacted. I want this position only if the announcement I expect is something like a one-shot dividend that will impact the relative value of the spot and forward. If I put on a long forward and short spot position, I'm taking a view on the interest rate.

Let me cite a real example. On June 24, 1998, a trader was holding a long forward position in Telecom stock against which he was short the stock. AT&T announced plans to purchase Telecom at a sizable premium, but the trader wound up with a sizable loss. Why? His outright position in Telecom stock was even, so he didn't gain from the rise in the stock price. Telecom had never paid a dividend, so the forward traded at a large premium to the cash. As soon as the market anticipated that the stock could be traded for a dividend-bearing AT&T stock, this forward-to-cash premium shrunk significantly since it was now less expensive to hold a cash position in the stock for delivery into a forward sale.

The difference between an outright position and a borrowing or lending position is the difference between wanting to hold an asset as a good investment (you expect it to gain value) versus wanting to make use of an asset. Consider a house. When you buy it, you get a combination of an investment

and a place to live. You might want to split the two. If you like it as an investment but don't want to live there, you can buy it and lend it to someone (rent it out). If you want to live in it but don't like it as an investment, you should borrow it (rent it) rather than buy it.

Similarly, a firm that is in the business of milling wheat and is running short of wheat supply to keep its production process going but does not like wheat as an investment (does not believe it will go up in price) will seek to borrow wheat rather than buy it (although borrowing may take the form of buying spot wheat while selling forward wheat). Likewise, a firm that likes wheat as an investment but does not need it for any production process will buy wheat and then lend it out (possibly combining the two steps into one by buying forward wheat).

Although a clear distinction can be made between an outright spot position and a borrowing or lending position, they also share close relationships. As we saw in the discussion of spot risk management in Chapter 9, maintaining a spot risk position over a longer period than a single trading day requires some form of borrowing or lending. In some markets, the use of borrowing or lending to maintain outright spot risk positions becomes such a dominant force that it is the principal driver of interest rate movements in the market. In many trades, such as forward purchases and sales, spot and forward risk are bound together, so it will be necessary to study the interactions between these two risks to fully understand the dynamics of forward risk management. It is important for the risk management function to clearly separate spot risk from forward risk in transactions in which they are bundled to ensure that all the firm's spot risk in a given asset is reported and managed in a unified fashion.

The borrowing and lending markets in currencies and gold started as a means for businesses and individuals to adjust the timing between when income is earned and when purchases are made. Borrowing and lending in other commodities started with users and suppliers of the commodity satisfying short-term needs, as in the previous milling example. Borrowing in stocks and bonds started with the need for short sellers, who want to act on the view that an asset will decline in value, needing to first borrow what they wanted to sell short. Borrowing to support short selling is also a feature of all the other borrowing markets.

Once borrowing and lending markets are established, they begin to attract investors, speculators, and hedgers who have views on the borrowing rate rather than on the asset price. So one trader who believes that a particular borrowing rate will soon decline will lend at that rate solely in hopes that he can match that lending with a borrowing at a lower rate when the rate declines. Another trader might believe that the borrowing rate for May 2015 is too high relative to the borrowing rate for April 2015 so she

will borrow for April and lend for May, hoping to reverse the transactions when rates return to a more normal relationship. Another trader might believe that borrowing rates for a particular corporation will decline relative to those of another corporation or the government, so he will lend to the former by buying its bond and borrow to support a short sale of the latter's bond. A business firm worried about the possible impact of high borrowing costs on its financial health in 2017 will borrow funds now that do not become available until 2017.

The emphasis I am placing on borrowing and lending rates as the foundation of forward risk is somewhat nonstandard; but see Williams (1986) for an incisive economic analysis of forward, futures, and lending markets for commodities using this approach; also see Brown (2012, Chapter 10) for an excellent discussion along similar lines. A more conventional exposition, such as Hull (2012, Chapter 5), would focus on borrowing rates only for currencies and would analyze forward risk on commodities and securities through forward contracts that involve exchanging the commodity or security for currency. The borrowing rate on the commodity or security still comes into play as one of the inputs determining the price of the forward or implied by the price of the forward.

The two methods are mathematically equivalent, so choosing between them is a matter of deciding which is the most convenient and supplies the greatest financial insight. My choice of emphasis is based on the following considerations:

- Direct borrowing and lending markets exist for many assets—such as gold, stocks, and government bonds—that do not require any involvement with borrowing/lending risk on currencies. Let's look at an example. Suppose that the rate for borrowing gold for three months is 2 percent annualized. If I want to borrow 1,000 ounces of gold today, I must be prepared to return $1,000 \times (1 + 2\% \times 3/12) = 1,005$ ounces of gold in three months. No mention has been made of any currency—there is simply an equivalence of a certain amount of gold on one date and some other amount of gold on another date.
- A uniform approach to all underlying instruments makes for easier exposition of some concepts. For example, Section 10.2 on mathematical models for forward risk is built around a single discount curve that could represent borrowing costs for a currency, but could represent borrowing costs for a security or commodity equally well.
- It is consistent with a risk management viewpoint in which, for example, it is natural for a gold trader to be taking risk with regard to gold borrowing rates, but not with regard to dollar borrowing rates. Gold borrowing costs are primarily impacted by economic factors unique to

the gold market, including the supply and demand for gold, so it would be a sound risk management practice for the same trading desk to run risks in the gold spot and borrowing rates. However, there is little linkage between gold and dollar borrowing rates. A gold trader running dollar borrowing risks through the vehicle of positions in gold/dollar forwards requires serious management scrutiny. At a minimum, dollar interest rate exposures taken in this way need to be reported and aggregated together with other dollar rate risks throughout the firm. Similar comments apply to borrowing risk on other commodities and securities.

The primary argument against a borrowing rate focus is that for some assets, such as oil, no borrowing market exists, requiring forward risk to be managed through forward contracts. Even for some assets for which a borrowing market does exist, the borrowing market has considerably less liquidity than the comparable forward contract. However, it is always possible to take spot and forward prices and currency interest rates and derive implied asset borrowing rates that can then be used just as if they had been obtained by a direct quote. Indeed, even in some currency markets, the most liquid source for rate quotes is to combine forward foreign exchange (FX) prices with dollar rates to derive interest rates for the currency. This is no bar to developing discount curves for the currency or combining directly obtained rates that are the most liquid price source for some maturity segments with implied rates for other maturity segments and using them to form a single discount curve.

Within the fixed-income departments of investment banks, it is customary to find separate trading desks for interest rate and credit products, with interest rate trading focused exclusively on changes in credit risk-free rates and credit trading focused exclusively on changes in the credit spread to risk-free rates. Clearly, some products cut across this boundary—a fixed-rate bond issued by a corporation will change in value because of both changes in risk-free rates and changes in credit spreads. But interest rate swaps that convert fixed-rate into floating-rate payments can be used to transform a fixed-rate corporate bond into an instrument that is almost totally dependent on credit spread, so trading desks can utilize internal transfers to almost completely separate the two types of exposure.

We will want to follow this division in studying risk. While there is a certain amount of overlap between interest rate risk and credit risk measurement and modeling, particularly in extracting term structure from market prices, the differences are greater than the similarities:

- Option products are very important instruments in interest rate trading, requiring the modification of traditional option models to more

complex multiple-tenor environment. Option products are currently of negligible importance in credit trading.

- Credit modeling focuses on correlation between debt and equity within a firm and between debt of different firms. There are no comparable issues for interest rate models.

Consequently, we will focus only on products free of credit risk in this chapter, reserving the study of credit risk management for Chapter 13.

Strictly speaking, it is only bonds issued by the central government (for example, U.S. Treasury bills and bonds in the United States) that are (nearly) completely free of credit risks. It is only the central government that has unlimited power to issue its own currency and so can (nearly) certainly meet any obligations to pay that currency. But fixed-income trading desks of investment banks generally also trade a variety of instruments whose credit risk is extremely low: bonds issued by agencies of the central government, mortgages guaranteed by such agencies, and derivatives tied to bank indexes such as the London Interbank Offered Rate (LIBOR). This latter case is a particularly important class for interest rate products; indeed, the largest interest rate risk exposures of financial institutions is usually to LIBOR products: LIBOR futures, forward rate agreements, swaps, caps, floors, and swaptions. So we ought to examine closely why credit risk on these products is considered negligible and why it predominates over Treasury rates as the basis for derivative products.

First, we need to distinguish between the credit risk to the counterparty on a derivative and the credit risk on the derivative instrument itself. Consider the example of a typical derivative tied to LIBOR, a 10-year interest rate swap of fixed coupon payments against three-month US Dollar LIBOR reset each quarter. Certainly there is credit risk that the counterparty will default on its obligations under this contract, with the severity of risk tied to the creditworthiness of the counterparty. But this has nothing to do with credit risk of the swap itself. An equity option or foreign exchange swap or option on a Treasury bond would also entail counterparty credit risk but no underlying credit risk. By contrast, a default swap in which you must pay Company A an agreed amount based on the default of Company B against fixed payments to you from company A entails both counterparty credit risk of losing your fixed payments if company A defaults and underlying credit risk of Company B defaulting.

So we need to see whether there is any underlying credit risk on a bank index product. Let us continue with our example of the 10-year swap based on three-month US Dollar LIBOR resets. There is clearly some credit risk—a severe economic downturn will raise concern about potential bank defaults and therefore raise the rate that banks need to pay on three-month deposits

relative to three-month Treasury bill rates. But to see just how small this element of credit risk is, let us contrast it with the credit risk on a 10-year bond issued by one of the banks whose deposit rates form the LIBOR index, noting that the credit risk on this bond is very small to begin with, since all banks in the index are of very high credit quality, generally Aa. The credit spread on the 10-year bond needs to reflect the probability of default over a 10-year period, which includes scenarios in which the creditworthiness of the bank declines severely prior to the default. But these scenarios will have little impact on the average LIBOR index over the 10 years, since a bank that declines in creditworthiness will be replaced in the panel that determines the LIBOR index. For example, Moody's data for a 20-year period shows that 0.81% of Aa-rated firms defaulted within 10 years of the rating, but only 0.02% of Aa firms defaulted within one year of an Aa rating. Furthermore, even in the event of default, the chances of depositors losing money are very small since bank regulators are primarily concerned with protecting depositors and take steps to ensure that losses will be absorbed by stockholders and bondholders but not depositors. Spreads between LIBOR-based rates and Treasury rates therefore primarily reflect the superior liquidity of Treasuries and their value as collateral. Under the very extreme conditions of the global banking crisis of 2007–2008, there was a period in which a spread between LIBOR and Treasury rates based on credit concerns came to play a major role (see Tuckman and Serrat 2012, 431–432), but this is a very rare occurrence.

Individual government issues have idiosyncratic characteristics (liquidity, borrowing rates, country of issue for euros) getting in the way of creation of a single discount curve against time. This is closely tied to government bonds being in fixed supply as opposed to swaps, which can be freely created. Government rates represent only investment rates for most firms and not borrowing rates (you can only borrow at government rates if you have government bonds available as collateral) while deposit rates are two-sided, at least for the large financial institutions that serve as market makers. This has led to LIBOR being the risk-free rate generally used to price derivatives, such as futures, forwards, interest rate swaps, interest rate options, and credit swaps, and as a target against which to measure borrowing rates. This explains why LIBOR is far more popular than government rates as a basis for derivatives used to hedge interest rate risk. As this book is going to press, news stories about manipulation of the LIBOR rate setting process are raising questions that could threaten the popularity of LIBOR as a basis for derivatives. As this story develops, its consequences will be addressed on this book's website.

A good summary of the issues raised by this manipulation of LIBOR rate setting is the article "The Rotten Heart of Finance" in the July 7, 2012,

issue of the *Economist* magazine. The website of the British Bankers Association (www.bbalibor.com), the organization in charge of determining LIBOR, has many articles giving details of the LIBOR-setting process. When derivative contracts based on bank deposit rates were designed, a significant worry was that if a derivative referenced the rate set by a particular bank, that bank might manipulate the rates at which it bid for deposits in order to generate profits in its derivatives holdings. The decision was made to tie derivatives products to an index of bank deposit rates, which would be harder for any one bank to manipulate. A large panel of banks is selected (currently 16 for US Dollar LIBOR), based on criteria of expertise and prominence in the market and the highest degree of credit worthiness (any bank no longer meeting these requirements would be replaced in the panel). The highest and lowest quartiles of submitted rates are dropped, to minimize any potential for manipulation, and only the middle two quartiles averaged (in addition, any bank not operating within the spirit of the rules would be dropped from the panel). But when the global banking crisis of 2007–2008 caused a large decline in the use of interbank deposits, the lack of market liquidity may have opened the door to potential manipulation.

Given the complexities of forward risk management, we will need to carefully organize our study into the following sections:

- **Section 10.1.** This is a study of the variety of instruments that entail forward risk and that can be used to manage forward risk. The large variety of structures in which spot and forward risk (and occasionally implicit options risk) are woven together means that an important part of risk analysis is often just making sure that all the risks of a particular trade have been properly identified. In addition to the market risks, slight variations in structure, which may result in virtually identical spot and forward risk, can have large differences in credit risk, legal risk, and funding liquidity risk.
- **Section 10.2.** This section provides a study of the mathematical models used to value and measure forward risks. Although these models have been used heavily for many years and a great deal of consensus has been built up around them, enough subtle issues remain to merit a careful understanding of the residual risks of model uncertainty.
- **Section 10.3.** This section takes a brief look at the factors that impact borrowing and lending costs. Although this is not primarily a book about economics, at least some familiarity with the determinants of forward prices is necessary to properly understand the requirements for designing a risk management structure for forward risks.
- **Section 10.4.** This section provides a study of how to design a risk management reporting system for forward risk.

10.1 INSTRUMENTS

The management of forward risk can involve a number of different instruments that can be used to take on the same market risk position. These instruments may differ in legal form, with different regulatory consequences and standing in bankruptcy proceedings, and have different implications for credit risk and funding liquidity risk. They also differ in the extent to which they bundle together spot and forward risk.

We consider each of the following categories:

- Direct borrowing and lending
- Repurchase agreements.
- Forwards.
- Futures.
- Forward rate agreements (FRAs)
- Interest rate swaps
- Total return swaps
- Asset-backed securities

10.1.1 Direct Borrowing and Lending

Suppose a trader wants to sell a given asset short. In a number of asset markets—such as stocks, bonds, currencies, and gold—the asset can be borrowed directly in order to sell short. Other markets, such as most physical commodities, have not developed direct borrowing products.

One drawback to using borrowing as the means of obtaining an asset to short is that it creates a sizable credit risk and funding liquidity risk for the asset lender, who could lose the entire value of the asset if the borrower defaults and who has to finance the asset that has been lent. The borrower may be paying for credit usage that is not really needed, since the cash raised by selling the asset short is incidental to the original objective of selling the asset short to position for a price drop. One solution is to use the cash raised as collateral against the borrowing. This reduces the credit risk for the asset lender, who can hold on to the cash collateral in case of borrower default, and reduces the funding liquidity risk, since the cash collateral received by the asset lender can be used to fund the asset purchase.

Providing cash collateral to the asset lender creates credit risk for the asset borrower, even though this is mitigated by the value of the asset, which does not need to be returned if the recipient of the cash collateral defaults.

10.1.2 Repurchase Agreements

In the previous example, one party borrows the asset and provides cash collateral to the other party. An equivalent way of describing the same trade is to say that one party borrows the asset and lends cash, whereas the other party borrows cash and lends the asset. Yet another equivalent way of describing the same trade is a transaction in which the first party purchases the asset for cash and, at the same time, contracts to sell the asset back to the second party at an agreed forward date for an agreed cash price. Table 10.1 demonstrates that all three ways of describing this transaction are equivalent in terms of the flows of cash and asset.

TABLE 10.1 Alternative Descriptions of an Asset Borrowing Collateralized by Cash

Description 1	
Today	A borrows \$1 million par amount of a Treasury bond from B. A sells the bond in the market and receives \$980,000. A places the \$980,000 as collateral with B.
One month from today	A buys the \$1 million par bond in the market and returns it to B. B returns the \$980,000 collateral to A. A pays \$1,000 in interest for borrowing the bond to B. B pays \$5,000 in interest for the use of the cash to A.
Net effect	A delivers \$1 million in par amount of the Treasury bond to A. B pays $\$980,000 + \$5,000 - \$1,000 = \$984,000$ in cash to A.
Description 2	
Today	A borrows \$1 million par amount of a Treasury bond from B. B borrows \$980,000 in cash from A.
One month from today	A repays the \$1 million par Treasury bond loan to B plus \$1,000 cash in interest on the loan. B repays the \$980,000 in cash to A plus \$5,000 in interest on the loan.
Net effect	A delivers \$1 million par amount of the Treasury bond to A. B pays $\$980,000 + \$5,000 - \$1,000 = \$984,000$ in cash to A.
Description 3	
Today	A purchases \$1 million par amount of a Treasury bond from B for \$980,000 in cash.
One month from today	B buys the \$1 million par amount of the Treasury bond from A at the prearranged price of \$984,000.

The third description, which is known as a *repurchase agreement*, possesses some legal advantages in the event of default. If the party lending the asset defaults, the other party technically owns the asset, since it purchased it rather than just borrowing it, so it has fewer legal restrictions on its ability to use the asset. If the party borrowing the asset defaults, the party lending the asset, since it technically sold the asset and received cash as payment rather than just as collateral for the borrowing, has fewer legal restrictions in its ability to use the cash.

10.1.3 Forwards

A *forward contract* is an agreement to pay a fixed price on a set forward date for a specified amount of an asset. As such, it combines into a single transaction borrowing the asset and then selling the asset in the spot market. The seller of the forward needs to deliver the asset at a fixed forward date and price, exactly as a borrower of the asset must do. The seller of the forward is at risk for increases in the asset's price and will gain from decreases in the asset's price, just like a borrower of the asset who sells it in the spot market. The buyer of the forward is in the same position as a buyer of a spot who lends out the underlying asset but does not need to fund the currency to purchase the asset. Since no cash will change hands until the forward date, it does not have the credit and funding liquidity risks that an uncollateralized borrowing of the asset would have. From a credit risk standpoint, a forward transaction is very similar to a borrowing of the security that has been collateralized by cash.

Credit risk on either a forward or a borrowing collateralized by cash starts close to zero, but can build up as the market price of the underlying asset goes up or down. Managing this counterparty credit risk can be quite complex, as the amount of credit exposure varies through time and is correlated with movements in a market price. We will not be fully prepared to address this issue until Chapter 14 on counterparty credit risk. For now, we will just note that a frequently used approach to mitigate this credit exposure is the collateral call, in which the borrower and lender agree in advance that upward moves in the asset price will require the asset borrower to increase the amount of collateral placed with the lender, and downward moves in the asset price will require the asset lender to increase the amount of collateral placed with the borrower. This cross-collateralization agreement is an automatic feature of futures contracts, which are very closely related to forward transactions.

10.1.4 Futures Contracts

Futures contracts are identical to forward contracts in their market price exposures. They also specify the payment of a fixed price on a set forward date

for a specified amount of an asset. They differ from forward transactions in two primary dimensions: the management of counterparty credit risk and the degree to which they are tailored to trade off basis risk versus liquidity risk. We briefly discuss both aspects.

Unlike forward transactions, which are direct transactions between two firms or individuals, futures contracts are always arranged to have a futures exchange as one of the counterparties to each contract. So if Firm A agrees to sell 100,000 barrels of oil for delivery on June 15, 2015, to Firm B in exchange for \$3 million, it is technically broken up into an agreement for A to deliver 100,000 barrels of oil on June 15, 2015, to the futures exchange in exchange for \$3 million and an agreement for B to pay \$3 million to the futures exchange for the delivery of 100,000 barrels of oil on June 15, 2015. This greatly simplifies credit risk management for the firms, which do not need to worry about the creditworthiness of one another but only need to evaluate the creditworthiness of the futures exchange. This would involve enormous credit management problems for the futures exchange since it has credit exposure to every firm trading on the exchange, but it is managed by strict insistence on continuous cash payments to and from the futures exchange as the prices of the futures transactions rise and fall (details can be found in Section 14.2). This also requires sufficient initial collateral to reduce credit risk to a minimum. This has several significant implications.

Because of the continuous collateral calls, a firm using futures contracts will have constant inflows and outflows of cash as asset prices rise and fall. This has important consequences for both funding liquidity risk and market risk. The funding liquidity risk consequence is that if a firm is using futures contracts to offset transactions that do not have this cash settlement feature, it may lead to substantial funding needs. For details, refer back to the Metallgesellschaft (MG) case in Section 4.2.2. The market risk implication is that the constant cash payments create risk to the extent that payment amounts are correlated with the time value of the payments—see the discussion on convexity risk in Section 10.2.4.

This system of credit risk control requires that the terms of futures contracts be standardized with only a few possible delivery dates and assets that can be contracted for. This contrasts with forward transactions, which, as agreements between two firms, can be tailored to very specific forward dates and assets to be delivered. This freedom is permitted by the firms performing their own management of the credit considerations of the transaction. However, the futures exchange must have the ability to quickly close out any contract on which a counterparty cannot meet a collateral call. The ability to quickly close out contracts without a substantial risk of loss requires the liquidity derived from a few standardized contract terms.

The liquidity that results from this standardization can also be attractive to potential counterparties who may welcome the reduction in liquidity risk this offers. With only a few standardized contracts, it is easier to find good price valuations and close out positions that are no longer desired. The price of lower liquidity risk is, as always, heightened basis risk. A firm might desire to hedge flows on particular dates but need to accept a hedge with nearby standardized dates. It may also desire to sell short a particular asset—a particular grade of wheat, say—but need to accept a hedge with a related standardized grade. The maintenance of necessary liquidity may require the provision that several possible grades be deliverable, which requires formulas for determining how much of each grade must be delivered. However, changes in actual market conditions will differ from any set formula, resulting in profit opportunities and basis risks that may need quite complex modeling. For an example, see the discussion of the modeling of delivery options on Treasury bond futures in Hull (2012, Section 6.2).

10.1.5 Forward Rate Agreements

A *forward rate agreement* (FRA, pronounced “fra”) is a particular type of forward contract in which the asset to be delivered on the forward date is a borrowing with a specified maturity date, interest rate, and borrower. For example, it might be an agreement to deliver two years from today a \$200 million one-year deposit with Bank of America paying an interest rate of 6.50 percent. This means that in two years the buyer of the FRA will be able to deposit \$200 million with Bank of America in exchange for receiving \$213 million back at the end of the third year: $\$200 \text{ million} \times (1 + 6.50\%) = \213 million .

The standard practice for FRAs is to *cash settle*, meaning that no actual deposit of cash with Bank of America is expected; instead, a cash amount equal to the value of the deposit will change hands. In our example, if Bank of America is offering 5.00% on one-year deposits at the end of two years, the right to place a deposit at 6.50% is worth $1.50\% \times \$200 \text{ million} = \3 million . Therefore, the FRA seller owes \$3 million to the FRA buyer, which will be paid at the end of the one-year deposit period. (In most, but not all, cases, the payment will be made when the FRA settles, which is the beginning of the deposit period, not the end. However, the settlement price will be determined by the present value of the payment due, using the prevailing discount rates at the time of settlement. Economically, this is no different in value from receiving the payment at the end of the deposit period, but it has the beneficial effect of reducing credit exposure.) If Bank of America is offering 7.50% on one-year deposits at the end of two years, the requirement to place a deposit at 6.50% has a cost of $1.00\% \times \$200 \text{ million} = \2 million .

Therefore, the FRA buyer owes \$2 million to the FRA seller, which will be paid at the end of the one-year deposit period.

FRAs are valuable tools for managing forward risk since they can be used to lock in borrowing and lending costs for future time periods or take positions on rates rising or falling. They are almost wholly confined to rates offered on currency borrowings by very high-credit-grade banks, since they have developed primarily as tools for managing the cost of borrowing and lending currencies rather than tools for speculating on changes in credit quality. The most popular instruments are those tied to the deposit rates averaged over a panel of high-grade banks, such as the London Interbank Offered Rate (LIBOR), thereby reducing the link to credit quality even further. Some interest rate futures, such as LIBOR futures, are essentially FRAs contracted using futures rather than forward structuring.

10.1.6 Interest Rate Swaps

Standard *interest rate swaps* are equivalent to a package of FRAs. A very typical example would be a five-year swap for \$200 million settled quarterly with one party paying U.S. dollar LIBOR and the other party paying a fixed rate of 6.50%. This is equivalent to a package of 20 FRAs that are settled on each quarterly date for the next five years.

Interest rate swaps are extremely flexible instruments that can be tailored to specific customer needs. Although it is typical that the terms of each FRA that constitutes the package will be the same on all terms except the forward date, it is quite possible for customers to arrange swaps with rates, deposit lengths, and notional amounts that differ by period. It is also quite common to combine FRAs in different currencies into a single swap and combine FX forwards along with FRAs into a single swap. To better understand the customer motivation for these features, it is important to understand the relationship between bonds and interest rate swaps.

The primary initial demand for interest rate swaps, and much of the demand to this day, comes from issuers of and investors in corporate bonds. Most corporate bonds pay fixed coupons, as this represents the form popular with most investors. However, bond issuers may prefer to borrow at a floating interest rate rather than a fixed rate, either because they believe rates are likely to fall in the future or because they believe floating-rate borrowings are a better match to their overall asset-liability position. Some investors may prefer lending at a floating rate rather than at the fixed coupon on a bond, either because they believe rates are likely to rise or because they are primarily looking to take a position in the creditworthiness of the firm and don't want to combine this with a position on whether risk-free rates will rise or fall. For such clients, a fixed-for-floating single-currency interest

rate swap, which is just a package of FRAs for a single currency, can transform a fixed-rate bond position into a floating-rate one.

Another instance of interest rate swap demand arising out of the corporate bond market occurs when the currency a firm would prefer to owe debt in is different than the currency that is preferred by a segment of investors in the firm's bonds. A typical example would be a European firm that wanted to tap into investor demand in the U.S. market. The firm might prefer to have all its debt in euros, but most U.S. investors prefer to invest in dollar-denominated bonds. One solution is to have the firm issue a dollar-denominated bond, but then enter into a cross-currency interest rate swap in which the firm receives fixed dollar payments equal to the coupon payments it owes on the dollar-denominated bond and pays fixed euro cash flows. The firm would probably also want to combine this with an FX forward contract to exchange the euro principal it wants to pay on the maturity date of the bond for the dollar principal that it owes on the dollar-denominated bond. This combination is a standard product, a cross-currency interest rate swap with the exchange of fixed principal. The firm might also want to make floating-rate euro coupon payments rather than fixed-rate euro coupon payments for the reasons given in the previous discussion on single-currency swaps. Rather than execute two separate swaps, this can all be accomplished in a single fixed dollar for floating euro cross-currency swap. A cross-currency swap can therefore be a combination of a bundle of FRAs and a bundle of FX forwards. As such, it combines the spot FX risk of FX forwards, the forward risk of FX forwards, and the forward risk of FRAs.

10.1.7 Total Return Swaps

We have already seen how a cross-currency swap can combine FRA and forward positions in a foreign currency asset. *Total return swaps* are instruments that generalize this approach to enable forward positions to be taken in any asset. Instead of having an agreement to exchange a fixed amount of euros for a fixed dollar price on an agreed forward date, as might be the case in a cross-currency swap, an agreement might be made to exchange a fixed amount of an asset such as a bond or stock for a fixed dollar price on an agreed forward date.

The most common form of total return swap calls for the following. Party A makes a series of intermediate payments to Party B, usually tied to intermediate coupon payments or stock dividends. Party A delivers an asset to B on a fixed date for a fixed price. Finally, Party B makes a series of intermediate payments to A, usually tied to a funding index such as LIBOR. This form of total return swap is economically equivalent to a forward transaction. Like a forward, it combines into a single bundle the spot sale of an

asset and the borrowing of that asset for a fixed term. However, although a forward bundles together the sale price and borrowing costs into a single final fixed price, the total return swap makes the intermediate borrowing costs more explicit.

One major contractual difference between total return swaps and forwards is that a total return swap can be used by a party that might otherwise find it difficult, for legal or institutional reasons, to invest in a particular asset class. Although the forward contract generally calls for the actual delivery of the asset on the specified forward date, the total return swap will often call for a cash settlement based on the value of the asset on the specified date. This can be a necessity for a party that cannot legally own the asset (for example, the party may not have a subsidiary in the country in which the asset trades). This can still be a great convenience for a party that can legally own the asset but may not be well positioned to trade it. In effect, it is contracting out to the other party the actual sale, which makes sense if the other party is a major market player in this asset or if the asset is actually a participation in a portfolio of assets.

The downside of this arrangement is that it can lead to disputes as to what the actual settlement price should be in cases where there is not a publicly available and reliable pricing source. So although it may be easy to agree that the settlement of a basket of stocks traded on the New York Stock Exchange (NYSE) will end up at the published closing exchange prices for the day, it may be necessary to build elaborate legal processes for the settlement of a bond of limited liquidity. (For example, the processes could involve an appeal to a panel of other market makers or the right of the party receiving the value of the bond to take physical delivery in the event of failing to agree on a cash price.)

The primary initial impetus for the total return swap market came from parties wanting to purchase assets they would have difficulty holding. They were assets they either could not legally hold or would have difficulty trading, leading to the demand for a cash settlement discussed previously, or assets they would have difficulty funding. A firm wanting to purchase a bond with a higher credit grade than that of the firm could face the negative carry costs of having to fund at a higher credit spread than it can earn on the bond. To avoid this situation, one must find a way to borrow against the collateral of the security, as discussed previously in Sections 10.1.1 and 10.1.2. The total return swap offers the convenience of bundling purchases together with a locked-in borrowing cost for a fixed period. Collateralization is not required since the asset will not be delivered until the end of the borrowing period.

Another example of funding difficulty would be a firm with a sufficiently high credit grade that is under regulatory pressure to reduce the size of its

balance sheet. If it can find another high-credit-grade firm that is not under similar regulatory pressure, it can “rent the balance sheet” of the other firm by entering into a total return swap, although it must expect to pay for the service.

Many of the suppliers of total return swaps to parties wanting to purchase assets they would have difficulty holding simply purchase the asset and hold it until the scheduled delivery or cash settlement. They are being paid to provide a service, as a market maker able to skillfully execute purchases and sales at good prices, an efficient provider of a desired portfolio of assets, a firm having the legal standing to hold assets of a desired country, or a firm with a higher credit standing and lower funding costs or more balance sheet room. However, as the market has evolved, many suppliers are also using this market as an efficient means of borrowing assets in which they want a short spot position. As with a forward, the total return swap provides convenient bundling of the asset borrowing and short sale into a single transaction. For example, a firm wanting to gain on a price decline of a specific bond, either because of a market view or because this offers a hedge against the credit exposure the firm has to the bond issuer, can enter into a total return swap in which it needs to deliver the bond forward (or equivalently cash settle) and then simply does *not* hold a cash bond against this forward obligation.

10.1.8 Asset-Backed Securities

In general, an *asset-backed security* can be viewed as an alternative instrument to total return swaps in taking on exposures to asset classes where the actual management of the exposure is desired to be left to another party. The reasons why this might be desirable could be copied almost verbatim from Section 10.1.7. The use of asset-backed securities rather than total return swaps in a particular situation is largely a matter of how documentation and collateralization of the swap agreement are handled.

When a particular total return exposure is expected to have a fairly broad appeal to a class of investors, it may be desirable to standardize the terms and offer the exposure through a security rather than a swap. This eliminates the need to individually negotiate swap terms since a single document covers the terms of the security, but the trade-off is a loss of flexibility in fitting terms to an individual investor. The use of a security structure forces investors to invest cash up front, a convenient solution to collateralization concerns, particularly when the number of investors is potentially too large to make the negotiation of individual credit coverage attractive. Of course, the disadvantage is that investors must tie up cash in the transaction,

but in return they get a standardized security that can be sold or pledged as collateral. By contrast, it is hard to transfer ownership of a swap position since your counterparty on the swap, which did not place cash up front, may object to the creditworthiness of the new party to which you want to transfer ownership.

The cash nature of the investment protects the party managing the assets from credit concerns. Investors get credit protection by having the assets on which the exposure will be taken placed in some form of trust, thereby immunizing their payoff from default on the part of the party managing the assets. This leads to a potential problem in the flexibility of asset-backed securities in relation to total return swaps. If assets need to be walled off in a trust, then how can a market maker use this as a vehicle for taking a short sale position in the asset, as we showed can be done with a total return swap? The solution is to have a total return swap as an asset placed with the trust and sufficiently collateralize or protect it by third-party insurance.

Asset-backed securities lend themselves to pooling positions in a large number of similar assets, such as mortgages, credit card outstandings, loans to businesses, or bonds. The standardized nature of the documentation is well suited to the sharing of exposure by a large number of investors in a large pool of assets, thereby decreasing the event risk of each investor owning a particular block of assets. However, this is more a matter of convenience than necessity, and virtually any financial position that can be achieved through an asset-backed security can also be achieved through a total return swap.

When mortgages and loans are pooled to create an asset-backed security, it can be structured so as to virtually eliminate credit exposure to the underlying mortgages and loans by investors in the asset-backed security, or it can be structured to have this credit exposure be part of the risk borne by the investors. Risk management aspects of asset-backed securities for which credit exposure has been eliminated will be discussed in Section 12.4.6. Risk management aspects of asset-backed securities that involve credit exposure will be discussed in Section 13.4.3.

Table 10.2 summarizes the difference in risk between the different types of instruments through which forward risk can be taken. Spot risk refers to the underlying asset. Forward risk is always present for the underlying asset, but may or may not also involve forward risk in a currency (in the case where the underlying asset is a currency, the question is whether forward risk in a second currency is involved). Credit risk refers only to credit risk to the counterparty on the instrument, not to any credit risk embedded in the underlying asset. The distinction between the lender and borrower refers to their position relative to the underlying asset.

TABLE 10.2 Comparison of Risks in Forward Transactions

Instrument	Currency Forward Risk			Credit Risk for Lender	Credit Risk for Borrower	Other Risks
	Spot Risk	Forward Risk	Other Risks			
Direct borrowing and lending, including loans, bonds, and deposits	No	No	Yes, but can be mitigated by collateral.	Only if collateral needs to be posted.	Funding liquidity risk for the lender, unless collateralized.	
Repurchase	No	Yes	Small, only to the extent collateral is inadequate.	Small, only to the extent collateral is in excess of borrowing.		
Forward contract	Yes	Yes	Starts at zero, but can build up if not mitigated by collateral.	Starts at zero, but can build up if not mitigated by collateral.		
Futures contract	Yes, except for some interest rate futures	Yes	Credit exposure is to the futures exchange and is small due to cash settlement.	Credit exposure is to the futures exchange and is small due to cash settlement.	May have delivery option. Convexity risk caused by cash settlement. Possible funding liquidity risk if hedging positions that do not have cash settlement.	
FRAs	No	No	Starts at zero, but can build up if not mitigated by collateral.	Starts at zero, but can build up if not mitigated by collateral.		

TABLE 10.2 (Continued)

Instrument	Spot Risk	Currency Forward Risk	Credit Risk for Lender	Credit Risk for Borrower	Other Risks
Interest Rate Swaps	Single-currency swaps	No	No	Starts at zero, but can build up if not mitigated by collateral.	Starts at zero, but can build up if not mitigated by collateral.
	Cross-currency swaps	Typically yes	Yes	Starts at zero, but can build up if not mitigated by collateral.	Starts at zero, but can build up if not mitigated by collateral.
	Total return swaps	Yes	Yes	Starts at zero, but can build up if not mitigated by collateral.	Starts at zero, but can build up if not mitigated by collateral.
	Asset-backed securities	Yes	Yes	Yes, but can be mitigated by collateral and by trust structure.	Typically no, fully collateralized.

10.2 MATHEMATICAL MODELS OF FORWARD RISKS

The most important fact about the mathematical models used to manage forward risk is that they rely on one very simple principle: a flow on a given date owed by a particular entity should be regarded as absolutely equivalent to any other flow of the same quantity on the same date owed by the same entity (the term *flow* is used rather than cash flow, since we want to consider more general cases than cash payments, such as an entity owing an amount of gold or oil).

At first glance, this may look like a tautology, a statement true by definition. And it is close to one, which helps to explain why practitioners agree so widely on the models used to manage forward risk. However, it is not quite a tautology—a reminder that mathematical finance deals with market realities, not theoretical abstractions. When the product a trader is dealing with is actually a complicated bundle of flows on a large number of different dates, it is not immediately clear that breaking the valuation apart into many different pieces, few of which can independently be priced in the market, is the best way to proceed. Indeed, a few decades ago, objections to this practice were still being raised along the lines that it would be very expensive in terms of transaction costs for a trader to actually hedge the instrument in this way. By now, everyone involved has come to appreciate that the principle, far from causing traders to try to aggressively re hedge each piece of a deal, is actually a powerful analytical tool that enables very complex transactions to be managed in a way that permits a maximum amount of netting of risks before resorting to aggressive hedging.

We need to examine where the complexities in this approach lie in order to see what residual risks still need to be managed. Before turning to the hard part, however, let's first take a few moments to appreciate some of the benefits entailed by the simplicity of this approach.

One benefit is the computational simplicity of the method. The actual bundles of forward transactions that trade in the market can have very complex structures. Our fundamental principle says to ignore all these complexities; just calculate the individual flows that have been bundled together, calculate the present value of each flow in isolation from the others, and then sum the present values. It is not necessary to devise special methods that apply to particular cases, a feature that hobbled many of the methods that were used before the fundamental principle was generally adopted.

A second benefit is the generality of the principle. The same computational method can be used for cash flows, commodity flows, bonds, swaps, forwards, futures, risk-free debt, risky debt, and obligations to deliver stock. (Please note carefully that this is not saying that this method

can be used for valuing stocks, since stocks involve unknown future obligations rather than known flows; what is being said is that an obligation to deliver a fixed number of shares of a stock in the future can be translated into an equivalent amount of shares of the stock to be delivered today.) The same computational method can be used to value individual transactions or portfolios of transactions, since each can be reduced to the summation of a set of individual flows, and therefore can also be used for valuing total return swaps or asset-backed securities tied to portfolios of transactions.

A third benefit is that when all transactions are completely reduced to their respective constituent obligations, you are free to describe transactions in whatever manner is most convenient in a given context. When discussing a physical commodity such as gold, it is often convenient just to think in terms of equivalent quantities; for example, you are willing to trade 100 ounces of gold for delivery today for 102 ounces of gold for delivery in one year. When discussing a currency, you might prefer to talk about the interest rate to be paid—say, 6 percent for one year. Although this is just a different way of saying that you are willing to trade \$100 for delivery today for \$106 for delivery in one year, the interest rate view is often easier to understand using economic theory. When doing computations, it is usually best just to think of discount factors to be multiplied by each flow and then summed to get a net present value equation. When checking the reasonableness of a given set of input parameters to the model, it is often most convenient to think in terms of interest rates that apply to distinct forward time periods—the rate that applies to a particular month, week, or day. Improbable inputs can be more easily spotted if you can see that they imply that a rate of 20 percent on one day will occur in between a 7 percent rate on the immediately preceding and following days.

Formulas for translating from discount factors to zero-coupon interest rates, par-coupon interest rates, or forward interest rates and back again are readily available, as should be expected from our general principle. You are probably already familiar with these formulas. If not, consult Hull (2012, Chapter 4).

The **Rates** spreadsheet on the website for this book illustrates the techniques for valuing a portfolio of flows based on a given set of forward rates. The forward rates are translated into equivalent zero coupon rates, par rates, and discount factors. Each set of flows, which might correspond to a forward, a swap, a bond, or any of the other instruments discussed, is broken up into its individual flows. For each individual flow, a discount factor is determined based on interpolation from the discount factors derived from the given forward rate. Each individual flow is then multiplied by its discount factor, and these values are summed over the portfolio.

The interpolation methodology used in this spreadsheet for illustrative purposes is a simple linear interpolation. In practice, more complex interpolation methodologies are often used. Tuckman and Serrat (2012, Chapter 21) presents a good introduction.

The same spreadsheet will be used elsewhere in this chapter to illustrate how to derive a set of forward rates that can match a given set of observed market prices and to demonstrate the calculation of risk statistics for a portfolio of flows.

Having postponed looking at complexities, it's time to face up to the task. Basically, this discussion can be divided into four topics:

1. **Section 10.2.1.** Models are needed to perform interpolation from flows for which market prices are available to other flows.
2. **Section 10.2.2.** Models are needed to extrapolate prices for longer-dated flows.
3. **Section 10.2.3.** In some cases, going from flow prices to bundle prices is not as simple as the general approach. This is because some products involve flows representing a promised delivery that is actually a promise to deliver a future flow (for example, a forward purchase of a bond). Untangling these flows involves some complexities.
4. **Section 10.2.4.** Although the method is designed to handle fixed obligations, it can be applied to a very important class of nonfixed obligations with just a bit of work—flows that will be determined by certain types of indexes. However, this extension must be performed with care; otherwise, a significant source of risk can slip in unidentified.

10.2.1 Pricing Illiquid Flows by Interpolation

As was pointed out at the beginning of this chapter, the large number of days on which future flows can occur makes it almost certain that liquid quotations will be available from the market for only a small portion of possible flows. Creating price quotes for all possible flows will require some theory that enables the inference of prices of illiquid flows from prices of liquid flows. We will present two theories:

1. The interpolation of the price of an illiquid flow from prices of liquid flows that are both earlier and later than it.
2. A stack-and-roll methodology for pricing flows that have longer maturities than any flows with liquid prices.

The mathematics of interpolation is so simple that it can be easy to lose sight of the fact that interpolation is a financial model to the same

degree as more complex options models. It shares the same characteristics of being a methodology for predicting future financial events, requiring well-thought-out assumptions about the financial markets as grounds for choosing one possible methodology over another, and being a source of potential earnings loss to the extent future events diverge from predictions.

When the modeling nature of interpolation is not kept clearly in mind, the choice of interpolation method can be made based on aesthetic criteria, as if it is just a matter of individual taste with no financial consequences. So let us be very specific about financial assumptions and the financial consequences of choices.

Consider the following example, which is typical of the circumstances in which interpolation needs to be employed in pricing forward flows. You need to price a forward flow occurring in $6\frac{1}{2}$ years in a market in which liquid prices can be obtained for 6- and 7-year flows, but nothing in between. Let us suppose you choose to price the $6\frac{1}{2}$ -year flow as the average of the prices of the 6- and 7-year flows. If you put on a hedge that consists of 50 percent of the 6-year flow and 50 percent of the 7-year flow, you will be perfectly hedged in the short run, since at first changes in the daily mark of the $6\frac{1}{2}$ -year flow will just reflect the average of changes in the daily mark of the 6- and 7-year flows. The same would be true of any other interpolation method chosen (for example, 25 percent of the 6-year flow and 75 percent of the 7-year flow) as long as you match the hedge to the chosen interpolation method.

The test of the hedge's effectiveness will come through time. How well will it hold up as flows come closer to maturity, encountering the denser price quotations that exist (in all forward markets) for nearby flows? If, in this example, liquid prices are available for 2-, $1\frac{1}{2}$ -, and 1-year flows, then the hedge will prove effective to the extent that the $1\frac{1}{2}$ -year flow is priced at the average of the 1- and 2-year flows at the time five years from today, when it will be possible to unwind the trade and its hedge at these liquid prices. To the extent the interpolated value differs from the actual value at unwind, an unexpected loss, or gain, will result. Note that the unwind values are determined by the relationship between 1-, $1\frac{1}{2}$ -, and 2-year flow prices five years from now. The current relationship between 1-, $1\frac{1}{2}$ -, and 2-year flow prices cannot be locked into and plays no role other than serving as a historical observation to use in forecasting future relationships.

An interpolation methodology needs to be judged by the stability of the valuations it will lead to. Trading desks develop a feel over time for how stable the valuations produced by particular interpolation techniques are in a particular market. Historical simulation can be used as a quantitative check on these judgments. (Exercise 10.1 takes you through a test of some possible interpolation methods judged by their degree of instability around historical

price quotes.) The potential valuation errors determined by simulation can be controlled through limits and reserves. The most important lessons to be drawn are:

- Interpolation, like any other model, represents a judgment about what is most likely to occur in the future. To the extent the judgment is wrong, unanticipated future losses and gains will result.
- The key event that needs to be projected by an interpolation model is determining the actual relations between prices for flows at a future date when more liquid unwinds are possible.
- Historical relationships between these liquid flows can be used as inputs to and tests of judgments about future relationships. Limits and reserves can be based on measured historical instability.

The preference usually shown for interpolations that produce smooth pricing curves can be explained by two complementary facts: historical relationships between most liquid flows tend to show smooth pricing curves, and economic intuition about future events tends toward long-term trends without a belief that at some specific future date a sharp change in conditions will occur. However, these are only general trends, not rules. If some specific dates may be believed to have forecastable effects, you should expect to see patterns, such as seasonal patterns, reflected in the interpolations. For a discussion of the impact of seasonal patterns on different forwards markets, see Section 10.3.4.

The choice of which variables to interpolate, whether they are discount prices, zero rates, or forward rates, is in one sense arbitrary since we know that each way of representing prices of forward flows is mutually translatable. However, interpolation using one representation may turn out to be more natural than interpolation using a different representation based on the economic motivation supporting the interpolation method chosen.

One approach that would follow naturally from our discussion would be to choose an interpolated value that minimizes a selected smoothness measure for forward rates or zero coupon rates. Methods that are utilized on many trading desks, such as cubic splines, have been justified on formal or informal arguments along these lines. Another approach that is fairly widely used is to interpolate the logarithm of the discount factor. Table 10.3 shows how this works, with the resulting zero coupon rates and forward rates.

As shown in Table 10.3, the impact of this interpolation method is to use a constant forward rate in all subperiods of the period between two already determined discounts. This method is generally favored by traders with backgrounds in the forwards and futures markets who believe that “all you really know is the quoted forward.” So if you have a forward rate

TABLE 10.3 Interpolation Based on the Logarithm of the Discount Factor

Time	Zero Rate	Discount Factor	Forward Rate
T_1	R_1	$e^{R_1 T_1}$	
T_2	R_2	$e^{R_2 T_2}$	$\frac{R_2 T_2 - R_1 T_1}{T_2 - T_1}$
$T_1 + K(T_1 - T_1)$		$e^{(1-K)R_1 T_1 + KR_2 T_1}$	$\frac{[(1-K)R_1 T_1] + KR_2 T_2 - R_1 T_1}{K(T_2 - T_1)}$ $= \frac{KR_2 T_2 - KR_1 T_1}{K(T_2 - T_1)}$ $= \frac{R_2 T_2 - R_1 T_1}{T_2 - T_1}$

agreement that runs from the end of month 9 to the end of month 12 of 7 percent and no other market observations in this vicinity, this method would assign forward rates of 7 percent to the subperiods from the end of month 9 to the end of month 10, the end of month 10 to the end of month 11, and the end of month 11 to the end of month 12.

But what do you do if you have a 7 percent deposit maturing at the end of month 3 and 8 percent FRA from the end of month 3 to the end of month 6, and you are looking to price a FRA from the end of month 3 to the end of month 4? The methodology says use 8 percent, but most practitioners' economic intuition says the rate should be lower than 8 percent, since it seems as if the market is anticipating rising rates over the period. Most traders make some kind of exception when rates are changing this sharply, but an interpolation methodology tied to a smoothness measure has the advantage of building on this approach in a more general setting.

Computationally, it would be convenient if a definitive set of flows was available for which liquid prices could be obtained on the basis of which prices for all other flows could be interpolated. This is rarely true for two reasons:

1. Price liquidity is a matter of degree. Some instruments have prices that are less liquid than others but still show some liquidity. Therefore, these should be given less weight in determining the discount curve, but should not be completely ignored in setting the curve.
2. Prices are often not available for single flows, but are available for bundles of flows—for example, coupon-paying bonds and fixed-for-floating swaps. If enough liquid flow prices are available to interpolate prices for

all but the last of the flows in a bundle, then the common technique of *bootstrapping* (see Hull 2012, Section 4.5) can be used to first price all the flows except the last and then derive the price of the last flow from these prices and the price of the bundle. However, often not enough prices are available to value all but the last flow. For example, many bond markets have a liquid price for a 7- and 10-year bond, but have no liquid prices in between. To derive a value for the flows occurring in the eighth, ninth, and tenth years, it is necessary to combine interpolation and price fitting into a single step.

The **Rates** spreadsheet on the book's website provides a sample discount curve-fitting methodology that is very general in allowing the optimization of a weighted mixture of the accuracy of fitting known liquid prices and determining a forward rate curve that fits closely to an expected smoothness criterion.

The optimization method simultaneously determines all the discount rates needed to match all of the market prices of instruments that can potentially be priced off a single discount curve. All these discount rates are taken as input variables in the optimization. The objective function of the optimization is a combination of two measures. The first is a measure of how closely the derived price comes to the market-quoted price for each instrument, and the second is a measure of how smooth the discount curve is.

The measure of closeness of fit of the derived price to the market quote can take several forms. The spreadsheet uses a very simple measure, a summation of the square of the differences between the derived price and the market quote summed over all instruments. Each is multiplied by a selected weight. Higher weights are assigned to more liquid prices, and lower weights are assigned to less liquid prices. This places a greater premium on coming close to the more reliable prices while still giving some influence to prices that have some degree of reliability. Greater complexity can be introduced, such as placing a higher weight on differences that are outside the bid-ask spread. The most extreme form of this approach would be to introduce constraints that require that the fit be within the bid-ask spread (this is equivalent to placing an extremely high weight on differences outside the bid-ask spread). The desirability of putting such a high weight on the bid-ask spread depends on your opinion of the quotations you are obtaining, how prone they are to error, and whether you really can count on being able to get trades done within the bid-ask range.

The measure of the smoothness of forward rates used in the spreadsheet is also a very simple one: to minimize the squares of second differences of the forward rates. This measures smoothness based on how close the forward rates come to a straight line, since a straight line has second

differences equal to zero. (For example, the sequence 7, 7.5, and 8, which forms a straight line, has first differences of $7.5 - 7 = 0.5$ and $8 - 7.5 = 0.5$, and therefore a second difference of $0.5 - 0.5 = 0$. The sequence of 7, 7.25, and 8, which is not linear, has first differences of 0.25 and 0.75 and therefore a nonzero second difference of 0.5.) Practitioners may use more complex measures of smoothness, such as minimizing second derivatives.

Different weights can be specified for how important the closeness of price fit is relative to the smoothness of the discount curve. This is just one more appearance of the trade-off between basis risk and liquidity risk. The lower the weight put on smoothness and the more even the weight put on fitting each of the instruments, the greater the assurance that the discount curve produced matches exactly the observed market prices of all instruments. This minimizes basis risk, but increases liquidity risk. If it turns out that you really cannot close out one of these positions at the price obtained from the market, you could have significant losses, since the price you used for valuation was based only on the assumed market price, even if this differed a great deal from the price that could be obtained by hedging with more liquid instruments. Conversely, the higher the weight put on smoothness and the more weight put on more liquid instruments, the greater the assurance that you are pricing off hedges that are based on liquid, achievable prices. This minimizes liquidity risk but increases basis risk, since you are now pricing off hedges that can be achieved with combinations of nearby instruments in the market.

The same guidance we have given for testing the financial impact of interpolation rules carries over to testing the financial impact of procedures for extracting a discount curve from a set of liquid prices. Historical simulation should be used to estimate the stability of valuations that will result from a candidate procedure. Figure 10.1 illustrates the degree to which greater smoothness of forward curves can be achieved with the optimization procedure just discussed than with a simple version of the bootstrapping technique. This simple version is used in Hull (2012, Section 4.5) and is used on a number of trading desks; the **Bootstrap** spreadsheet gives details of this comparison. To repeat, the degree of importance of the greater smoothness resulting from the optimization is not to be found in aesthetic pleasure, but should be measured quantitatively in financial impact.

We have been assuming that all the instrument prices can be completely be determined by discount prices. However, some instruments could have option features, such as callable bonds or futures, that have a nonlinear component to their price. This can be handled by subtracting the option component of the price, leaving a pure nonoption portion that can be priced off the discounts. A complexity is that the option component price may depend in part on the discount curve. An iterative process might be needed.




FIGURE 10.1 Comparison of Forward Rates from Bootstrapping and Optimal Fitting

Option components based on a first approximation to discounts can be used to get the inputs to the optimization, which yields a discount curve. This is then used to reprice the option components. These can be used as inputs to a second round of optimization. This cycle can be repeated until the discount curves produced stabilize.

When developing discount factors, it is important to remember that every obligor will have its own set of factors; a promise to deliver a flow on a given date will be worth something different, depending on how reliable the promise is. There even need to be multiple sets of discount factors for promises of the same obligor since some debts are senior to others and will more likely be paid in case of a default condition.

Before this multiplicity of sets of discount factors seems too overwhelming, let's introduce a note of simplification. It is rare that the type of flow owed plays any role in determining the probability of payment. If you can observe a set of discount factors for a firm relative to the discount factors for the assured delivery of one asset, you can infer the discount factors for that firm for delivery of any other asset. However, this rule has exceptions. If the government of Mexico owes you a debt denominated in its own currency, the peso, you would certainly apply a different discount factor than to its promise to pay a debt denominated in another currency. Mexico has control over the supply of its own currency and can create new currency to meet its

payments. It has no such ability in another currency, and although it could create new pesos and exchange them for the currency owed, this might have a severe enough impact on the exchange rate between the currencies to call into question the country's willingness, and even its ability, to do so.

As the procedures for minimizing credit exposure to a counterparty become more complex, involving collateralization, netting, and margin calls, among other techniques, it becomes more difficult to represent the credit exposure in discounting procedures. Any oversimplification should definitely be avoided, such as discounting the flows owed by A to B on a swap at a discount rate appropriate for A's obligations and the flows owed by B to A on the same swap at a discount rate appropriate to B's obligations. This treats the gross amounts owed on the swap as if they were independent of one another, completely ignoring a primary motivation for structuring the transaction as a swap—the netting of obligations.

As credit exposure mitigation techniques grow in sophistication, they demand a parallel sophistication in valuation technology. This consists of initially treating all flows on a transaction to which credit exposure mitigation has been applied as if they were flows certain to be received. The actual credit exposure must then be calculated separately, taking into account the correlation between the net amount owed and the creditworthiness of the obligor. This methodology is more complex than we can tackle at this point in the book. We will return to this topic in Section 14.3.

10.2.2 Pricing Long-Dated Illiquid Flows by Stack and Roll

An issue that arises frequently for market-making firms is the need to provide value to customers by extending liquidity beyond the existing market. This need arises not only for bonds and single-currency swaps and forwards, but also for FX forwards and commodity forwards. A concrete example would be a firm trying to meet customer demand for 40-year swaps in a market that has liquidity only for swaps out to 30 years.

To see the actual profit and loss (P&L) consequences of a methodology for pricing these longer-term flows, we need to consider a well-known trading strategy: the *stack-and-roll hedge*. In our example, a stack-and-roll hedge would call for putting on a 30-year swap in the liquid market as a hedge against a 40-year swap contracted with a customer. Then, at the end of 10 years, the 20-year swap to which the 30-year swap has evolved will be offset and a new 30-year swap in the liquid market will be put on, which will completely offset the original 40-year swap, which is, at this point, a 30-year swap.

This stack-and-roll strategy can be characterized as a quasistatic hedge in that it requires one future rehedged at the end of 10 years. The results

of this re hedge cannot currently be known with certainty, either as to the transaction costs (that is, bid-ask spread) or as to the impact without transaction costs. However, the fact that only a single re hedge is required allows for great simplification in estimating the expected cost of the hedging strategy and its statistical uncertainty. These features recommend using the methodology to quantify the cost and risk of the longer-term position.

To carry out a numerical example, assume that today's 30-year yield is 6.25 percent. Since you are planning to roll at the end of 10 years from a 20- to a 30-year yield, to the extent you expect yield curve shifts to be predominantly parallel, you should enter into a duration-weighted hedge of 1.192 30-year swaps for every 40-year swap you are trying to create. The number 1.192 is the ratio of a 30-year swap duration (13.40) to a 20-year swap duration (11.24), assuming a 6.25 percent annual par swap rate. To estimate the impact of the roll, you should look at the history of the relationship between 20- and 30-year swap rates. If 30-year swap rates tend to be 5 basis points higher on average than 20-year swap rates plus or minus a standard deviation of 7 basis points, and if you want to keep a reserve against a two-standard-deviation adverse move, you could mark the 40-year swap to a rate of $6.25\% + 0.05\% = 6.30\%$ and set up a 14-basis-point reserve. If historical analysis shows that 30-year swap rates minus 105% of 20-year swap rates have a lower standard deviation (say, 5 basis points) than an unweighted spread, because 20-year rates are more volatile than 30-year rates, you could set up a hedge ratio of $1.192/1.05 = 1.135$ and set up a 10-basis-point reserve.

The actual hedging practice on a trading desk might be to initiate a stack and roll, but it would probably be flexible as to the time at which the roll was actually carried out. The roll could take place at the end of 10 years, but the trading desk might, at that time, decide it was more favorable to defer the roll, since the roll could just as well be carried out at other times using equally liquid instruments—for example, roll at the end of 20 years from a 10-year swap into a 20-year swap. The trading desk might also decide, opportunistically, to roll into a less liquid instrument. For example, after two years, the opportunity might arise in which a bid is available for a 38-year swap that would close out the remaining term of the 40-year swap. In this case, the desk would also need to look for a 28-year swap to close out the remaining term of the 30-year swap it was using as a hedge.

Although a trading desk will want to retain flexibility in managing a stack-and-roll strategy once it is entered into, modelers and risk managers can best achieve their aims by assuming a fixed-roll strategy that involves liquid instruments. By considering a strategy that involves liquid instruments, it should be possible to get very reasonable data history that bears on the probable cost of the strategy. If 20- and 30-year swaps have liquid

market quotes available, it may be possible to obtain several years' worth of daily data on the cost of rolling out of a 20-year swap into a 30-year swap. This data can be used not just to decide on an expected roll cost, but also to determine a probability distribution of roll costs. The probability distribution can give reasonable estimates of the uncertainty of results, which can serve as an objective basis for establishing limits and reserves.

The advantages of this method for risk management are:

- Appropriate hedge ratios can be based on historical data since different possible hedge ratios can be judged based on the relative degree of historical uncertainty of roll cost.
- The method makes a clear distinction between the portion of the expected cost of creating a long-term instrument that can be locked into at current market prices versus the portion that requires projections. In this example, the portion that can be locked into is the current 30-year rate, and the portion that requires projection is the spread between the 30- and 20-year rates at the time of the roll (in 10 years).
- This approach gives a solid financial foundation for what is often a loose intuitive argument along the following lines: "The current 20- to 30-year portion of the yield curve is flat to slightly upward sloping, so to price the 40-year swap at the same yield as the current 30-year swap is conservative relative to extrapolating the 20- to 30-year upward slope out to 40 years." This approach makes clear that what matters is not the current 20- to 30-year relationship, but the projected one, which can probably best be estimated based on a longer history of this relationship.
- Estimates of uncertainty for establishing limits and reserves can be based on readily observable historical market data.
- Future liquidity costs, such as the potential payment of the bid-ask spread, are confined to a single point in time.

Exercise 10.2 takes you through some sample calculations using the stack-and-roll methodology.

10.2.3 Flows Representing Promised Deliveries

Let us consider a typical example of a product involving a flow that represents a promised delivery of future flows. A market maker is asked to quote a price for a three-year U.S. Treasury bond to be delivered in seven years (let's assume we are working with zero coupon instruments for the sake of simplicity—the principles for coupon-paying instruments are the same). If the U.S. Treasury were trying to create such a forward, it would be easy. The Treasury would value the forward as a reduction in its need for 10-year

borrowing and an increase in its need for seven-year borrowings, both of which can be valued off the standard U.S. Treasury discount curve. However, a market maker has a lower credit rating and hence higher borrowing costs than the U.S. Treasury has. If the market maker tries to create the forward by buying a 10-year instrument, the price it would need to charge for the forward would be burdened by seven years' worth of the credit spread between the Treasury and the market maker.

To avoid this, the market maker needs to find a way to borrow for seven years at essentially a U.S. Treasury rate. Since it has the 10-year Treasury purchased to put up as collateral against its seven-year borrowing, this should be feasible. However, it is an institutional fact that a liquid market does not exist for borrowing against Treasury collateral at a fixed rate for seven years. It is certainly possible to borrow against Treasury collateral for short periods with great liquidity, and the market maker should feel no fear about the ability to continuously roll over this borrowing. However, this introduces a large variance in the possible funding costs due to uncertainty about the direction short-term repurchase rates will take over a seven-year period.

The way around this impasse is for the market maker to buy a 10-year Treasury, borrow a seven-year Treasury, and sell the seven-year Treasury. The 10-year Treasury is financed for seven years by a series of overnight repurchase agreements (RPs). The borrowing of the seven-year Treasury is financed by a series of overnight RPs. The market maker has succeeded in achieving the same cost of creating the forward that the Treasury would have, except for any net cost between the overnight RP rates at which the longer Treasury is financed and the overnight RP rate at which the borrowing of the shorter Treasury is financed.

In general, these two RP rates should not differ; on any given day, each represents a borrowing rate for the same tenor (overnight) and with the same quality collateral (a U.S. government obligation). However, the RP market is influenced by supply-and-demand factors involving the collateral preferences of the investors. Some of these investors are just looking for an overnight investment without credit risk, so they don't care which U.S. Treasury security they purchase as part of the RP. Other investors, however, are looking to receive a particular U.S. Treasury that they will then sell short—either as part of a strategy to create a particular forward Treasury or because they think this particular Treasury issue is overpriced and they want to take advantage of an anticipated downward price correction. The higher the demand by cash investors to borrow a particular security, the lower the interest rate they will be forced to accept on their cash. When RP rates on a particular Treasury issue decline due to the demand to borrow the issue, the RP for the issue is said to have *gone on special*.

So the market maker in our example will not know in advance what the relative RP rates will be on the shorter security on which it is receiving the RP rate and the longer security on which it is paying the RP rate. To properly value the Treasury forward created by a market maker, it is necessary to make a projection based on past experience with RP rates for similar securities. This source of uncertainty calls for risk controls, which could be a combination of limits to the amount of exposure to the spread between the RP rates and reserves on forward Treasuries, with reserve levels tied to the uncertainty of RP spreads.

Constant-maturity Treasury (CMT) swaps (see Hull 2012, Section 32.4) are popular products with rate resets based on U.S. Treasury yields. They are therefore valued in the same way as U.S. Treasury forwards. Control of the risk for this product focuses on creating long and short cash Treasury positions and managing the risk of the resulting RP spreads.

10.2.4 Indexed Flows

We will now examine how to extend our methods for handling fixed flows to handling nonfixed flows tied to certain types of indexes. Let's start with a simple example. To keep this clear, let's label all the times in our example with specific dates.

Let's say the current date is July 1, 2013. Bank XYZ is due to pay a single flow on July 1, 2015, with the amount of the flow to be determined on July 1, 2014, by the following formula: \$100 million multiplied by the interest rate that Bank XYZ is offering on July 1, 2014, for \$100 million deposits maturing on July 1, 2015. Since this interest rate will not be known for one year, we do not currently know the size of this flow. However, we can determine a completely equivalent set of fixed flows by the following argument and then value the fixed flows by the methodology already discussed.

We write our single flow as the sum of two sets of flows as follows:

	July 1, 2014	July 1, 2015
Set 1	-\$100 million	+\$100 million \times (1 + Index rate)
Set 2	<u>+\$100 million</u>	<u>-\$100 million</u>
Contracted flow	0	+\$100 million \times Index rate

We will argue that the flows in set 1 should be valued at zero. If this is true, then the present value of our contracted flow must be equal to the present value of the second set of flows, which is a set of completely fixed flows.

It can be argued that the present value of the flows in set 1 should be zero because the very meaning of the interest rate that Bank XYZ will be offering on July 1, 2014, for \$100 million deposits maturing on July 1, 2015, is the rate at which customers of XYZ are willing on July 1, 2014, to pay XYZ \$100 million in order to receive a cash flow of \$100 million \times (1 + Rate) on July 1, 2015. So why would we currently value the right to enter into a transaction that will, by definition, be available on that date at anything other than zero?

A second argument can be given for why the present value of the flows in set 1 should be zero. Mathematically, it is equivalent to the argument already given, but it differs in institutional detail and can deepen intuition, so I will provide it.

Let's say we are considering offering a FRA with the following flows:

	July 1, 2015
Set 3	+\$100 million \times Index rate
	-\$100 million \times Fixed rate

At what fixed rate would you be willing to enter into this FRA at an up-front cost of zero (which is equivalent to saying it has a discounted present value of zero)? You should be willing to do this only if the fixed rate is one that you can lock into today at zero cost. The only such rate is the one that makes the following set of flows have a discounted present value of zero; see Hull (2012, Section 4.7) for a detailed example.

	July 1, 2014	July 1, 2015
Set 4	-\$100 million	+\$100 million \times (1 + Fixed rate)

Since both sets 3 and 4 have discounted present values of zero, their sum must also have a discounted present value of zero. The fixed rate is the same in sets 3 and 4 by construction, so the sum is just:

	July 1, 2014	July 1, 2015
	-\$100 million	+\$100 million \times (1 + Index rate)

This is the set of flows we wanted to prove has a discounted present value of zero.

One major caveat exists for this approach: it works only when the timing of the index payment corresponds exactly to the index tenor. If, in our example, the payment based on the one-year index, set on July 1, 2014,

had taken place on July 1, 2014, rather than July 1, 2015, the argument would not have worked in eliminating the index rate from the cash flows and we would have ended up with an additional term consisting of the receipt of $\$100 \text{ million} \times \text{the index rate on July 1, 2014}$, and the payment of $\$100 \text{ million} \times \text{the index rate on July 1, 2015}$. The value of this early receipt of payment depends on what the level of the one-year interest rate will be on July 1, 2014, and the size of the early receipt also depends on what the level of the one-year interest rate will be on July 1, 2014. This nonlinearity gives rise to convexity, which is very similar to an options position in that no static hedge is possible (a dynamic hedge is required), and the value of the position rises with higher rate volatility.

Other situations also lead to convexity:

- **Positions that have up-front cash settlement without discounting, such as futures.** The value of receiving gains up front is dependent on future rate levels. If changes in the value of the future correlate with changes in rate levels, as they certainly will for an interest rate future, the value will be a nonlinear function of rate levels.
- **Positions where a payment is linearly based on the future rate, rather than the future price, of a bond or swap.** The value of payments based on the future price can be determined by discounted cash flows. However, the future rate is a nonlinear function of the future price.

Hull (2012, Section 6.3 and Chapter 29) discusses the issue of convexity adjustments in the valuation of forward risk. Although complete models of convexity adjustments require term structure interest rate options models, Hull offers some reasonable approximation formulas for convexity adjustment in these sections. We will examine a more precise technique for convexity adjustments in Section 12.1.3.

Now that we have found the set of fixed cash flows that are equivalent to an indexed flow, it is important to remember that these fixed flows need to be identified with the same obligor as the indexed flows. Indexed flows are almost exclusively determined for a panel of highly creditworthy banks. For example, LIBOR is determined by a set formula from offering rates of a panel of banks determined by the British Bankers' Association. Panels of firms are used because they minimize the danger of firms manipulating the index. If many contractual rates were tied to the rate at which a certain individual bank was offering to pay for deposits, the bank could set its rate a bit higher if it knew this would impact the amount it owed on a large number of contracts. By using a panel of banks and having rules that throw out high and low offers and average those in between, the impact of any one bank on the index rate is lessened.

So the index flows need to be translated into fixed flows representing the average credit discount of a panel of banks. This can lead to risk in four different directions, all of which need to be properly accounted for:

1. Different panels are used for different currencies within the same location. There are more Japanese banks in the panel that determines yen LIBOR than in the panel that determines dollar LIBOR; therefore, if Japanese banks are perceived to decline in creditworthiness, it will lead to a higher credit spread applied to the fixed flows that yen LIBOR is equivalent to than for the fixed flows that dollar LIBOR is equivalent to.
2. Different panels are used for the same currency within different locations. There are more Japanese banks in the panel that determines the yen Tokyo Interbank Offered Rate (TIBOR) than in the panel that determines the yen LIBOR. Fluctuations in the perceived creditworthiness of Japanese banks lead to fluctuations in yen LIBOR-TIBOR spreads. Firms that have taken the shortcut of valuing all yen index flows the same have suffered significant losses from overlooking this exposure.
3. The panel of banks determining an index has a different credit rating than that of an individual obligor. It is important to discount indexed flows at a different set of discount factors than fixed flows of a specific obligor and make sure that exposure to changes in the relationship between these discount factors is kept under control. During the global banking crisis of 2007–2008, this became particularly important, as credit concerns caused wide gaps to appear between the funding costs of individual banks; see Tuckman and Serrat (2012, Chapter 13) for an analysis of how modeling of interest rate products needed to adjust to these events.
4. There can be differences in the pricing of index flows for different frequencies. For example, if there is an expectation that six-month LIBOR will average 5 basis points more than three-month LIBOR over a five-year period, you would expect a five-year swap against six-month LIBOR to have a 5-basis-point higher fixed rate than a five-year swap against three-month LIBOR quoted at the same time. There is a swap product, *basis swaps*, that trades LIBOR at one frequency against LIBOR at another frequency. Usually, basis swap pricing shows a slightly higher expectation for LIBOR that is reset less often (e.g., six-month LIBOR would be greater than three-month LIBOR). This is both because banks in raising funds prefer to lock in rates for a longer time period, guarding against temporary periods of illiquidity, and because swap investors receiving LIBOR have a slight preference for more frequent resets; it gives them an advantage if a bank in the LIBOR panel has to drop out due to a deteriorated credit outlook and is replaced

on the panel by a bank with lower funding costs. This basis difference is normally quite small, but rose during the banking liquidity crisis of 2007–2008. For details, see Tuckman and Serrat (2012, 449–450).

10.3 FACTORS IMPACTING BORROWING COSTS

When designing stress tests and setting limits for forward risk for a given product, risk managers must understand the economics of the borrowing costs for that product in order to gauge the severity of stresses the borrowing cost can be subject to. Four key characteristics, which differ from product to product, should be distinguished:

1. **Section 10.3.1.** How large and diversified is the borrowing demand for the product?
2. **Section 10.3.2.** To what extent does *cash-and-carry* arbitrage place a lower limit on borrowing costs?
3. **Section 10.3.3.** How variable are the storage costs that impact the cash-and-carry arbitrage?
4. **Section 10.3.4.** To what extent are borrowing costs seasonal?

We also discuss the relationship between borrowing costs and forward prices in Section 10.3.5.

10.3.1 The Nature of Borrowing Demand

A source of borrowing demand that exists for all products comes from traders wanting to borrow in conjunction with short selling. For some products—such as stocks, bonds, and gold—this is the only significant source of borrowing demand. At the other extreme are currencies, where there is strong credit demand by businesses and households to finance purchases and investments. Intermediate cases include most physical commodities, such as oil or wheat, where borrowing demand exists to meet immediate consumption needs.

Products for which borrowing demand comes almost exclusively from short sellers tend to have very low borrowing rates most of the time, since there is little competition for the borrowing. This may not be immediately obvious in market quotes if the quotes are made as forward prices rather than borrowing rates. For example, a one-year gold forward might be quoted at 314.85, a 4.85 percent premium to a \$300 spot price. However, when this is broken apart into a borrowing cost for cash and a borrowing cost for gold, it almost always consists of a relatively high borrowing cost for cash,

say 6 percent, and a relatively low borrowing cost for gold, say 1 percent. As a result, \$300 today is worth $\$300 \times 1.06 = \318 received in one year, and 1 ounce of gold received today is worth $1 \times 1.01 = 1.01$ ounces of gold received in one year, giving a forward price of gold of $\$318/1.01$ ounces = \$314.85 per ounce. Borrowing rates rise as short-selling activity increases. The major risk for short sellers in these products is the *short squeeze* in which borrowing costs are driven sharply upward by a deliberate policy of a government or of holders of the assets seeking to support prices by restricting the supply of available borrowing. The resulting increase in borrowing rates pressures short sellers to abandon their strategy and close out their positions.

Short squeezes are possible for any asset class, but are more difficult to achieve for assets where borrowing demand has a broader base. A government wanting to support the price of its currency may be tempted to tighten the money supply in order to place borrowing cost pressures on the short sellers of the currency, but it will be limited by the fact that these increased borrowing costs will also hurt business firms and consumers who borrow. Even so, a government faced with a run on its currency will still decide on occasion that the desire to pressure short sellers outweighs other considerations, and will either take steps to sharply raise rates or put in place legal measures that discriminate against certain classes of borrowers who are believed to be selling the currency short. An example of the former is the Irish central bank driving short-term rates to 4,000 percent in 1992 in an attempt to teach speculators a lesson (Taleb 1997, 212). An example of the latter is Malaysia in 1997 closing its currency borrowing markets to foreign investors.

The possibility of a short squeeze on borrowing rates acts as a brake on those who want to take a position on an asset declining in value, since they are faced with a risk to which those taking a position on the asset price increasing are not subject. Those wanting to position for price increases in a particular asset have the freedom to borrow any other asset (most probably, but not necessarily, a currency) relative to which they believe it will increase in price and exchange one for the other in the spot market. However, those wanting to position for a price decrease in a particular asset must borrow that particular asset. The consequence of this asymmetry for rate scenarios is that the possible short squeeze means the risk of very high borrowing rates needs to be guarded against.

10.3.2 The Possibility of Cash-and-Carry Arbitrage

When available, the possibility of a *cash-and-carry* arbitrage position acts as a lower limit on borrowing costs. A cash-and-carry arbitrage is one in which an asset is either purchased or borrowed at one date and repaid or sold at a

later date, being held or stored in between the two dates. The example most often cited is a limit on currency-borrowing rates not to go below zero, since if they did, a trader could borrow the currency at a negative interest rate, hold the currency, and then use the currency held to pay back the borrowing, collecting the negative interest payment as guaranteed profit. A more general result says that a lower limit on negative borrowing costs is the storage cost of the asset. This generalization makes it clear that the specific result for a currency rests on the assumption that currency storage costs are zero (by contrast, a physical commodity such as gold has handling and insurance costs of storage that can lead to negative borrowing costs). Although currency storage costs are almost always zero for large amounts (retail depositors may be charged transactions fees), there have been a few historical exceptions. For example, governments wanting to slow the pace at which foreign deposits are driving up the value of their currency have imposed transaction fees or taxes on large deposits, permitting negative borrowing costs.

Cash-and-carry arbitrage is not feasible for all asset classes. Perishable physical commodities, such as live steers or electricity, cannot be stored, so cash-and-carry arbitrage does not place a lower limit on borrowing costs in such markets. Although arbitrage is not available as a limit, some pressure on borrowing costs getting too low will still result from economic incentives for consumers to change the patterns of demand. So if current prices get too high relative to those six months forward, beef consumption will be postponed to the point that the spot price will start to decline relative to the forward price, thereby raising the borrowing rate.

10.3.3 The Variability of Storage Costs

Storage costs on physical commodities tend to be reasonably stable, since they are the cost of physical processes such as handling and transportation. Coupon payments on bonds, a storage benefit, are also stable. However, the storage benefit on stocks, the receipt of dividends, can be quite unstable. A financial setback could lead to a sudden drop in dividends. A merger could lead to a sudden increase in dividends, as in the example at the beginning of this chapter. Changes in tax laws applied to dividends have also resulted in substantial sudden changes in stock-borrowing costs. Note that it does not matter whether the contractual borrowing terms call for the stock borrower to receive the dividends or pass them through to the stock lender. If the borrower receives the dividend, then an increase in dividend will cause the lender to demand a higher borrowing rate. When the borrower must pass the dividend through to the lender, then a borrower who has sold the stock short (which is the only economic rationale for borrowing stock) must pay the increased dividend out of the borrower's own pocket.

10.3.4 The Seasonality of Borrowing Costs

Interpolation methodology for discount factors and the evaluation of the risk of incorrect interpolation must take into account the seasonality of borrowing costs, which can lead to patterns that would be missed by simple interpolation from adjoining prices. To illustrate this with an extreme example, suppose a stock pays a dividend on exactly every July 15. The value of the dividend to be received on July 15, 2014, will be reflected in the borrowing cost to July 15, 2014, but not in the borrowing cost to July 14, 2014. Without knowing this, no conventional methodology for interpolating between borrowing costs to January 1, 2014, and January 1, 2015, will pick up the sharp difference between the borrowing costs to these two dates.

Most borrowing markets do not hinge on such specific scheduling. However, markets for physical commodities such as oil and other energy products and agricultural products often reflect seasonal supply and demand factors such as a stronger demand for heating oil as winter approaches and stronger supply of wheat immediately following harvesting months. The seasonality of borrowing costs for physical commodities is closely tied to the possibility of cash-and-carry arbitrage. Commodities capable of storage that permit cash-and-carry arbitrage will have a smaller seasonal component since the storage of supply can be used as a response to seasonal demand. Perishable commodities that do not permit cash-and-carry arbitrage show a stronger seasonal component, since pricing differentials need to become large enough to start shifting demand. In the extreme case of electricity, which cannot be stored for even very short periods of time, seasonality effects can be seen within a single day, with different forward prices for different times of the day based on differing demand by the time of day.

Borrowing rates for gold, stocks, bonds, and currencies generally show far less of a seasonal effect than borrowing rates for physical commodities, both because of the possibility of storage and because the seasonality of supply and demand is weaker than that for physical commodities. However, some seasonal effects can be observed—most prominently turn-of-the-quarter effects in currency borrowing. This effect is a sharp spike in demand for borrowing currency on the last business day of each quarter and particularly the last business day of each year. A more detailed discussion can be found in Burghardt and Kirshner (1994).

A particularly pronounced seasonal borrowing effect for currencies was experienced throughout 1999 as fears of computer operational problems starting on January 1, 2000—the Y2K problem—caused a large demand for liquidity over the first few weeks of January 2000. Since firms wanted to lock in the currency availability for this period, they were willing to pay much higher borrowing rates for this period than for any period preceding it or succeeding it.

10.3.5 Borrowing Costs and Forward Prices

As emphasized in Section 10.2, every statement made about borrowing costs can be translated into an equivalent statement about forward prices, and vice versa. In market convention, statements about currencies are usually made in terms of borrowing costs, and statements about physical commodities are usually made in terms of forward prices. Since currencies generally have more widespread borrowing demands than physical commodities, as discussed in Section 10.3.1, the borrowing costs for physical commodities will usually be lower than the borrowing costs for a currency. This is usually expressed in forward price terms by saying that the forward price of a physical commodity is generally higher than its spot price—a condition known as *contango*. However, when a strong demand exists for the availability of a particular physical commodity, its borrowing cost may be driven above the borrowing cost of a currency, resulting in forward prices being lower than spot prices—a condition known as *backwardation*. An example of this relationship is shown in Table 10.4. (This terminology has considerable history behind it. In the 1893 Gilbert and Sullivan operetta *Utopia, Limited*, a character is introduced as a financial wizard with the phrase “A Company Promoter this, with special education, Which teaches what Contango means and also Backwardation.”)

A similar situation arises when the borrowing costs are quoted on a net basis in a situation where collateral is being lent to reduce the credit risk of the borrowing. For example, if a security is being borrowed and cash is being lent as collateral, there may be no explicit quote on the borrowing cost of the security. Instead, a net rate is quoted as an interest rate on the cash. If a short squeeze develops on the security, making it expensive to borrow, this will manifest itself as a low (possibly negative) interest rate to be paid

TABLE 10.4 Examples of Contango and Backwardation

Contango Example	Currency	Commodity	Price
Spot	\$100.00	1 unit	\$100.00/unit
1 month	\$100.50	1 unit	\$100.50/unit
2 months	\$101.00	1 unit	\$101.00/unit
3 months	\$101.50	1 unit	\$101.50/unit
Backwardation Example	Currency	Commodity	Price
Spot	\$100.00	1 unit	\$100.00/unit
1 month	\$100.50	1.05 units	\$95.71/unit
2 months	\$101.00	1.10 units	\$91.82/unit
3 months	\$101.50	1.15 units	\$88.26/unit

for the loan of the cash. An identical trade, from an economic viewpoint, is a repurchase agreement. An expensive-to-borrow security will manifest itself through a low to negative rate being paid on the cash side of the transaction.

10.4 RISK MANAGEMENT REPORTING AND LIMITS FOR FORWARD RISK

Risk management reports for forward risk must be more detailed than those for spot risk. Not only do the reports involve an extra dimension of time, but they also involve a dimension of credit quality, since the same flow owed to you on the same day has different risks depending on who owes it to you. We'll examine the time dimension first and then the credit quality dimension.

The basic principle of breaking all forward instrument exposures apart into individual flows has already done a lot of the necessary work for risk management reporting. A complete risk report would just show the amount of net flow exposure for each forward date. The remaining question is what types of date groupings make sense in giving a trading desk and then senior managers a more concise picture of this exposure.

One issue that can lead to some confusion when designing and using risk management reports for forward risk is the overlap in the usage of many close-to-equivalent measures. This starts with disagreement over the simple convention of what is meant by a long position and a short position. In spot markets, *long* clearly means to own an asset, benefit by a rise in the asset price, and lose from a decline in the asset price, while *short* means exactly the opposite in each respect. In forward markets, some practitioners who think about owning a bond use *long* and *short* in the same way—the long position benefits from bond prices rising and therefore from interest rates falling, and the short position benefits from bond prices falling and therefore from interest rates rising.

Other practitioners with backgrounds in instruments such as swaps and FRAs, where no natural concept of an asset being owned is available, use *long* to mean a position benefiting by interest rates rising and *short* to mean a position benefiting by interest rates falling. Often, all you can do is remind yourself which trading desk you're talking to in order to know which way the term is being used, but insist that everyone must agree to use a firmwide convention, no matter how much they hate it, when talking to the chairman of the board.

A similar set of differences in convention is at work when describing the size of a position. Some traders have grown up using the term *value of a basis point* (or equivalently *value of an 01*), whereas others refer to a *5-year equivalent*, *10-year equivalent*, or *duration*. Tuckman and Serrat (2012, Chapters 4 and 5) is highly recommended for a detailed and intuitive

explanation of these concepts. Table 10.5 illustrates this with a numerical example in which we'll consider a position with just two components: a 5-year flow and a 10-year flow.

As shown in Table 10.5, the different position size measures differ only by a constant factor. The five-year equivalent of a position is just the value of a basis point of that position divided by the value of a basis point of a five-year instrument. Any other instrument could be used as a similar common denominator (also known as a *numeraire*). Table 10.5 also shows that the weighted duration is essentially just the value of a basis point divided by minus 1 basis point (−0.01 percent). However, note that duration needs to be weighted by the price value of the position, whereas all the other measures are weighted by the par value of the position, reflecting the definition of duration as the price change per dollar of portfolio value. See Tuckman and Serrat (2012, 130); see also Tuckman and Serrat (2012, 145–147) for a proof that the duration of a cash flow is simply equal to its tenor.

You can check that if the position held was +100 of the five-year flow and −74.536562 of the 10-year flow, using the ratio between values of a basis point, the five-year equivalent, 10-year equivalent, and duration measures for the portfolio would all come out equal to 0. However, the impact of a 100-basis-point increase would not be 0; it would be $-3.613013 + (0.74536562 \times 4.725634) = -0.090688$. So a position that is completely hedged for a 1-basis-point rate move is not completely hedged for a 100-basis-point move. This nonlinearity stems from the fact that the formula for converting interest rates to prices is not a linear formula. Risk exposure to the size of a move in input variables is known as *convexity risk*. This is a risk that does not exist for spot exposures, which are linear, and is a major issue for options exposures; it will be a principal topic of the next chapter.

The convexity of forwards is much less severe than for options, and it is rare for risk managers to focus much attention on it. In addition to not being a very large effect, it is directly tied to hedging longer positions with shorter positions (since the nonlinear effects grow with time to maturity), and risk reporting will already be directed at the degree of maturity mismatch.

Convexity is an important issue for one type of forward risk—credit exposure. Because a credit event, such as the downgrade of a credit rating or, at the extreme, a default event, can cause credit spreads to jump by hundreds or even thousands of basis points, the degree of hedge exposure can be enormous. Reconsider our previous example with the hedge ratio of 100:74.536562, making the position neutral to a 1-basis-point change in the credit spread. In the event of default, there will no longer be any difference between a 5- and 10-year flow—both will just represent claims in a bankruptcy proceeding. If a 30 percent recovery occurs on these claims, the hedged position will show a loss of $70\% \times (-100 + 74.536562) = -17.8244066$.

TABLE 10.5 Sample Computation of Forward Risk Positions

	5-Year Flow	10-Year Flow	Portfolio
Amount	+100	-100	
Current zero coupon rate	6.00%	7.00%	
Current discount factor	$1/\exp(0.06 \times 5)$ = 0.74081822	$1/\exp(0.07 \times 10)$ = 0.49658530	
Current value of positions	74.081822	-49.658530	24.423292
Discount factor given 1-basis-point increase in zero coupon rates	$1/\exp(0.0601 \times 5)$ = 0.74044790	$1/\exp(0.0701 \times 10)$ = 0.49608897	
Value of portfolio given 1-basis-point increase in zero coupon rates	74.044790	-49.608897	24.435943
Impact of 1-basis-point increase in zero coupon rates	74.044790 - 74.081822 = -0.037032	-49.608897 - (-49.658530) = 0.049633	0.01260
5-year equivalents	+100	-100 \times (0.049633/ 0.037032) = 134.0305	-34.0305
10-year equivalents	+100 \times (0.037032/ 0.049633) = 74.536562	-100	-25.390117
Duration	5 years	10 years	
Weighted duration	+74.081822 \times 5 = 370.40911	-49.658530 \times 10 = -496.58530	-126.17619
Discount factor given 100-basis-point increase in zero coupon rates	$1/\exp(0.07 \times 5)$ = 0.70468809	$1/\exp(0.08 \times 10)$ = 0.44932896	
Value of portfolio given 100-basis-point increase in zero coupon rates	70.468809	-44.932896	25.535913
Impact of 100-basis-point increase in zero coupon rates	70.468809 - 74.081822 = -3.613013	-44.932846 - (-49.658530) = 4.725634	1.112621

This is a risk that investors need to be aware of. It explains why investors in bonds issued by firms with high default risk (known as *high-yield debt*, or, less politely, as *junk bonds*) tend to deal directly with prices and avoid reference to interest rates. For a further discussion of the impact of convexity on credit exposure, see Section 13.1.2.2.

Firm-level risk management for forward risk requires decisions about the degree of detail with which exposure to changes in yield curve shape will be represented. Senior management almost certainly needs to be informed of only a few parameters that represent the rate exposure. Many studies have been performed on the historical changes in the shape of many different rate curves, and almost all have shown that about 80 to 90 percent of all changes can be explained by just two parameters, and close to 95 percent of all changes can be explained by just three parameters. Although statistical methods can be employed to determine the best two or three principal components, it makes for better intuitive understanding if parameters can be chosen that convey a concrete meaning. Fortunately, almost all studies of yield curve movement show that intuitively meaningful parameters perform almost as well as parameters selected by statistical means (see, for example, Litterman and Scheinkman 1988). The three parameters that explain most of the change, in order of importance, are:

1. A parallel shift parameter.
2. A parameter to measure the degree of linear tilt of the yield curve.
3. A parameter to measure yield curve twist, the degree to which the middle of the curve changes relative to the two ends of the curve.

The **Rates** spreadsheet illustrates the calculation of the impact of parameter shifts on a portfolio.

The parallel shift parameter certainly represents a nondiversifiable risk in the sense of Section 6.1.1, and a case could be made for considering the linear tilt parameter having an element of nondiversifiable risk as well. It is therefore particularly important that these exposures be highlighted to management.

Nonstatistical limits on yield curve shape exposure also often start with such overall parameters, but it is usually found to be necessary to have more refined limit measures as well. The debate is often between bucket measures based on groupings of forward risks (for example, zero- to one-year forwards, one- to two-year forwards, two- to three-year forwards, and so on) versus bucket measures that break the yield curve exposure down to exposure to yield changes in the most liquid hedging instruments (such as futures contracts out to five years and then 7-, 10-, and 30-year swaps). The primary argument in favor of the latter

approach is that these are the actual hedging instruments most likely to be used; therefore, limits expressed in these terms are immediately operational (a trader knows what action needs to be taken to close a position) and can more easily be judged as to the viability of limit size relative to customer order flow and market liquidity for that instrument. The primary argument against this approach is that the translation of cash flow exposures into liquid hedging instrument equivalents is not completely determined, and very small changes in the choice of algorithm can lead to large changes in how a position is distributed between different instruments. For further discussion of this choice, see Tuckman and Serrat (2012, 158–159).

The decision of which currencies, commodities, and equities should be grouped together rests on very similar considerations for yield curves as for spot risk (refer to the discussion in Chapter 9). Within a grouping, limits are needed by obligor. You would, at a minimum, want to have limits on the government's curve and the interbank rate curve (also known as the swap curve or LIBOR curve), but would probably want to group together rate curves for other obligors, probably by credit rating and possibly by industry and country.

EXERCISES

10.1 Interpolation

For this exercise, make use of the **RateData** spreadsheet. Suppose you are making a market in 16-, 17-, 18-, and 19-year swaps. Liquid swaps are available at 15 and 20 years. Try out some different interpolation methods and test their effectiveness when using them to derive unwind values.

Here are some suggestions:

- There isn't enough data on the spreadsheet to see what the impact of initiating a hedge at one point and unwinding in 10 years would be, so let's make the reasonable assumption that the long-term distribution of rate curve shapes is reasonably stable. So, for example, we'll judge the effectiveness of interpolating the 18-year rate from $40\% \times$ the 15-year rate $+ 60\% \times$ the 20-year rate by looking at the long-term distribution of unwind costs of

an eight-year rate relative to $40\% \times$ the five-year rate $+ 60\% \times$ the 10-year rate.

- Standard deviation can be used as a reasonable summary statistic for the uncertainty of unwind cost, although you should feel free to explore other possible measures such as the 99th percentile.
- To keep the math easier, ignore any compounding effects; that is, treat the par swap rates as if they were zero coupon rates. So the gain from buying an eight-year swap at 6% and selling a five-year swap at 5.70% $+ 60\%$ of a 10-year swap at 6.10 percent is just the following: $(40\% \times 5.70\% + 60\% \times 6.10\%) - 6\% = -0.06\%$.
- You can look at the impact of interpolating with different percentages than those suggested by maturity; for example, consider a 50% five-year, 50% 10-year interpolation for an eight-year swap as well as the standard 40% five-year, 60% 10-year interpolation.
- You can consider the impact of factoring the 30-year rate into the interpolation; this will lead to the use of a 20-year rate in the unwind.
- Explore how much improvement in reducing hedge uncertainty comes about by interpolation rather than just assuming a flat curve by looking at the degree to which uncertainty is reduced by using both the five- and 10-year rates in the unwind rather than just the five-year rate (or just the 10-year rate).

10.2 Stack and Roll

Use the sample stack-and-roll computation in Section 10.2.2 and the rate data history from the **RateData** spreadsheet to calculate two standard deviation reserves for the following products:

- 40-year swap
- 35-year swap
- 33-year swap
- 50-year swap

As in Exercise 10.1, assume that the par swap rates are actually zero coupon rates to keep the math simpler.

10.3 Rates

Use the **Rates** spreadsheet to calculate risk exposure for a portfolio of forward instruments:

1. Begin by creating a discount curve that can be used in subsequent calculations. Enter a set of benchmark instruments and market prices into the **Instruments** worksheet and solve for a discount curve that fits these prices, following the spreadsheet instructions. You might, for example, select a set of U.S. Treasury bonds with one-, two-, three-, four-, five-, seven-, and 10-year maturities. A reasonable set of parameters is to put an equal weighting of 1 on each of your benchmark instruments and to place a weight of 90 percent on fitting prices and 10 percent on the smoothness of the resulting forward curve, but you are encouraged to try different parameters and see their impact on the resulting discount curve.
2. After creating the discount curve, select a portfolio of instruments for which to calculate risk exposure by placing weights on each instrument (you can also add other instruments beyond the benchmark instruments). Look at the resulting risk exposure by forward bucket and summary exposure to forward shifts, tilt shifts, and butterfly shifts, and try to make intuitive sense of them.
3. By trial and error (or by creating an optimization routine with the Solver), find modifications to your portfolio weights that make parallel shift exposure close to zero, but retain roughly the same tilt exposure and butterfly shift exposure as your original portfolio.
4. Follow the same instructions as for part 3, but make tilt exposure close to zero and leave parallel shift exposure and butterfly shift exposure roughly the same as in your original portfolio.

Managing Vanilla Options Risk

Every book should have a hero. The hero of this book is not a person but an equation: the Black-Scholes formula for pricing European-style options. Like every hero, it has its flaws and no shortage of detractors ready to point them out. But with help from some friends, it can recover to play a vital role in integrating all options risk into a unified, manageable framework. This is the theme of this chapter and the next.

Options risk may be subdivided into two categories: the risk of relatively liquid options, termed *plain-vanilla* or *vanilla options*, and the risk of less liquid options, termed *exotic options*. Managing options risk for vanilla options is quite different from managing options risk for exotic options, so we will discuss them in two separate chapters.

Almost without exception, the only relatively liquid options are European-style calls or puts, involving a single exercise date and a simple payoff function equal to the difference between the final price level of an asset and the strike price. As such, vanilla options can be priced using either the Black-Scholes formula or one of its simple variants (see Hull 2012, Section 14.8, Chapter 16, and Sections 17.8 and 25.13). The only notable exception to the rule that all vanilla options are European style is that some American-style options on futures are exchange traded and liquid. However, the early exercise value of such options—the difference between their value and that of the corresponding European option—is quite small (as discussed in Section 12.5.1). So treating all vanilla options as European-style calls and puts is a reasonable first approximation.

To simplify our discussion of European options, we will utilize the following three conventions:

1. All options are treated as options to exchange one asset for another, which enables us to only consider call options. So, for example, we treat an option to put a share of stock at a fixed price of \$50 as being a call option to exchange \$50 for one share of stock. This is a more natural

way of treating foreign exchange (FX) options than the usual approach, since whether an FX option is a call or a put depends on which currency you use as your base.

2. Options prices and strikes will often be expressed as percentages of the current forward price, so a forward price of 100 (meaning 100 percent) will be assumed.
3. All interest rates and costs of carry are set equal to zero. This means that the volatilities quoted are volatilities of the forward, not the spot; the hedges calculated are for the forward, not the spot; and option payments calculated are for delivery at the option expiry date. Although almost all options traded are paid for at contract date rather than expiry, discount curves derived from market prices, as shown in Section 10.2, can always be used to find the current spot price equivalent to a given forward payment.

With these three conventions, we can use the following formula for Black-Scholes values:

$$BS(K, T, \sigma) = N(d_1) - KN(d_2) \quad (11.1)$$

where K = strike as percentage of current forward to time T
 T = time to option expiry in years
 N = cumulative normal distribution
 σ = annualized volatility of the forward
 $d_1 = [\ln(1/K) + 1/2 \sigma^2 T] / \sigma \sqrt{T}$
 $d_2 = d_1 - \sigma \sqrt{T}$

This is similar to Equation 25.5 in Hull (2012, Section 25.13). Technically, we are using a model in which the zero coupon bond price is the *numeraire* (see Hull 2012, Section 27.4).

Stating the equation in terms of the forward price rather than the spot price is important for reasons other than formula simplification. First, it follows the principle stated and justified in Section 6.2 that all forward risk should be disaggregated from options risk. Second, this has the advantage of not assuming constant interest rates; the volatility of interest rates and their correlation with spot price are all imbedded in the volatility of the forward. The historical volatilities of forwards can often be measured directly. If they cannot be measured directly, they can easily be calculated from the spot volatility, interest rate volatilities, and correlations. Hedges with forwards are often the most liquid hedges available. If a spot hedge is used, then the appropriate interest rate hedges should be used as well, since interest rates

and carry costs cannot be assumed to be constant. This combined hedge will be synthetically equivalent to a hedge with a forward.

11.1 OVERVIEW OF OPTIONS RISK MANAGEMENT

Even when we limit our discussion to vanilla options, the vast variety of instruments available makes it unlikely that liquidity of any single instrument will be large. For the options on just a single asset, not only do we face the multiplicity of dates we encountered for forward risk products, but each date also has a multiplicity of possible strikes. Once we take into account that options involve an exchange between pairs of assets, the number of possible contracts expands even more rapidly. For example, if a desk trades 10 different currencies, the number of currency pairs of FX options is $10 \times 9 = 90$. In fact, the degree of liquidity available for option products is significantly smaller than that for spot or forward products.

When options market trading first began and, to a more limited extent, as options markets continue to develop for new assets, initial market-maker hedging strategies were often a choice between acting as a broker (attempting to find a structure for which a simultaneous buyer and seller could be found) or relying on an initial static hedge with the underlying instrument until a roughly matching option position could be found. The broker strategy is very limiting for business growth. The static hedge strategy can only convert call positions into put positions, or vice versa; it cannot reduce the nonlinear nature of the option position. As such, it can be used only by trading desks that are willing to severely limit the size of positions (thereby limiting business growth) or to take very large risks on being right about the maximum or minimum levels to which asset prices will move. Static hedging with limited position size remains a viable strategy for a proprietary desk, but not for a market-making desk.

The development of dynamic hedging strategies was therefore a major breakthrough for the management of options market making. Consider Table 11.1, which extends an example that Hull (2012, Tables 18.1 and 18.4) presents, using Monte Carlo simulation to evaluate the performance of dynamic hedging strategies.

Table 11.1 shows that even a very naive dynamic hedging strategy, the stop-loss strategy, which calls for a 100 percent hedge of a call whenever the forward price is above the strike and a 0 percent hedge whenever the forward price is below the strike, results in a large reduction in the standard deviation of results—76 percent of option cost relative to 130 percent of option cost for a static hedge. However, an increased frequency of rehedging can only improve stop-loss results up to this point. By contrast, the dynamic

TABLE 11.1 Performance of Dynamic Hedging Strategies

Price = \$49, interest rate = 5 percent, dividend rate = 0, forward price = \$50
 Strike = \$50
 Volatility = 20 percent
 Time to maturity = 20 weeks (0.3846 years)
 Drift rate = 13 percent
 Option price = \$240,000 for 100,000 shares

		Performance Measure (Ratio of Standard Deviation to Cost of Option)		
		Delta Hedge		
Frequency of Rehedging	Stop Loss	No Vol. of Vol.	10% Vol. of Vol.	33% Vol. of Vol.
5 weeks	102%	43%	44%	57%
4 weeks	93%	39%	41%	52%
2 weeks	82%	26%	29%	45%
1 week	77%	19%	22%	47%
½ week	76%	14%	18%	43%
¼ week	76%	9%	14%	38%
Limit as frequency goes to 0	76%	0%	11%	40%

With no hedging, the performance measure is 130 percent.

hedging strategy corresponding to the Black-Scholes analysis enables the standard deviation to get as close to zero as one wants by a suitable increase in the frequency of rehedging. You can see why the Black-Scholes approach had such an impact on options risk management.

But almost immediately, this was followed by a backlash, focusing on the unrealistic nature of the Black-Scholes assumptions. Principally, these assumptions and the objections are:

- Trading in the underlying asset can take place continuously. (In fact, a practical limit exists on how frequently trading can occur, which places a lower limit on the standard deviation that can be achieved.)
- No transaction costs are involved when trading in the underlying asset. (In practice, transaction costs place an even tighter limit on the frequency of rehedging.)
- The volatility of the underlying asset is a known constant. (If we make the more realistic assumption that volatility is uncertain, with a

standard deviation around a mean, we get results like those in the last two columns of Table 11.1, placing a lower limit on the standard deviation that can be achieved.)

- The underlying asset follows a Brownian motion with no jumps. (In practice, discontinuous jumps in asset prices can occur, even further limiting the degree to which standard deviation can be lowered.)

Trading desks that have tried pure Black-Scholes hedging strategies for large positions have generally found that unacceptably large risks are incurred. A related example is the *portfolio insurance* strategy. Many equity portfolio managers were using this strategy in the mid-1980s to create desired options positions through dynamic hedging. In October 1987, the global stock market crash caused liquidity to dry up in the underlying stocks, leading to trading discontinuities that resulted in large deviations from planned option payoff profiles.

As a result, vanilla options market makers have generally moved in the direction of a paradigm in which they attempt to match the options positions bought and sold reasonably closely, enabling basis risk to be taken both over time while waiting for offsetting trades to be available and with regard to strike and tenor mismatches. The Black-Scholes model is relied on as an interpolation tool to relate observed market prices to prices needed for the residual risk positions left after offsetting closely related buys and sells. Black-Scholes dynamic hedging is used to hedge these residual risk positions.

Three key tools are needed for managing a vanilla options book using this paradigm:

1. A reporting mechanism must be available to measure the amount of basis risk exposure resulting from mismatches in the strike and tenor of options bought and sold. Although summary measures such as *vega* (exposure to a move in implied volatility levels) and *gamma* (the sensitivity of delta to a change in underlying price level) can be useful, the two-dimensional (strike and tenor) nature of the exposure requires a two-dimensional risk measure to be really effective. This measure is the *price-vol matrix* that depicts portfolio valuation sensitivity to the joint distribution of two variables: underlying asset price and implied volatility. It therefore measures exposure to both jumps in underlying asset price and changes in implied volatility. It also measures simultaneous changes in both. We will examine illustrative examples and discuss the use of price-vol matrices in Section 11.4.
2. Dynamic delta hedging of the portfolio of bought and sold options needs to be performed. Guidance for this process comes from the Black-Scholes formula. The targeted hedge for the portfolio is a simple

summation of the targeted hedges of each individual option position, as determined by Black-Scholes. However, given the reality of transaction costs for executing the delta hedges in the underlying, a set of guidelines about how often to hedge is necessary. It has been shown, both by theory and trader experience, that hedging guidelines based on the distance between the current delta hedge and the target delta hedge are more effective than guidelines tied to the frequency of hedging. The degree of tolerance for deviation from the target delta determines a trade-off between higher transaction costs (for lower tolerances) and higher uncertainty of results (for higher tolerances). Section 11.5 discusses these delta-hedging guidelines in more detail along with related issues such as what implied volatility to use to determine the target hedge.

- Options for which liquid market prices are not available are valued based on interpolation from options that do have liquid market prices available. The interpolation methodology translates prices of liquid options into implied volatilities using the Black-Scholes formula, interpolates these implied volatilities to implied volatilities for less liquid options (interpolation is based on both strike and tenor), and then translates implied volatilities to prices of the less liquid options, again using the Black-Scholes formula. Limits and reserves are needed to control uncertainty in the interpolation process. Section 11.6 gives a detailed account of this interpolation method.

Note how closely bound together the three operative legs of this paradigm are. The Black-Scholes formula serves as the glue that binds them together:

- The price-vol matrix shows how the portfolio valuation will change based on a joint distribution of changes in underlying asset price and implied volatility. However, many (probably most) of the options in the portfolio lack liquid market prices, so their valuation depends on the interpolation step. Furthermore, the calculation of the change in option value for a change of asset price and implied volatility is calculated using the Black-Scholes formula.
- As will be seen in the detailed discussion of the price-vol matrix, all calculations are done under the assumption that exposure to small changes in underlying asset price have been delta hedged with a position in the underlying asset, so the validity of the price-vol matrix depends on the execution of this dynamic delta hedging.
- The need for this approach to options risk management is based on the flat rejection of the key assumptions of the Black-Scholes model: continuous rehedging, no transaction costs, no price jumps, and known and

constant volatility. How, then, can we continue to rely on the Black-Scholes model to calculate the impact of changes in underlying asset price, calculate the target delta hedges, and play a critical role in value interpolation? The answer is that position limits based on the price-vol matrix are being counted on to keep risk exposures low enough that deviations from the Black-Scholes assumptions will not have that large an effect. Small risk exposures mean that the size of required delta hedges will be small enough that transaction costs will not be that significant. Small risk exposures mean that the differences between the Black-Scholes model and the presumably much more complex true model (whatever that may be) are small enough to hold down the errors due to valuing and hedging based on a model that is only an approximation to reality.

It is important to be aware of the degree to which this paradigm depends on the availability of market liquidity for hedging instruments. The paradigm works best when reasonable liquidity in vanilla options is available for at least some combinations of strike and tenor. This enables risks to be hedged by actively pursuing the purchase and sale of options to lower exposures as measured by the price-vol matrix. As we will see in Exercise 11.1, price-vol matrix exposures can be held reasonably flat even if only a small number of strike-tenor combinations provide significant liquidity. The valuation of options with other strike-tenor combinations can be interpolated from the liquid set.

If a particular options market does not have liquidity, the paradigm can still work reasonably well as long as the underlying asset has liquidity. The price-vol matrix now serves primarily as a measure of position imbalance. It can serve as a signal to marketers to encourage customer business at some strike-tenor combinations and discourage it at others. It can be used to place limits on new customer business when this would cause risk to exceed management guidelines. It can be used as input to setting limits and determination of reserves against illiquid concentrations of risk. It can also be used as input to calculations of portfolio risks such as value at risk (VaR) and stress tests. Price interpolation, in the absence of liquid market quotations, becomes primarily a mechanism to enforce the consistency of valuations. Delta hedge calculations continue to serve the function of directing dynamic hedging and ensuring the proper representation of options positions in firmwide reports of spot and forward risk.

It is far more questionable to employ this paradigm in the absence of liquidity in the underlying asset. In this case, it is doubtful that dynamic delta hedging can be carried out in any systematic way, and it probably becomes preferable to analyze positions based primarily on how they will behave under longer-term scenarios, with limits and reserves calculated

from this scenario analysis. An example where this may apply is for options written on hedge fund results where there are restrictions on the ability to buy and sell the underlying, which is an investment in the hedge fund. A specific case to illustrate this point is the option Union Bank of Switzerland (UBS) wrote on Long-Term Capital Management (LTCM) performance (see Section 4.1.5).

How well does this paradigm work? Trading desks that have years of experience using it have generally been satisfied with the results. But this is insider knowledge and may be specific to conditions in particular markets. How can outsiders get comfortable with these assumptions, and how can these assumptions be tested in new options markets to which they might be applied? The best tool available is Monte Carlo simulation, in which all of the Black-Scholes assumptions can be replaced with more realistic assumptions, including limits on hedge frequency, transaction costs, uncertain volatility, nonlognormal changes in the underlying price, and price jumps. In Section 11.3, we examine the results of a typical Monte Carlo simulation to see what it indicates about the feasibility of this risk management paradigm.

11.2 THE PATH DEPENDENCE OF DYNAMIC HEDGING

To understand options pricing, an important distinction must be made between path-independent and path-dependent options. A path-independent option's payout depends only on what the price of some underlying asset will be at one particular point in time and does not depend on the actual path of price evolution between the current date and that future date. All European-style options are path independent. Exotic options are divided between path-independent and path-dependent options. In Chapter 12 on managing exotic options risk, we will see that path-independent options are generally much easier to risk manage than are path-dependent options.

Although, when considered in isolation, European-style options are path independent, once we start to evaluate the impact of dynamic hedging, we find that dynamic hedging makes "every option become path dependent." (This is quoted from Taleb [1997, Chapter 16]. I strongly recommend reading Taleb's Chapter 16 along with this chapter.) This is a direct consequence of the limitations of the Black-Scholes assumptions, since continuous hedging at a known constant volatility would result in a definite value with no variation (hence, you would achieve not just path independence, but independence of the final underlying asset value as well). Sporadic dynamic hedging and stochastic volatility make the realized value of a dynamic hedging strategy dependent on the full price history of the underlying asset. Let's illustrate this with a few examples.

The first example is based on one presented in Taleb (1997, 270). It is an out-of-the-money call on \$100 million par value of a stock with 30 days to expiration that is purchased for \$19,000. If no dynamic hedging is attempted, then the option will expire either out-of-the-money for a total loss of the \$19,000 premium or in-the-money with upside potential. The amount of return will be completely dependent on where the underlying asset price finishes in 30 days. Suppose a trader wanting to reduce the uncertainty of this payoff attempts to dynamically hedge her position. Taleb demonstrates a plausible price path for the underlying asset that results in a loss of \$439,000, not even counting any transaction costs. The **NastyPath** spreadsheet provided on the course website enables you to see the details of this path and experiment with the impact of other possible paths. What is it about the path that leads to a loss that is so large relative to the option's cost? Try to reach your own conclusion. I will provide my answer at the end of Section 11.5.

The second example is drawn from my own experience. In early 1987, I was part of a team at Chase Manhattan that introduced a new product—a term deposit for consumers that would guarantee a return of principal plus a small interest payment, but could make higher interest payments based on a formula tied to the closing price of the Standard & Poor's (S&P) stock index on the maturity date of the deposit. Although the stock market had been showing very good returns in the mid-1980s, stock market participation among smaller investors was still not well developed. Therefore, a product that would be Federal Deposit Insurance Corporation (FDIC) insured, would guarantee against loss, and would provide some upside stock participation quickly attracted a sizable amount of investment.

Our hedging strategy for this product was to invest part of the proceeds in standard deposit products, ensuring the ability to return principal plus guaranteed minimum interest, and use the remainder to fund an S&P index call position. As might be anticipated by those who remember the financial events of 1987, this product suffered an untimely demise in the autumn of that year. After the stock market crash of October 19, consumer interest in possible stock market participation sharply diminished, so new funds stopped coming in. We also experienced severe losses on our hedging of the existing product, and the postmortem we conducted to determine the reason for these losses produced some interesting results.

The equities options markets were at a very early stage of development in 1987, so there was virtually no liquidity for options with tenors beyond a few months. Since our market research had determined that there would be little interest in a deposit product with tenors shorter than a year or two, we had decided to initially rely entirely on a dynamic hedging strategy, using a Black-Scholes-determined delta hedge. We were certainly aware of the

vulnerability of this approach to high volatility, but we had done extensive research on the historical patterns of stock market volatility and concluded that we could price the product at an implied volatility that allowed a margin for error that would result in hedging losses only in extremely rare cases.

Not surprisingly, our postmortem showed significant losses due to our inability to carry out the delta-hedging strategy during the period of October 19 and the following few days. The cash and futures equities markets during that period were highly illiquid in the face of panicky selling, and there were even some short periods in which the markets were closed in an attempt to restore stability to chaotic trading. Illiquid markets in the underlying during large price moves result in gapping losses to options sellers employing dynamic hedging strategies. We were not alone in this vulnerability. In October 1987, a substantial number of asset managers following portfolio insurance strategies in which they attempted to achieve the payoff profiles of an option through delta hedging experienced heavy losses as a result of this gapping.

What was less expected, though, was our finding that a considerable part of our loss would have been experienced even if the markets had not gapped. Our loss was due to higher-than-anticipated volatility. This was despite the fact that when we looked over the tenor of our deposit product the average realized volatility was well within the range we had anticipated in pricing the product. Here's where path dependence comes in. The average realized volatility consisted of very high volatility during a short period when the market was plunging sharply, which was preceded and followed by periods of much lower volatility. However, exposure to volatility depends on the relationship between the price level and strike. The higher-than-average volatility during the period when prices were falling sharply cost us much more than we saved from the lower-than-average volatility during the other periods.

This phenomenon can be easily illustrated with some simple Black-Scholes calculations. Suppose you have written a one-year call option with a strike equal to the current forward price. You intend to delta hedge and expect volatility to average 20% over the year. If you are wrong and volatility averages 30%, your expected losses will be $BS(100\%, 1, 30\%) - BS(100\%, 1, 20\%) = 11.923\% - 7.966\% = 3.957\%$. Suppose one-tenth of a year goes by and the forward price is at the same level as when you wrote the option. Your remaining exposure to volatility averaging 30% is $BS(100\%, 0.9, 30\%) - BS(100\%, 0.9, 20\%) = 11.315\% - 7.558\% = 3.757\%$. So $3.757\%/3.957\% = 94.9\%$ of your volatility exposure comes in the last 90% of the option's life and only 5.1% comes in the first 10% of the option's life (a consequence of the fact that $\sqrt{.9}/\sqrt{1} = .949$). However, if the price at the end of one-tenth of a year has fallen by 30%, the remaining exposure to volatility averaging

30% is $BS(70\%, 0.9, 30\%) - BS(70\%, 0.9, 20\%) = 1.188 - 0.184 = 1.004$. So $(1.004\%/3.957\%) = 25.4\%$ of your volatility exposure comes in the last 90% of the option's life and 74.6% comes in the first 10% of the option's life. A very similar effect will be seen for a large rise in underlying price.

With the benefit of experience, we concluded that we had badly underestimated the risk of the product. First, we had not taken into account the potential losses from pricing gaps. Second, the chances of volatility being very high during a short time period are much larger than the chances of it being very high during a long time period, so we had not properly calculated our vulnerability to a short period of high volatility combined with a large price move. Third, we had not looked at the impact of other market participants pursuing strategies similar to ours, thereby decreasing liquidity by competing with us for hedges in the underlying when we most needed them.

What would have been a more prudent way of managing this risk? We had been considering, but had not implemented, a proposal from a broker in exchange-traded, shorter-term S&P options for a hedge of our longer-term options with these shorter-term options. See Section 11.6.3 for a discussion of the risk characteristics of this hedge.

11.3 A SIMULATION OF DYNAMIC HEDGING

In the immediately preceding section, we established that, under realistic economic assumptions, dynamically hedged options are path dependent. In the section before that, we observed the need for testing how well the paradigm of managing options risk using Black-Scholes theory works. Both sections point toward using Monte Carlo simulation to see what the probability distribution of results can be for dynamically hedging an options portfolio.

Using Monte Carlo simulation for dynamic hedging options is an invaluable tool for understanding how the management of an options trading book works in practice. When new options products or hedging strategies are proposed, traders and risk managers alike will want to look at simulation results to assess potential pitfalls. This is an example of the use of simulation in model testing recommended in Section 8.4.3. Simulation gives the flexibility to take into account the impact on hedging results of real-life constraints such as liquidity constraints on the size of changes in hedges that can be performed in a given time period (or the impact of larger changes on the price at which the hedge can be executed).

Simulation also provides a vital learning tool for people who are unfamiliar with the workings of options markets. Theoretical demonstrations

of the power of dynamic hedging rarely carry the conviction that can be provided by observing hundreds of simulation paths that, despite wild gyrations in underlying prices, produce almost identical hedging results. Nothing short of actually suffering through a losing options strategy can convey the pain of an unsuccessful hedge as will observing the losses pile up on a simulation path.

In the course I teach, on which this book is based, I have always insisted that each student personally program and run simulations of a dynamic hedge. I lack a comparable power of persuasion over readers of this book, but I urge each of you to do as much of Exercise 11.2 as you can. Even if you lack the time to program your own simulation, you should at least do parts 4 and 5 of this exercise using the provided spreadsheets.

What features do we want a Monte Carlo simulation of dynamic hedging to contain?

- The simulation must be over a sufficiently large number of possible price paths to produce stable statistics. Prices for the underlying variable must be sampled at enough points on each path to allow for rehedging.
- Since volatility of the underlying price is not constant, but is a stochastic variable, a random process should drive it. Data to determine reasonable values of volatility can be obtained by looking at historical distributions of realized volatility for separate time periods. A separate volatility should be chosen for each path generated.
- The distribution of the underlying price does not necessarily need to be lognormal. Different mixtures of normal and lognormal processes should be tried.
- Rehedges should be allowed only at periodic intervals, and transaction costs of the hedge should be calculated explicitly. Different rules for determining hedge amounts, as discussed in Section 11.5, should be considered.
- When calculating Black-Scholes deltas for rehedging, you generally do not want to take advantage of knowing what volatility is being used for the path, since this would not be available in making actual hedging decisions. Either you want to use the same implied volatility to calculate rehedges on all paths or you want to use some adaptive rule tying volatility used to the history of price moves on the path up to the time of the hedge.
- A random process of significant price jumps, where no rehedging is permitted until after the jump is completed, can be used as a simulation of periods of illiquidity.
- When simulating a portfolio of options for one particular expiry date, it is usually convenient to assume that all hedges are performed with

a forward with the same expiry to avoid needing to keep track of discounting rates. When simulating options with different expiry dates, some assumptions about discounting rates must be used to arrive at relative prices between forwards.

In effect, we are testing the performance of the Black-Scholes model as a hedging tool by running a Monte Carlo simulation based on a more complex, and presumably more accurate, model of underlying price behavior than Black-Scholes utilizes. Why not just value and hedge options by directly using this more complex and complete model? For two reasons:

1. **Computational complexity.** The speed of the computation of the Black-Scholes model for valuation and the fast and direct computation of the target underlying hedge are enormous advantages in providing timely risk information on portfolios of options that may have many thousands of deals outstanding at any given time. By contrast, more complex models can be orders of magnitude slower when computing valuations and often lack a direct computation of target hedges, requiring multiple runs of the valuation algorithm to determine the appropriate hedge. This advantage can particularly be seen in Monte Carlo testing of hedge effectiveness. At each potential rehedg point, the Black-Scholes target hedge is a simple equation; a more complex model may require full recalibration to compute each hedge (see Section 12.3.2 for a discussion of this point in conjunction with hedging barrier options).
2. **Validity.** We don't necessarily know what the correct model is. For testing hedge performance with Monte Carlo, we can make different runs with alternative candidates for the correct model.

As a first example of a simulation, let's look at a comparison between hedging an option using a pure Black-Scholes hedge and hedging using a combination of Black-Scholes delta hedging and hedging with other options. We may suppose that an option has been sold at a strike for which no liquidity is readily available. We can either utilize a dynamic hedging strategy or buy some options at strikes for which liquidity is available and then utilize a dynamic hedging strategy for the residual risk.

For this example, we will assume that a one-year option has been sold at a strike 5 percent in-the-money and that one-year options are available for purchase at strikes at-the-money and 10 percent in-the-money. For the second case, we will consider purchasing the same notional amount of options as has been sold, but split 50–50 between the at-the-money option and 10 percent in-the-money option. The reason for thinking that this might be a good hedge will be shown in Section 11.4. There we will see that the

TABLE 11.2 Monte Carlo Simulation Comparing Pure Dynamic Delta Hedging with Combined Static Option and Dynamic Delta Hedging

Number of Rebalancing	Standard Deviation of P&L		Standard Deviation of P&L		Transaction Costs	
	Given 0% Standard Deviation of Volatility		Given 33% Standard Deviation of Volatility			
	Unhedged	Two-Sided Hedge	Unhedged	Two-Sided Hedge	Unhedged	Two-Sided Hedge
10	25.7%	6.4%	50.6%	6.3%	1.5%	0.1%
20	19.8%	5.6%	41.5%	6.7%	2.2%	0.2%
50	12.4%	4.6%	40.9%	5.5%	3.5%	0.4%
100	8.5%	3.6%	42.6%	4.9%	5.0%	0.6%
200	6.3%	2.5%	41.6%	4.8%	7.1%	0.9%
300	5.1%	1.9%	39.9%	3.8%	8.5%	1.1%
400	4.3%	1.8%	40.1%	4.1%	9.9%	1.2%
500	3.9%	1.6%	38.4%	3.9%	11.2%	1.4%
600	3.5%	1.4%	35.9%	3.4%	12.0%	1.5%
700	3.3%	1.3%	41.0%	3.5%	13.3%	1.6%
800	3.2%	1.3%	39.2%	3.7%	14.4%	1.7%
900	2.9%	1.5%	40.0%	3.8%	15.0%	1.9%

All results are shown as a percentage of the cost of the option to be hedged.

The option is a one-year call struck 5 percent in-the-money.

The expected volatility is 20 percent, and all hedges are calculated based on a 20 percent implied volatility.

The two-sided hedge has half a call struck at-the-money and half a call struck 10 percent in-the-money.

Transaction costs are based on a bid-ask spread of one-fourth point per \$100.

price-vol matrix for this portfolio (Table 11.9) shows very little sensitivity to changes in either the price level or implied volatility. This does not, by itself, prove that the hedge will work well over the life of the option, since it only shows a snapshot at one point in time. In fact, you will be able to see from Tables 11.10 and 11.11 in Section 11.4 that although this portfolio does continue to show low sensitivity to price on volatility shifts for a substantial time period, this sensitivity increases at some point in its evolution. So we need the Monte Carlo simulation to get a statistical measure of the sensitivity. Table 11.2 shows the results of the simulation.

In the context of the discussion of model risk in Section 8.4, the 50–50 mixture of at-the-money option and 10 percent in-the-money option constitutes the liquid proxy that would be used to represent the 5 percent in-the-money option in standard risk reports, such as VaR and stress tests. The Monte Carlo simulation would be used to generate a probability distribution of how much extra risk there is in holding the 5 percent in-the-money option than there is in holding the liquid proxy. The assumption that the 50–50 mixture will constitute a good hedge all the way to the expiration of the option is a simplifying assumption that makes the Monte Carlo simulation easier. In reality, a trading desk would change this mixture through time, particularly as time to option expiry was close. But while a Monte Carlo simulation that included changes in the mixture would be more realistic, it would also be far more difficult to perform. Changes in the volatility surface would need to be simulated, since changes in the mixture will require purchases and sales of options at future dates; transaction costs for purchases and sales of options would need to be included; behavioral rules for trading decisions would be needed on the trade-off between these transaction costs and the desirability of changing the mixture.

What conclusions can we reach?

- If the standard deviation of volatility is zero, then both the pure dynamic hedging and the mixed-option/dynamic hedging strategies can achieve as low a standard deviation of results as you like by increasing the frequency of rebalancing the dynamic hedge, although the mixed strategy achieves a given level of standard deviation with far fewer rebalancings than the pure strategy. For either strategy, there is a trade-off between higher expected transaction costs with more frequent rebalancing and lower standard deviations of results. (Standard deviations of total results, including transaction costs, don't differ significantly from the standard deviations without transaction costs, which are shown in Table 11.2.) However, the mixed strategy can achieve a desired level of standard deviation at a far lower transaction cost level than the pure strategy. For example, achieving a 3% standard deviation with the pure strategy requires about 900 rebalancings with an associated transaction cost of 15.0%. Achieving a 3% standard deviation with the mixed strategy requires about 150 rebalancings with an associated transaction cost of about 0.8%.
- If the standard deviation of volatility is 33%, then there is a lower bound on how much the standard deviation of results can be decreased. For both the pure and mixed strategies, this lower bound is reached at about 250 rebalancings. The lowest level of standard deviation of

results that can be achieved by the mixed strategy is about one-tenth of what can be achieved by the pure strategy, roughly 4% compared to roughly 40%.

- The inability to reduce the standard deviation of results below a lower bound is due to both the uncertainty of volatility and the use of incorrect volatility inputs in forming hedge ratios. However, the first effect is many times larger than the second. A Monte Carlo run with 33% standard deviation of volatility, but with hedge ratios on each Monte Carlo path based on the actual volatility of that path, results in a lower bound on the standard deviation of results that is only reduced from 40 to 36%

Please note that although we are using standard deviation as a convenient summary statistic to give a rough feel for relative levels of uncertainty, both in this example and others in this book, more detailed analysis would be needed before arriving at any precise conclusions. For example, if a measure was being developed for a risk versus return trade-off as input to a decision on a trading strategy, a more complete set of measures of the probability distribution of returns should be used. The discussion of measures of portfolio risk in Section 7.1.2 gives more of a flavor for these considerations.

These results will not be surprising when we examine the price-vol matrix in Table 11.9 in Section 11.4. From the relative insensitivity of portfolio value to a shift in implied volatility we will see there, you would expect low sensitivity to the standard deviation of volatility. The small size of the portfolio's convexity translates into small changes in the delta when prices move, so transaction costs should be low. A reasonable inference, which is supported by experience with Monte Carlo simulations, is that a trading desk can estimate its vulnerability to uncertain volatility and transaction costs by forecasting how large its price-vol matrix positions are likely to be given the anticipated flows of customer business and the availability of hedges with liquid options. Management can keep these vulnerabilities under control by placing limits on the size of price-vol matrix positions.

It is important to recognize the distinction between the two aspects of dynamic hedging costs—transaction costs that arise from bid-ask spreads and gamma hedging costs from buying high and selling low that would be present even if all trades were at midmarket. Transaction costs are a direct function of the frequency of rehedging, and a trade-off occurs between higher transaction costs and lower variability of profit and loss (P&L) with less frequent rehedging. By contrast, there is no a priori reason to believe

that the level of gamma hedging costs will vary in any systematic way with the frequency of rehedging.

A good way to see this latter point is to look at how P&L is related to the gap between actual hedges held and the theoretical hedge called for by the Black-Scholes formula. The expected value of this P&L under the standard Black-Scholes assumption is given by the formula:

$$\sum_{\substack{\text{small} \\ \text{time} \\ \text{periods}}} (\Delta_{\text{actually held}} - \Delta_{\text{theoretical}}) \times \text{expected price change of underlying forward} \quad (11.2)$$

A full mathematical derivation of this formula can be found in Gupta (1997). I will give an alternative derivation using a simple financial argument. In the presence of the Black-Scholes assumptions, use of the theoretical delta will lead to an expected return of zero, so any holdings above or below the theoretical delta can be regarded as proprietary positions that will lead to the same expected return as an outright position in the underlying forward.

The consequence of this formula for the relationship between gamma hedging costs and the frequency of rehedging is that as rehedging becomes less frequent, it widens the gap between $\Delta_{\text{actually held}}$ and $\Delta_{\text{theoretical}}$. However, unless a correlation between the sign of this gap and the sign of the expected price change in the underlying forward is expected for some reason, the expected value of the incremental P&L should be zero. (Although this formula is strictly correct only in the case where the Black-Scholes assumptions hold, Monte Carlo simulation with stochastic volatility shows similar results.)

Are there cases where we might expect a relationship between the sign of the delta gap and the sign of expected price changes in the underlying forward? Let's consider a case that will cast an interesting light on a long-standing debate among practitioners. The debate is over what options pricing is appropriate for a market in which the underlying process shows *mean reversion*, resulting in a narrower dispersion of future price levels than would be implied by a pure random walk with the short-term volatility of the underlying process. One group argues that delta-hedging costs are completely a function of short-term volatility, so mean reversion is irrelevant to pricing. The opposing group argues that risk-neutral valuation principles should result in the same pricing of options as would be implied by the probability distribution of final prices; compare the discussion here to Rebonato (2004, Sections 4.7 and 4.8).

Some of this dispute reflects a failure to distinguish between the short-term volatility of spot prices and forward prices. If the market is pricing

TABLE 11.3 Impact of Drift and Mean Reversion on Dynamic Hedging Results

	All Paths	Upward Drift	Downward Drift	Mean Reversion
20 rehedges	-0.07%	-0.33%	-0.45%	+0.57%
100 rehedges	+0.01%	-0.06%	-0.10%	+0.20%
1,000 rehedges	-0.01%	0%	0%	-0.02%

the mean reversion process into the forward prices, we should expect to see a lower historical short-term volatility of forward prices than a historical short-term volatility of spot prices. Equivalently, this can be viewed as a correlation between changes in spot prices and changes in the discount rate of the forwards, a pattern that can be seen in the market for seasonal commodities. When seasonal demand is high or seasonal supply is low, spot prices rise, but so does the discount rate, dampening the rise in forward prices. When seasonal demand is low or seasonal supply is high, spot prices fall, but so does the discount rate, dampening the fall in forward prices. Since the option can be delta hedged with the forward, replication costs will be tied to the volatility of the forward, so we should expect implied option volatilities to reflect the impact of mean reversion relative to the volatility of the spot price.

Suppose that a trader believes that the market has not adequately priced in mean reversion, so he expects that forward prices will show mean reversion. In this case, we cannot resolve the controversy between the two differing views on options pricing by an appeal to the difference between short-term volatility of spot and forward prices. Let us look at the results of a Monte Carlo simulation in which we ignore transaction costs and study the impact of rehedges at a fixed number of evenly spaced intervals. We will calculate statistics for the whole sample of paths, but also for three subsamples:

1. The third of paths having the highest finishing forward prices, which we can take as representing upward drift of the forward.
2. The third of paths having the lowest finishing forward prices, which we can take as representing downward drift.
3. The remaining third of the cases, which we can take as representing mean reversion.

Table 11.3 shows the resulting expected values of a delta-hedging strategy for a written (sold) option (for a purchased option, the signs would be reversed).

What conclusions can we draw?

- As you increase the frequency of rehedging, you get the same expected results regardless of drift or mean reversion. This is consistent with the theoretical result that, under the Black-Scholes assumptions, standard deviation of results goes to zero as the frequency of rehedging increases, so the P&L will be the same on every path. It is also consistent with Equation 11.2, since frequent rehedging drives the difference between the $\Delta_{\text{actually held}}$ and $\Delta_{\text{theoretical}}$ terms to zero.
- As you decrease the frequency of rehedging, you increase the losses from a sold option with drift or a purchased option with mean reversion, and you increase the gains from a sold option with mean reversion on a purchased option with drift. All of these results are consistent with Equation 11.2. For example, here's the reasoning for mean reversion on a sold option: It is likely that one period's up move will be followed by the next period's down move, and vice versa. After an up move, the $\Delta_{\text{theoretical}}$ on the sold option will increase, but if no hedge is performed, due to the infrequency of rehedging, this will make the $\Delta_{\text{actually held}} - \Delta_{\text{theoretical}}$ for the next period be negative. Since the expected price change in the next period is negative, the expected P&L is the product of two negatives, and hence positive.
- The consequence of the last point for hedging strategies is that if you anticipate mean reversion, you should try to decrease hedging frequency for a sold option (which also saves transaction costs, but increases the uncertainty of return) and try to increase hedging frequency for a bought option (but this needs to be balanced against the increase in hedging costs and uncertainty of return). This is intuitively correct. As the option seller, you want to hold off on rehedging since you expect the market to rebound; as the option buyer, you want to take advantage of the market move with a re-hedge prior to the expected rebound. Conversely, if you anticipate a drifting market, whether up or down, you should try to decrease hedging frequency for a bought option and increase hedging frequency for a sold option.
- If you cannot anticipate either drift or mean reversion, there is no difference in gamma hedging costs based on the frequency of rehedging, so the decision rests purely on the trade-off between transaction costs and the uncertainty of return.

11.4 RISK REPORTING AND LIMITS

The best tool for managing residual options risk on a trading desk is the *price-vol matrix*, which depicts valuation sensitivity to joint

distributions of two variables: the asset price and implied volatility. The **PriceVolMatrix** spreadsheet on the website for this book calculates a price-vol matrix for a small portfolio of options. See the accompanying documentation for details. We will note just three important points about the computation:

1. All boxes in the matrix represent full valuations using the Black-Scholes model utilizing the shifted volatility level and underlying price level. No approximations are being used in the computation.
2. Each box assumes that an underlying position has been put on to neutralize the initial delta position of the options.
3. Only the initial delta position is neutralized; no delta rehedging is allowed during a price shift. Therefore, the price-vol matrix represents the potential impact of price jumps that cannot be delta hedged.

For those who respond better to visual presentations than to numerical information, the spreadsheet produces two graphical representations of the price-vol matrix:

1. A three-dimensional surface of the P&L consequences of changes in the underlying price and implied volatility.
2. A chart showing changes in valuation, delta, vega, and gamma as price levels change.

The price-vol matrix enables a trading desk manager to see at a glance the *convexity* (the nonlinear impact of large price changes), *vega* (sensitivity to a small change in implied volatility), nonlinearities in vega, and interactions between convexity and vega. The price-vol matrix can pick up discontinuities caused by strikes in a portfolio clustering around certain levels. In order for the price-vol matrix to highlight nonlinear effects, it is best to assume that any linear delta position has already been hedged. To the extent that the trading book chooses not to hedge delta risk, the resulting underlying position should be reported separately and be subject to limits separate from those on options positioning for the reasons given in Section 6.2 concerning the need for clear separation of linear and nonlinear risks.

Traders have recently shown greater focus on the sensitivity of vega to changes in implied volatility and the sensitivity of vega to changes in spot. A sign of this increased focus is that these sensitivities have acquired their own mock Greek names, *vomma*, also known as *wisoo*, and *vanna*, also known as *DdelV*, respectively. Note that the price-vol matrix measures changes in P&L impact due to both *vomma* and *vanna*. Also note that the convexity measure goes well beyond a simple P&L impact of *gamma*, which is just the

second derivative of price changes, and hence determines the second-order term in the Taylor expansion of option price in terms of underlying price. Since the matrix is filled in by a full revaluation of the Black-Scholes model for each box, the impact of as many terms in the Taylor series as desired can be picked up by a sufficient refinement of the underlying price grid.

The price-vol matrix is a valuable tool both in the daily P&L reconciliation needed to control model risk (compare with Section 8.2.7.1) and in P&L approximations used in VaR and stress test calculations (compare with Section 7.1.1.2). When the price-vol matrix is used for making P&L approximations, it is often referred to as a *heat map*. For P&L reconciliation, it allows a quick first-cut calculation of P&L change from the close of one business day to the close of the next business day due to the combined change in underlying price and overall volatility level if no delta hedging had been performed during the day. It can then be supplemented by more detailed calculations of P&L changes due to changes in the shape of the volatility surface and due to delta hedging performed during the day. P&L due to changes in volatility shape can be calculated from a matrix that breaks down vega exposure by strike and by time to expiry (the **PriceVolMatrix** spreadsheet contains a sample computation of a vega exposure matrix).

Another valuable tool for P&L approximation is *dollar gamma*. Dollar gamma is calculated as one-half the gamma multiplied by the square of the current price level. Its use in P&L approximation is that when you multiply a portfolio's dollar gamma by the difference between the volatility at which positions have been marked and the actual price move for the day, you get a good first estimate of a delta-hedged portfolio's P&L for the day. A good explanation of dollar gamma and sample calculations can be found in Allen, Einchcomb, and Granger (2006, Section 4.1).

We will now use the price-vol matrix to examine some representative option positions as a way to learn about both risk characteristics of the positions and the analytic power of the price-vol matrix:

- **Short a call option.** This is the simplest possible options portfolio. We are short one unit of a one-year call struck at-the-money. Table 11.4 shows the price-vol matrix. Naturally, vega and gamma are both negative, and vega remains negative at all price levels. Negative vega is largest at-the-money and declines as prices rise and fall, reflecting the decline in the time value of an option as it goes into or out of the money. The negative gamma is reflected in large losses from either up or down price jumps at the current volatility.
- **Call spread.** We are short one unit of a one-year call option struck at the money and long 1.06 units of a one-year call option struck at 110 percent of the forward price. Table 11.5 shows the price-vol matrix. The

Discount	5.00%																								
Spacing	Volume																								
Price	call	-1																							
	100																								
Volatility	Strike	100																							
	Time	1																							
Portfolio	Implied vol	20.0%																							
	BS price	-7.58%																							
	Delta	-54.0%																							
	Vega	-0.38%																							
	Gamma	-2.0%																							
	Theta	0.015%																							
Spot-Vol Matrix																									
Price	-8%	-6%	-4%	-2%	0%	2%	4%	6%	8%																
	-25	-5.40%	-5.51%	-5.64%	-5.81%	-6.01%	-6.24%	-6.50%	-6.77%																
	-20	-3.10%	-3.30%	-3.55%	-3.82%	-4.12%	-4.45%	-4.80%	-5.17%																
	-15	-1.10%	-1.43%	-1.79%	-2.18%	-2.59%	-3.02%	-3.46%	-3.91%																
	-10	0.90%	0.00%	-0.48%	-0.97%	-1.48%	-1.99%	-2.51%	-3.03%																
	-5	2.01%	1.45%	0.89%	0.33%	-0.24%	-1.38%	-1.96%	-2.53%																
	0	2.42%	1.81%	1.21%	0.60%	0.00%	-1.21%	-1.81%	-2.41%																
	5	2.14%	1.56%	0.97%	0.37%	-0.23%	-1.43%	-2.04%	-2.65%																
	10	1.27%	0.76%	0.23%	-0.32%	-0.88%	-2.03%	-2.62%	-3.20%																
	15	-0.08%	-0.49%	-0.93%	-1.40%	-1.90%	-2.95%	-3.50%	-4.05%																
	20	-1.77%	-2.07%	-2.41%	-2.80%	-3.22%	-4.14%	-4.64%	-5.15%																
	25	-3.67%	-3.87%	-4.13%	-4.43%	-4.78%	-5.57%	-6.00%	-6.46%																

TABLE 11.4 Price-Vol Matrix for Being Short a Call Option

Discount	5.00%					
Spacing	Volume	-1	1.06			
Price	call/put	call	call			
	Price	100	100			
Volatility	Strike	100	110			
Portfolio	Time	1	1			
	Implied	20.0%	20.0%			
	vol					
	BS price	-7.58%	4.33%			
	Delta	-54.0%	37.4%			
	Vega	-0.38%	0.38%			
	Gamma	-2.0%	2.0%			
	Theta	0.015%	-0.014%			
Spot-Vol Matrix						
				Implied Volatilities		
				0%		
				2%		
				4%		
				6%		
				8%		
				Vega		
				Convexity		
Price	-25	-8%	-4%	0%	0.00%	
	-20	-0.74%	-0.87%	-0.95%	-1.04%	0.00%
	-15	-0.10%	-0.33%	-0.46%	-0.58%	0.00%
	-10	0.36%	0.03%	-0.12%	-0.26%	0.00%
	-5	0.56%	0.19%	0.04%	-0.08%	0.00%
	0	0.46%	0.17%	0.07%	-0.01%	0.00%
	5	0.14%	0.03%	0.01%	0.00%	0.00%
	10	-0.27%	-0.13%	-0.06%	0.01%	0.00%
	15	-0.61%	-0.41%	-0.08%	0.06%	0.00%
	20	-0.76%	-0.48%	0.00%	0.21%	0.00%
	25	-0.65%	-0.33%	0.22%	0.46%	0.00%
		-0.28%	0.02%	0.59%	0.85%	0.00%
			0.31%	1.09%	1.53%	0.00%
				1.31%	1.73%	0.00%
				0.89%	1.27%	0.00%
				0.89%	1.27%	0.00%
				0.57%	0.88%	0.00%
				0.32%	0.55%	0.00%
				0.15%	0.28%	0.00%
				0.01%	0.05%	0.00%
				-0.12%	-0.03%	0.00%
				-0.27%	-0.06%	0.00%
				-0.49%	-0.07%	0.00%
				-0.81%	-0.06%	0.00%
				-1.23%	-0.05%	0.00%
				-1.14%	-0.05%	0.00%
				-1.32%	-0.06%	0.00%
				-0.91%	-0.07%	0.00%
				-0.34%	-0.06%	0.00%
				-0.15%	-0.03%	0.00%
				0.03%	0.00%	0.00%
				0.21%	0.03%	0.00%
				0.44%	0.07%	0.00%
				0.73%	0.10%	0.00%
				1.09%	0.12%	0.00%
				1.53%	0.12%	0.00%

TABLE 11.5 Price-Vol Matrix for a Call Spread

TABLE 11.6 Center Boxes of Price-Vol Matrix

Price	Implied Volatility	-2%	0%	2%
-5			-0.01%	
0		0.01%	0.00%	0.00%
5			0.01%	

1.06 units have been deliberately selected to create a portfolio with zero vega, gamma, and theta. However, as the price-vol matrix shows, this is not the same as saying there is no options risk in the portfolio.

Focus on the center five boxes in the price-vol matrix of Table 11.5, representing the current price and implied volatility, as well as one shift up and down in price and implied volatility, as shown in Table 11.6.

You can see that this is consistent with vega and gamma being zero, since vega and gamma measure the sensitivity to small changes in volatility and price. However, as you widen your view to the whole matrix, you see both convexity and volatility exposure.

The convexity exposure is to a loss on downward price jumps for which the impact of the sold at-the-money option will outweigh the impact of the purchased option at a higher strike. The convexity impact of upward price jumps is a gain, since the effect of the purchased higher-strike option will outweigh the effect of the sold at-the-money option.

As prices rise, vega will be positive, reflecting the greater impact of the purchased higher-strike option. As prices fall, vega will be negative, reflecting the greater impact of the sold lower-strike options.

Option positions that display these characteristics—acting like a bought option to some price levels, with positive vega and gains from convexity, and acting like a sold option at other price levels, with negative vega and losses from convexity—are known as *risk reversals*, since the direction of risk exposure reverses itself with changes in price level (for further discussion of risk reversals, see Taleb [1997, 135, 275–276]).

Here are two stories that illustrate some of the characteristics of risk reversals. The first comes from the Japanese equity derivatives market in the mid-1990s. Many Japanese banks were selling warrants on their stock that had the price-vol profile of a risk reversal. The warrant buyer would have a positive vega and convexity at the stock price levels then prevailing, but would switch to a negative vega and convexity if stock prices were to fall significantly. Rumors in the market indicate that some trading desks purchased these warrants to provide a hedge against the negative vega and convexity exposure they had from other positions in Japanese equity derivatives, but

did not adequately plan for what would happen if stock prices plummeted, causing the now negative vega and convexity on the warrant to exacerbate the overall negative vega and convexity of the desk. When Japanese stock prices did experience a sharp decline in 1996, it was accompanied by a rise in implied volatility and a decline in the liquidity of underlying stock positions, so negative vega and convexity positions resulted in large trading losses. Some reports indicate that this was one of the events that contributed to the large losses at UBS (refer to the discussion in Section 4.1.5).

The second story goes back further in time to the early days of options trading. The business executive of a newly formed options business, for which I was in charge of analytics, came to me with a situation that was disturbing him. A recent series of large moves had occurred in this particular market, with large decreases in underlying prices and increases in implied volatility followed by large increases in underlying prices and decreases in implied volatility. The net effect was that prices and implied volatilities had pretty much finished up where they had started. Although the market had retained good trading liquidity throughout, the implied volatility moves were substantial enough to trigger material P&L swings. What was disturbing to the business head was that the trading book had been a loser in both the increase and decrease in implied volatility. The time period that was involved had been short enough that no significant change in the options position had taken place. So how could this pattern be explained?

This trading desk did not yet have a regular price-vol matrix, but my team was able to put one together, which quickly revealed a risk reversal pattern for the portfolio. At the price level that prevailed at the beginning of the period, the portfolio's vega was negative, leading to losses from rising implied volatilities. At the level to which prices then fell, the portfolio's vega was positive, leading to losses from falling implied volatilities. So far, so good. But underlying prices and implied volatilities ended where they began. In an unchanged portfolio, wouldn't the Black-Scholes valuation yield the same option prices at the end of the period as at the beginning of the period given that not enough time had elapsed to make a significant difference? It would be a good exercise to think this through yourself before seeing my answer.

The key to understanding what happened is that the portfolio was not really unchanged since delta hedging had gone on throughout the period. Since the markets had retained liquidity throughout, this delta hedging had been smooth and no gains or losses due to price jumps had occurred. If price jumps had occurred rather than smooth delta hedging, then the portfolio would have come back to its original value.

If this is not clear, follow the example in Table 11.7, which corresponds to being short one unit of a one-year at-the-money call and long one unit of

TABLE 11.7 P&L Consequences of a Cycle in Prices and Volatilities

Moves	With Price Jumps	With Smooth Delta Hedging
Volatilities up 8% (0, 0%) → (0, 8%)	-1.06%	-1.06%
Prices down 25% (0, 8%) → (-25, 8%)	+0.84% [-0.22% - (-1.06%)]	0
Volatilities down 8% (-25, 8%) → (-25, 0%)	-0.83% [-1.05% - (-0.22%)]	-0.83%
Prices back up to original level (-25, 0%) → (0, 0%)	<u>+1.05%</u>	<u>0</u>
Total	0	-1.89%

a one-year call at 80 percent of the current price. Assume that the following four moves take place in sequence: volatilities up 8 percent, prices down 25 percent, volatilities down 8 percent, and prices back up to the original level. Table 11.7 shows the P&L consequences, contrasting a case with price jumps and a case with smooth delta hedging. The computations for Table 11.7 can be found in the **PriceVolMatrixCycle** spreadsheet.

This is the most extreme case in which implied volatility moves completely precede price moves. When implied volatility and price moves are mixed together, the effect is attenuated but not lost. Altogether, this constitutes another example of the maxim that delta hedging makes all options path dependent.

- Calendar spread.** We are short one unit of a one-year call option struck at-the-money and long one unit of a six-month call option struck at-the-money. The price-vol matrix in Table 11.8 shows positive P&L from price jumps but negative P&L from an increase in implied volatility. This is also reflected in the positive gamma and negative vega measures for the portfolio. Shorter-term options generally have a greater impact on sensitivity to price jumps than longer-term options of the same size, but longer-term options generally have greater exposure to implied volatility than shorter-term options of the same size.
- Reduced risk portfolio.** We are short one unit of a one-year call option struck at 105 percent of the forward price and long 0.525 units of a one-year call option struck at-the-money and 0.5 units of a one-year call option struck at 110 percent of the forward price. The price-vol matrix is shown in Table 11.9. These weights have been deliberately selected to make gamma and vega zero. However, unlike the call spread

Discount	5.00%																					
Spacing	Volume	-1	1																			
Price	call/put	call	call																			
	5	100	100																			
Volatility	Strike	100	100																			
	2%	1	0.5																			
Portfolio	Implied vol	20.0%	20.0%																			
	BS price	-7.58%	5.50%																			
	Delta	-54.0%	52.8%																			
	Vega	-0.38%	0.27%																			
	Gamma	-2.0%	2.8%																			
	Theta	0.015%	-0.021%																			
Spot-Vol Matrix																						
Price		-8%	-6%	-4%	0%	2%	4%	6%	8%	Vega										Vega		
	-25	2.05%	1.99%	1.90%	1.79%	1.66%	1.51%	1.36%	1.03%		1.03%	1.19%	1.19%	1.36%	1.36%	1.03%	1.03%	1.03%	1.03%	1.03%	1.03%	0.00%
	-20	1.90%	1.77%	1.62%	1.46%	1.29%	1.11%	0.92%	0.74%		0.74%	0.74%	0.92%	0.92%	0.92%	0.74%	0.74%	0.74%	0.74%	0.74%	0.74%	0.00%
	-15	1.59%	1.41%	1.22%	1.03%	0.84%	0.65%	0.46%	0.27%		0.27%	0.27%	0.46%	0.46%	0.46%	0.27%	0.27%	0.27%	0.27%	0.27%	0.27%	0.00%
	-10	1.17%	0.98%	0.78%	0.60%	0.41%	0.23%	0.05%	-0.13%		-0.13%	-0.13%	0.05%	0.05%	0.05%	-0.13%	-0.13%	-0.13%	-0.13%	-0.13%	-0.13%	0.00%
	-5	0.80%	0.62%	0.45%	0.28%	0.11%	-0.06%	-0.23%	-0.39%		-0.39%	-0.39%	-0.06%	-0.06%	-0.06%	-0.39%	-0.39%	-0.39%	-0.39%	-0.39%	-0.39%	0.00%
	0	0.66%	0.50%	0.33%	0.17%	0.00%	-0.16%	-0.33%	-0.49%		-0.49%	-0.49%	-0.16%	-0.16%	-0.16%	-0.49%	-0.49%	-0.49%	-0.49%	-0.49%	-0.49%	0.00%
	5	0.83%	0.64%	0.46%	0.28%	0.10%	-0.07%	-0.25%	-0.42%		-0.42%	-0.42%	-0.07%	-0.07%	-0.07%	-0.42%	-0.42%	-0.42%	-0.42%	-0.42%	-0.42%	0.00%
	10	1.20%	0.99%	0.78%	0.58%	0.38%	0.18%	-0.01%	-0.21%		-0.21%	-0.21%	0.18%	0.18%	0.18%	-0.21%	-0.21%	-0.21%	-0.21%	-0.21%	-0.21%	0.00%
	15	1.65%	1.43%	1.20%	0.98%	0.76%	0.54%	0.33%	0.12%		0.12%	0.12%	0.54%	0.54%	0.54%	0.12%	0.12%	0.12%	0.12%	0.12%	0.12%	0.00%
	20	2.05%	1.85%	1.64%	1.41%	1.19%	0.96%	0.73%	0.51%		0.51%	0.51%	0.96%	0.96%	0.96%	0.51%	0.51%	0.51%	0.51%	0.51%	0.51%	0.00%
	25	2.37%	2.21%	2.03%	1.82%	1.61%	1.38%	1.15%	0.92%		0.92%	0.92%	1.38%	1.38%	1.38%	0.92%	0.92%	0.92%	0.92%	0.92%	0.92%	0.01%

TABLE 11.8 Price-Vol Matrix for a Calendar Spread

case, the zero gamma and vega are reflected throughout the price-vol matrix by low exposures at all combinations of price jump and volatility shift. This demonstrates the ability to achieve greater risk reduction by using positions that are symmetrical in strike price.

Tables 11.10 and 11.11 show how this position evolves through time. We can see that at the end of 0.5 years (Table 11.10), there is still not much risk exposure, but at the end of 0.9 years, with only 0.1 year left until option expiration (Table 11.11), there is some convexity, with gains if prices jump upward and losses if prices jump downward. This shows that even a hedge of options against options that works very well at first cannot be maintained as a purely static hedge. We have already explored the implications of this for options risk management using Monte Carlo simulation in Section 11.3.

The price-vol matrix has the great advantage of looking at precise sensitivity to many different values of two variables, but this carries the disadvantage of only being able to consider two variables. This has two consequences: the choice of which two variables to look at is an important one, and the price-vol matrix needs to be supplemented with risk measures that go beyond these two variables.

The selection of the best variables to use in the price-vol matrix can be based on economic insight or on statistical techniques, such as principal component analysis. On the side of asset prices, one question is whether to assume a parallel shift in forward prices. This is equivalent to assuming zero correlation between changes in the underlying asset price and changes in discount curves. Another question is whether to assume constant spreads between different variants of the asset, such as different grades for a physical commodity and different individual stocks relative to a stock market index. For volatilities, the question is whether to assume parallel changes in the volatility surface or whether to assume a statistical relationship based on historical experience.

Looking more closely at the issue of whether to assume a parallel shift in the volatility surface, let's break this down into a time-to-expiry component and a strike component. With regard to time to expiry, the first principal component of changes in volatility surfaces has less tendency to be flat than the first principal component of changes in interest rate curves. Longer-term volatilities often tend to move substantially less than shorter-term ones. Although a time-differentiated shift conveys less immediate intuitive meaning in discussions with senior management than a flat 1 percent shift, the increase in likelihood may outweigh the communications disadvantage. A possible compromise that is reasonably easy to express and often reasonably close to historical experience is a proportional rather than an absolute shift. So if one-year volatilities are currently 20 percent and five-year volatilities

Discount	5.00%	-1	0.525	0.5	
Spacing	Volume	call	call	call	
Price	call/put	100	100	100	
	Price	105	100	110	
Volatility	Strike	0.5	0.5	0.5	
	Time	20.0%	20.0%	20.0%	
Portfolio	Implied vol	-3.53%	2.89%	1.08%	
	BS price	-39.2%	27.7%	13.7%	
	Delta	-0.27%	0.14%	0.12%	
	Vega	-2.7%	1.5%	1.2%	
	Gamma	0.020%	-0.011%	-0.009%	
	Theta				
Spot-Vol Matrix	Implied Volatilities				
Price	-8%	-6%	-4%	-2%	0%
-25	0.10%	0.10%	0.11%	0.11%	0.12%
-20	0.00%	0.01%	0.02%	0.03%	0.05%
-15	-0.07%	-0.05%	-0.03%	-0.01%	0.00%
-10	-0.08%	-0.05%	-0.04%	-0.02%	-0.01%
-5	-0.02%	-0.01%	-0.01%	-0.01%	-0.01%
0	0.07%	0.04%	0.03%	0.01%	0.00%
5	0.10%	0.06%	0.03%	0.01%	-0.01%
10	0.04%	0.02%	-0.01%	-0.02%	-0.04%
15	-0.06%	-0.07%	-0.07%	-0.08%	-0.08%
20	-0.18%	-0.16%	-0.15%	-0.14%	-0.14%
25	-0.26%	-0.23%	-0.21%	-0.20%	-0.19%
	2%	4%	6%	8%	Vega
	0.13%	0.13%	0.14%	0.15%	0.00%
	0.06%	0.07%	0.08%	0.09%	0.00%
	0.02%	0.03%	0.04%	0.05%	0.01%
	0.00%	0.01%	0.01%	0.02%	0.00%
	-0.01%	-0.01%	0.00%	0.00%	0.00%
	-0.01%	-0.02%	-0.02%	-0.02%	0.00%
	-0.02%	-0.03%	-0.04%	-0.04%	-0.01%
	-0.05%	-0.06%	-0.07%	-0.07%	-0.01%
	-0.09%	-0.09%	-0.10%	-0.10%	0.00%
	-0.13%	-0.13%	-0.13%	-0.13%	0.00%
	-0.18%	-0.17%	-0.16%	-0.16%	0.01%
	Vega		Convexity		
	0.00%	0.00%	0.00%	0.00%	0.00%

TABLE 11.10 Price-Vol Matrix for the Reduced Risk Portfolio of Table 11.9 After 0.5 Years Have Elapsed

Discount	5.00%		-1		0.525		0.5		Implied Volatilities								Vega		
	Spacing	Volume	call	put	call	call	call	call	-8%	-6%	-4%	-2%	0%	2%	4%	6%	8%	Vega	Convexity
Price	5	Price	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	0.00%	0.00%
Volatility	2%	Strike	105	100	100	110	110	110	105	100	100	100	100	100	100	100	100	0.01%	0.00%
Portfolio	2%	Time	0.1	0.1	0.1	0.1	0.1	0.1	0.037%	0.037%	-0.026%	-0.008%	-0.008%	0.037%	0.037%	0.037%	-0.01%	0.00%	
	0.60%	Implied vol	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	-0.81%	1.32%	26.9%	3.5%	3.5%	3.5%	3.5%	3.5%	-0.02%	0.02%	
	7.4%	BS price	-0.81%	1.32%	26.9%	3.5%	3.5%	3.5%	-2.3.0%	26.9%	0.07%	0.02%	0.02%	0.02%	0.02%	0.02%	1.1%	1.1%	
	-0.01%	Delta	-0.10%	0.07%	0.07%	0.07%	0.07%	0.07%	-0.10%	0.07%	3.3%	1.1%	1.1%	1.1%	1.1%	1.1%	-0.008%	-0.008%	
	-0.4%	Vega	-4.8%	3.3%	3.3%	3.3%	3.3%	3.3%	0.037%	-0.026%	-0.008%	-0.008%	-0.008%	-0.008%	-0.008%	-0.008%	-0.008%	-0.008%	
	0.003%	Gamma	0.037%	-0.026%	-0.008%	-0.008%	-0.008%	-0.008%	0.037%	-0.026%	-0.008%	-0.008%	-0.008%	-0.008%	-0.008%	-0.008%	-0.008%	-0.008%	
	0.003%	Theta	0.037%	-0.026%	-0.008%	-0.008%	-0.008%	-0.008%	0.037%	-0.026%	-0.008%	-0.008%	-0.008%	-0.008%	-0.008%	-0.008%	-0.008%	-0.008%	
Spot-Vol Matrix																			
Price	-25	-8%	-6%	-4%	-2%	0%	2%	4%	6%	8%	1.25%	1.25%	1.25%	1.25%	1.25%	1.25%	1.25%	0.00%	0.00%
	-20	0.88%	0.88%	0.88%	0.88%	0.88%	0.88%	0.88%	0.88%	0.88%	0.88%	0.88%	0.88%	0.88%	0.88%	0.88%	0.88%	0.00%	0.00%
	-15	0.51%	0.51%	0.51%	0.51%	0.52%	0.52%	0.52%	0.52%	0.52%	0.52%	0.52%	0.52%	0.52%	0.52%	0.52%	0.52%	0.00%	0.00%
	-10	0.15%	0.15%	0.16%	0.18%	0.19%	0.19%	0.21%	0.21%	0.21%	0.21%	0.21%	0.21%	0.21%	0.21%	0.21%	0.21%	0.01%	0.00%
	-5	-0.12%	-0.08%	-0.05%	-0.02%	0.01%	0.01%	0.03%	0.03%	0.03%	0.03%	0.03%	0.03%	0.03%	0.03%	0.03%	0.03%	0.01%	0.00%
	0	0.03%	0.03%	0.02%	0.01%	0.00%	0.00%	-0.01%	-0.01%	-0.01%	-0.01%	-0.01%	-0.01%	-0.01%	-0.01%	-0.01%	-0.01%	0.01%	0.00%
	5	0.16%	0.08%	0.00%	-0.06%	-0.11%	-0.11%	-0.15%	-0.15%	-0.15%	-0.15%	-0.15%	-0.15%	-0.15%	-0.15%	-0.15%	-0.15%	-0.02%	0.00%
	10	-0.47%	-0.47%	-0.48%	-0.50%	-0.52%	-0.52%	-0.54%	-0.54%	-0.54%	-0.54%	-0.54%	-0.54%	-0.54%	-0.54%	-0.54%	-0.54%	-0.01%	0.00%
	15	-1.18%	-1.14%	-1.10%	-1.07%	-1.05%	-1.05%	-1.03%	-1.03%	-1.03%	-1.03%	-1.03%	-1.03%	-1.03%	-1.03%	-1.03%	-1.03%	0.01%	0.00%
	20	-1.56%	-1.55%	-1.53%	-1.51%	-1.48%	-1.48%	-1.46%	-1.46%	-1.46%	-1.46%	-1.46%	-1.46%	-1.46%	-1.46%	-1.46%	-1.46%	0.01%	0.00%
	25	-1.83%	-1.83%	-1.82%	-1.81%	-1.80%	-1.80%	-1.79%	-1.79%	-1.79%	-1.79%	-1.79%	-1.79%	-1.79%	-1.79%	-1.79%	-1.79%	0.01%	0.00%

TABLE 11.11 Price-Vol Matrix for the Reduced Risk Portfolio of Table 11.9 After 0.9 Years Have Elapsed

are 15 percent, a 5 percent proportional shift would move the one-year volatility up 1 to 21 percent and the five-year volatility up 0.75 to 15.75 percent. The **PriceVolMatrix** spreadsheet allows the user specification of either flat or proportional shifts.

With respect to the strike component, a frequently used alternative to a flat shift by instrument is a flat shift by delta. For example, assume that an at-the-money option currently has a 20 percent implied volatility and an in-the-money option with a delta of 75 percent currently has a 19 percent implied volatility, and assume that we are dealing with a currently at-the-money option. Then a volatility shift of down 2 percent combined with a price jump in the underlying asset that makes this option in-the-money with a 75 percent delta results in an implied volatility of $20\% - 2\% = 18\%$ if we are assuming a flat shift by instrument. It results in a $19\% - 2\% = 17\%$ implied volatility if we are assuming a flat shift by delta. The **PriceVolMatrix** spreadsheet allows the user specification of either flat shift by instrument or flat shift by delta.

The driving force behind the use of a flat delta shift is that the factors that generate the shape of the volatility surface by the strike, such as stochastic volatility and the structure of jumps, tend to remain static across changes in the underlying price level. We discuss these factors in Section 11.6.2. Taleb (1997, 138–142) provides a detailed exposition of a flat delta shift methodology and its consequences for hedging. Derman (1999) contrasts flat instrument shifts with flat delta shifts (“sticky-strike” versus “sticky-delta” in Derman’s terminology) along with a third possibility, “sticky-implied-tree.” Derman presents empirical evidence that differing market environments over time can result in a change in which shift patterns provide the greatest explanatory power.

No matter what selections are made for the price-vol matrix variables, there is clearly enough residual risk to require traders to also look at more detailed risk reports as supplements to price-vol matrices. Certainly, this will include exposure to changes in the shape of the volatility surface with respect to both time and strikes. The **PriceVolMatrix** spreadsheet includes a calculation of exposure to changes in the volatility surface. These more detailed reports usually focus only on the impact of small one-at-a-time changes, although a particularly significant residual risk might justify a price-vol matrix of its own. For example, an equity options trading desk might want to look at an overall price-vol matrix that considers parallel shifts in all stock market indexes as well as price-vol matrices for each individual country’s stock index, but would probably want only a simple delta and vega measure to reflect the exposure to each individual stock traded.

Senior management will want to see much less detail than the trading desk regarding options. The primary concern of senior management is

making sure that they are comfortable with large macro positions that may be an accumulation of the holdings of many trading desks. As such, the most important measure for senior management is outright exposure to spot positions (for example, JPY/USD FX, S&P index, and gold) or to forward positions (for example, exposure to a parallel shift in the USD interest rate curve). Since options desks hold delta-equivalent positions in these spot and forward markets, including these positions in reports of the total spot exposure of the firm is necessary in order to ensure an accurate summary. So senior management will generally just be interested in a single outright position number for each product, along with some measure of vega. For many options positions, the delta will fit the need for an outright position measure. Control of convexity risk around this delta is then left to the trading desk level, probably prescribed by limits on convexity. However, the positions of some complex trading books may not be at all accurately represented by the delta. If a book will gain \$100 million for the next 1-point rise in the S&P but lose \$2 million for each point rise after that, representing the position by a +\$100 million per point delta will be totally misleading. For senior management purposes, the delta needs to be defined not mathematically, as the instantaneous derivative, but economically, as a finite difference over a selected economically meaningful price movement (a one-standard-deviation daily price move might be a reasonable choice).

Limit-setting detail for options books lies somewhere between the level needed for trading desk control and that needed for senior management. Some form of limits on price-vol matrix positions is desirable, but separate limits for each matrix box would be overdetailed, whereas a single limit that no matrix box could exceed would be too broad. A limit set high enough to accommodate really unlikely combinations would be too liberal a limit for combinations close to the matrix center. A reasonable compromise is differentiated limits by groups of matrix boxes, where a similar likelihood of outcomes determines grouping. Limits on exposure to changes in the shape of the volatility surface can often be best expressed in terms of a few parameters that determine the shape. For details on possible parameters, see the discussion in Section 11.6.2.

The management of options risk is an inherently dynamic process. Unlike spot or forward risk, you can rarely just put on a hedge once and for all; you must constantly make adjustments. So options traders need measures to show them how their P&L and positions should change as a result of the passage of time or changes in prices. This enables them to prepare for the trading actions they will need to take and serves as a check against actual changes in P&L and positions to highlight anything that is happening that they don't understand. The best-known measures of this type are *theta* (the change in option values with time) and *gamma* (the change in delta with a

change in price). However, many other examples are available: for instance, *bleed* (see Taleb 1997, 191–199) and *Ddeltadvol* (Taleb 1997, 200–201).

By contrast, corporate risk managers are rarely interested in such measures. Theta cannot be a direct measure of risk since clearly you are not uncertain as to whether time will pass. It does measure the possibility of gain or loss if implied volatility fails to be realized over a given time period, but the same risk can be captured in a more comprehensive way by a time-bucketed vega measure. Gamma is of interest only to the extent that it can be used to compute convexity, which is a genuine P&L exposure, but gamma is a reliable indicator of convexity only for very simple portfolios. In general, corporate risk managers expect that trading desk heads will be able to deal with the operational issues of evolving positions. The only exceptions might be changes so large as to make liquidity questionable, which might require limits to be set.

11.5 DELTA HEDGING

In the presence of transaction costs, it is necessary to use optimization to determine a delta hedging strategy. A trade-off exists between achieving a lower standard deviation of results utilizing more frequent hedging, and achieving a higher expected return utilizing less frequent hedging leading to lower transaction costs. Whaley and Wilmott (1994) have shown that the efficient frontier for this problem consists of hedging policies with the following characteristics:

- Hedges will be triggered not by time intervals, but by the distance that the current delta hedge ratio differs from the theoretical delta hedge ratio required by the Black-Scholes formula.
- If transaction costs are only a function of the number of hedge transactions and not the size of the hedge transactions, then whenever a hedge transaction is triggered, the amount will be exactly enough to bring the hedge ratio in line with the desired theoretical ratio. Since the transaction cost is the same no matter how large the amount, you should go to the hedge ratio you would use in the absence of transaction costs.
- If transaction costs are only a function of the size of the hedge transaction, then whenever a hedge transaction is triggered, the amount of the transaction is only large enough to bring the difference between the actual and theoretical hedge ratios down to the trigger point. Since you don't care how many transactions you need to use, only the size of transactions, it makes sense that you will stay as close as possible to the point at which hedge inaccuracy exactly balances between the desire for low standard deviation of results and low transaction costs.

- If transaction costs are a function of both the number and size of hedge transactions, then the optimal rule will be a combination of these two cases, with an outer trigger distance between current and theoretical delta that institutes a trade to bring the difference down to an inner trigger distance.

Target delta hedges are determined by the Black-Scholes formula as $N(d_1)$, where $d_1 = [\ln(1/k) + \sigma^2 T / 2] / \sigma \sqrt{T}$. What value of σ , the volatility of the underlying asset, should be used to determine this target hedge? Options should be valued at the implied volatility that corresponds to the market price at which the position could be exited, but this does not provide any reason for using this implied volatility to determine delta hedges of positions that are not exited. Given that any misestimation of true volatility while determining the hedge will result in unintended proprietary positions in the underlying asset, as per our discussion in Section 11.3, it is best to give traders reasonable latitude to make their best estimate of future volatility as input to the target hedge.

This brings us to the suggested solution we promised to the question in Section 11.2. What causes the large losses from the nasty path? It is caused by the dramatic difference between actual realized volatility and implied volatility. You will see in the **NastyPath** spreadsheet that the option was priced at a 7 percent implied volatility, which was also used in creating the delta hedge. However, the actual price moves of 0.13 a day correspond to a realized volatility of 2 percent. Had the trader been able to foresee this and form the delta hedges based on a 2 percent volatility, P&L on the trade would have been close to 0 (try this out in the spreadsheet).

Continuing the theme from Section 11.3, concerning what actions to take if a trader believes the underlying price is mean reverting, simulations similar to those reported in Table 11.3 indicate that gains will result from delta hedges based on overestimates of actual realized volatility. If underlying prices are trending (either up or down) rather than mean reverting, then gains will result from delta hedges based on an underestimate of the actual realized volatility. So traders should consider biasing their volatility estimates if they have a view on mean reversion. To get an intuitive understanding of this result, consider what happens if you overestimate volatility. The higher volatility in the denominator of the formula for d_1 will cause the target delta to move less as price movements result in the option moving into or out of the money. If price moves tend to be followed by moves in the opposite direction, as they will be if the price process is mean reverting, then the difference between actual delta and theoretical delta will be in the right direction to create positive P&L.

11.6 BUILDING A VOLATILITY SURFACE

Building a volatility surface for pricing European options is similar to building a discount curve, but it operates in two dimensions rather than one, since volatilities will vary by strike as well as by time. However, the general principle is the same: Build a surface that balances the fitting of known options prices with a smoothness criterion. The smoothness criterion is designed to minimize the risk of loss from hedging options for which market prices are not known with options for which prices are known.

To build the surface in both dimensions simultaneously requires a stochastic volatility model to which you can fit parameters (for example, the Heston model—see Heston 1993). The more common approach is to build a volatility curve for at-the-money strikes by time period and separately build a volatility curve for a few selected time periods by strike. Arbitrary combinations of time and strike can then be interpolated from already determined points. We will look in turn at the issues of interpolating between time periods, interpolating between strikes, and extrapolating beyond the longest liquid time period.

11.6.1 Interpolating between Time Periods

We have a problem that's extremely similar to the one we faced for discount curves. We have a set of fitting conditions, wanting to choose underlying discount prices (implied volatilities), so that when they're plugged into pricing formulas, they come out with bond prices (option prices) that closely match those observed in the market, and a set of smoothness conditions, wanting to choose discount prices (implied volatilities) that lead to maximum smoothness of forward interest rates (forward volatilities) across periods.

The forward volatility, the amount of volatility expected to take place in some reasonably small time period in the future, is a natural analogy to the forward rate. With forward rates, we discussed whether to have an additional set of constraints stating that all forward rates must be nonnegative and examined economic arguments for and against this (refer to Section 10.3.2). With forward volatilities, there isn't any doubt—a negative standard deviation is not a mathematical possibility, so the constraints are necessary.

We can set up an optimization to solve for forward volatilities in a completely analogous manner to the optimization we set up to solve for forward rates, with different solutions corresponding to different trade-offs between the tightness of the fitting constraints and tightness of the smoothness constraints and different weights on different fitting constraints based on the liquidity of the price quotes. (Note that it is a more viable possibility

with options than with interest rates to just find forwards that exactly fit all available market prices and then interpolate between the forwards. Unlike bonds and swaps, options have no intermediate payments to require a bootstrap. However, optimization still might be desirable as a way of trading off between fitting and smoothness objectives.)

When fitting forward interest rates, we had to preprocess to adjust for the lack of smoothness that we were anticipating based on our economic theories, such as turn-of-the-quarter effects (see Section 10.3.4). In the same way, forward volatilities need preprocessing. Generally, the opinions of options traders regarding the patterns of forward volatility tend to be much more strongly held than the opinions of interest rate traders regarding forward rates. Opinions on forward volatility center on issues of the flow of information into the markets that will cause price fluctuations. If we look at daily forward volatilities (and traders of shorter-term options often do work at this level of detail), you might find a trader anticipating nearly zero volatility on weekends and holidays (markets are closed so no new prices can be observed), higher volatility on Mondays and days after holidays than on other weekdays (governments sometimes like to make surprise announcements when markets are closed), lower than normal volatility on days when most traders can be expected to be on vacation or leaving work early (such as the day before a three-day weekend), and higher than normal volatility on a day when a key economic statistic is scheduled to be announced. For more examples, see Taleb (1997, 98) and Burghardt and Hanweck (1993).

The website for this book has two spreadsheets to illustrate fitting a forward volatility curve to observed options prices. The first, **VolCurve**, can be used for all European options other than interest rate caps and floors, and emphasizes the adjustment for anticipated volatility patterns. The second, **CapFit**, is designed for use only for interest rate *caps* and *floors*, which are packages of individual options (known as *caplets* and *floorlets*, respectively). Since liquid prices are generally available only for the options packages and not for the underlying options, an optimization is needed to fit the observed prices of packages with as smooth a forward volatility curve as possible.

11.6.2 Interpolating between Strikes—Smile and Skew

Now let's turn to building a volatility curve by strike for a given time period. Market prices will be available for certain strikes that we will want to fit. Which variable should play the corresponding role to forward interest rates and forward volatilities as the one for which we try to achieve smoothness? A natural choice is the risk-neutral probability that the underlying variable finishes in a range between two prices. If these ranges are chosen small

enough, options at all strikes can be priced to as close a precision as you want based on such probabilities.

If S is the strike and p_i is the risk-neutral probability that the underlying will finish between price P_i and price P_{i+1} , the option price must be bounded by $\sum_i \max(P_i - S, 0)p_i$ from below, and bounded by $\sum_i \max(P_{i+1} - S, 0)p_i$ from above.

Like forward volatilities, probabilities must be constrained to be non-negative. Using this formula allows translation among cumulative probability, probability frequency, and implied volatility by strike as alternative, mutually translatable ways of describing a probability distribution, in much the same way that par rate, zero coupon rate, forward rate, and discount price are alternatives for describing the discount curve. See the **VolSurfaceStrike** spreadsheet for an illustration of this principle.

Jackwerth and Rubinstein (1996) illustrate an optimization setup to derive probability distributions based on a trade-off between the tightness of fitting constraints and smoothness constraints. When choosing a smoothness criterion, an alternative to just looking at how smooth the changes in probability levels are is to look at how closely the probabilities fit a distribution selected on theoretical grounds (for example, normal or lognormal) as the most likely prior distribution (prior, that is, to any knowledge of the actual options prices). This use of prior distribution ties closely to Bayesian statistical methods. In Section III.A of their paper, Jackwerth and Rubinstein explore several such smoothness criteria.

A fundamental problem often encountered when trying to derive volatility curves by strike is the relative paucity of market observations available by strike. It is not at all uncommon to find markets in which options prices are available for only three or four different strike levels at a given time period. In such circumstances, a smoothness criterion that does not utilize a prior distribution is of little use—you at least need to restrict your choice to some family of possible candidate distributions on theoretical grounds. Of course, any such choice is a model and should be analyzed for the degree of mispricing possible if the model is wrong by considering how different the volatility curve would be if another plausible model were chosen. Reserves and limits against model error should be considered.

A good discussion of candidate distributions and the theoretical basis for selecting between them can be found in Hull (2012, Sections 26.1–26.3). Let us first state some general facts about the shape of volatility surfaces observed in the markets; these comments can be compared with those in Hull (2012, Sections 19.2 and 19.3) and Rebonato (2004, Chapter 7). In this discussion, we use the term *smile* to refer to a pattern of volatility by strike where volatility rises as strikes move away from at-the-money in the

direction of either into or out of the money. We use *skew* to refer to a pattern of volatility by strike in which volatility either decreases or increases with increasing strike levels. So skew is primarily a linear relationship and smile is primarily a quadratic relationship. (Market practice from firm to firm, and even desk to desk within a firm, may differ in nomenclature. Sometimes *skew* is used to cover all aspects of volatility surface shape, and sometimes *smile* is used to cover all aspects of volatility shape.)

Using these definitions, the observed patterns are:

- Smiles tend to appear in all options markets.
- Equity options markets almost always show a pronounced skew, with volatility decreasing with increasing strikes. The combination of this skew with the smile produces a pattern that can be described as a sharp skew at strikes below at-the-money and relatively flat volatilities at strikes above at-the-money.
- No general skew pattern exists in markets for FX options between strong currencies (for example, between the dollar, euro, yen, sterling, and Swiss franc). However, there does tend to be a strong skew pattern (volatility decreases with increased strikes) for FX options between a strong currency and a weaker currency such as an emerging market currency.
- Skew patterns in interest rate options markets tend to vary by currency, with the strongest patterns of volatilities decreasing with increasing strikes appearing for currencies with low interest rate levels, particularly in yen.

What explanations have been offered for these observed patterns?

- The prevalence of volatility smiles can be explained in two different ways: stochastic volatility and jump diffusion. Stochastic volatility utilizes a probability distribution for the volatility that determines the probability distribution of underlying prices, whereas jump diffusion assumes that some price uncertainty is expressed through price jumps as opposed to a smooth random walk. Both assumptions result in a distribution of final prices with fatter tails than the lognormal distribution used by Black-Scholes. Fatter-tailed distributions have little effect on options at close to at-the-money strikes, which are primarily affected by the center of the distribution; however, they have greater effects the more an option is in-the-money or out-of-the-money, since these options are primarily affected by the size of the tail.

The pricing formula for options using either stochastic volatility or jump diffusion (see the equations in Hull 2012, Sections 26.1 and

26.2) consists of averages of option prices using the Black-Scholes formula across a range of volatilities. The difference between the two models is the probability weight used in averaging across these volatilities. Stochastic volatility results in a more pronounced smile as the time to option expiry increases, whereas jump diffusion results in a more pronounced smile as the time to option expiry decreases. It may be necessary to combine the two to obtain actual smile patterns observed in market options prices. See Matytsin (1999).

- The Black-Scholes model assumes a lognormal distribution of the underlying asset price. If the market is assuming a normal, rather than lognormal, price distribution, this will evidence itself as higher implied volatilities for lower-strike options and lower implied volatilities for higher-strike options when implied volatilities are computed using the Black-Scholes formula. So if the market is assuming that price changes are independent of market level rather than proportional to market level, implying normal rather than lognormal price distributions, this will lead to a skew with volatilities decreasing with increasing strikes. If the market is assuming a distribution intermediate between normal and lognormal, this skew pattern will still exist, but it will be less pronounced. Historical evidence shows support for interest rate movements that are sometimes closer to being independent of the rate level and other times closer to being proportional to the rate level. The skew for implied volatilities of interest rate options is generally believed to be driven primarily by the expectation that rate movements are not completely proportional to the rate level, with the expectation in low-rate environments that rate movements are close to independent of the rate level.
- The skew pattern in equity markets has sometimes been explained as the outcome of asymmetry of the value of investment in a corporation, which can suddenly collapse as a company approaches the bankruptcy point. Hull (2000, Section 17.7) discusses three alternative models based on this explanation—the compound option model, the displaced diffusion model, and the constant elasticity of variance model.

A more general explanation of skew patterns can be found in analyzing the degree of asymmetry in the structure of a particular market. For a thorough exposition of this viewpoint, see Taleb (1997, 245–252), on which much of my discussion here is based. This asymmetry can be described in two complementary ways: one that focuses on investor behavior and the other that focuses on price behavior.

From an investor behavior viewpoint, in some markets, investment has a structural bias toward one side of the market. Equity markets are a good

example. There are far more investors long equity investments than there are investors who have shorted the market; hence, more investors are seeking protection from stock prices falling than are seeking protection from stock prices rising. The reason is that corporate issuance of stock is a major source of supply, and corporations are not seeking protection against their stock rising; in fact, they welcome it. So you expect to see greater demand to buy puts on stock at strikes below the current market level, sought by investors protecting their long equity positions, than the demand for calls on stock at strikes above the current market level sought by short sellers to protect their short equity positions. This imbalance in demand drives up implied volatilities on low-strike options relative to high-strike options.

The complementary view from a price behavior viewpoint is that stock market crashes, in which large downward jumps occur in stock prices, are far more common than large upward jumps in stock prices. This can be seen as a consequence of the imbalance in investors who are long stocks relative to those who are short stocks. Falling prices can trigger a selling panic by investors faced with large losses forced to exit leveraged long positions supported by borrowings. There are fewer short sellers and leveraged short positions to cause a panic reaction when prices are rising. A bias toward downward jumps over upward jumps leads to a skew in the distribution of probabilities of price movements that will translate into higher implied volatilities at lower strikes. In addition, the anticipation of possible stock market crashes will exacerbate the demand for crash protection through puts at lower strikes.

A similar structural analysis can be constructed for FX markets for an emerging market currency versus a strong currency. These exchange rates are often maintained at artificially high levels by governments defending the value of the emerging market currency through purchases of the currency, high interest rates, or currency controls. When breaks in the FX rate come, they tend to be large downward jumps in the value of the emerging market currency. There is no similar possibility of upward jumps. This price behavior directly leads to a probability distribution that translates to higher implied volatilities at lower strikes (lower in terms of the value of the emerging market currency). Indirectly, this price behavior encourages holders of the emerging market currency to buy puts at lower strikes, bidding up the implied volatility at these strikes.

Other markets generally tend toward a more symmetrical structure. Exchange rates between two strong currencies are usually freer floating with less bottled-up pressure. Thus, no bias exists toward large upward jumps or large downward jumps. Most interest rate markets and commodity markets tend to be roughly evenly divided between longs and shorts—investors who would benefit from upward movement and those who would benefit

from downward movement. However, some particular asymmetries can be observed—for example, the large demand by U.S. mortgage investors for protection against falling interest rates leading to accelerated prepayments or a temporary imbalance of the suppliers of a commodity seeking put protection against falling prices relative to the consumers of the commodity seeking call protection against rising prices.

The **VolSurfaceStrike** spreadsheet illustrates both ways in which a probability distribution can be fit to a set of option prices at different strikes. With input on prices at a number of different strikes, it trades off the smoothness of the probability distribution and closeness of price fit. With input on prices at only a few strikes, it fits two parameters: one representing standard deviation of volatility and one representing the degree of proportional versus absolute price change to assume.

11.6.3 Extrapolating Based on Time Period

When we were looking at forward risk, we saw how to create valuation and reserves for a forward that had a longer tenor than any liquid instrument (see Section 10.2.2). The technique was to assume you were going to hedge the longer-term forward with a liquid shorter-term forward and later roll the shorter-term forward into a longer-term forward. The expected cost of the roll needs to be added into the initial cost of the hedge to obtain a valuation, and a reserve can be based on the historical standard deviation of the roll cost.

A similar approach suggests itself for valuing and reserving for long-term options that have a longer tenor than any liquid option. For example, if you want to create a 10-year option in a market that has liquid quotes only out to seven years, you could begin by hedging with a seven-year option and, at the end of five years, roll out of what will then be a two-year option into the five-year option you need to exactly match your actual position. Expected differences in implied volatility between five- and two-year options determine expected roll costs. Reserves can be based on the historical standard deviation of differences in two- and five-year implied volatilities.

However, options are more complicated because they depend on strike level as well as the time to expiry. The price-vol matrix in Table 11.12 shows that a ratio of seven-year options to 10-year options selected so as to minimize roll-cost uncertainty when the prices are at 100 leaves large roll-cost uncertainty when prices are above or below 100.

To minimize roll-cost uncertainty over a wide range of prices, you need to hedge with a package of options that differ by both tenor and strike. The price-vol matrix in Table 11.13 shows the impact of selecting a hedge from a set of six- and seven-year options at various strike levels, using the **OptionRoll** spreadsheet to select weightings of these options that will achieve

Discount Spacing	5.00%		Years		5		5		5		5		5		5	
	Volume	call/put	call	call	call	call	call	call	call	call	call	call	call	call	call	call
Price	-1	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100
Volatility	5	100	80	100	90	100	110	120	80	90	100	110	120	100	110	120
	2%	5	2	2	2	2	2	2	1	1	1	1	1	1	1	1
Portfolio	Implied vol	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%
	BS price	-13.78%	-14.42%	38.49%	-6.74%	4.54%	1.95%	3.65%	5.93%	-17.78%	6.53%	-1.78%				
	Delta	-58.8%	-56.9%	180.5%	-36.9%	28.4%	13.7%	-16.1%	33.7%	-126.7%	56.5%	-18.2%				
	Vega	-0.68%	-0.23%	1.16%	-0.33%	0.34%	0.20%	-0.03%	0.14%	-0.89%	0.57%	-0.24%				
	Gamma	-0.9%	-0.6%	3.2%	-0.9%	0.9%	0.6%	-0.2%	0.8%	-4.7%	3.0%	-1.3%				
	Theta	0.005%	0.004%	-0.022%	0.006%	-0.006%	-0.004%	0.001%	-0.006%	0.034%	-0.022%	0.009%				
	Current	-15.05%	-14.78%	46.04%	-9.75%	8.16%	4.49%	-3.91%	8.09%	-33.65%	18.60%	-8.22%				

Spot-Vol Matrix Price	Implied Volatilities										Vega	
	-8%	-6%	-4%	-2%	0%	2%	4%	6%	8%	Vega	Convexity	
-25	0.24%	0.17%	0.10%	0.05%	0.01%	-0.02%	-0.05%	-0.06%	-0.06%	-0.02%	0.00%	
-20	-0.03%	-0.03%	-0.05%	-0.06%	-0.07%	-0.07%	-0.07%	-0.06%	-0.05%	-0.00%	0.00%	
-15	-0.08%	-0.07%	-0.07%	-0.08%	-0.07%	-0.06%	-0.05%	-0.03%	-0.01%	0.00%	0.00%	
-10	0.00%	-0.02%	-0.04%	-0.05%	-0.04%	-0.03%	-0.01%	0.01%	0.04%	0.00%	0.00%	
-5	0.10%	0.03%	0.00%	-0.01%	-0.01%	0.00%	0.02%	0.05%	0.08%	0.00%	-0.01%	
0	0.12%	0.04%	0.01%	0.00%	0.00%	0.02%	0.04%	0.07%	0.10%	0.01%	-0.01%	
5	0.07%	0.01%	-0.01%	-0.02%	-0.01%	0.01%	0.04%	0.07%	0.10%	0.01%	-0.01%	
10	0.00%	-0.03%	-0.04%	-0.04%	-0.03%	-0.01%	0.02%	0.05%	0.08%	0.01%	0.00%	
15	-0.02%	-0.04%	-0.06%	-0.06%	-0.05%	-0.03%	-0.01%	0.02%	0.06%	0.01%	0.00%	
20	0.04%	0.00%	-0.04%	-0.05%	-0.06%	-0.05%	-0.03%	-0.01%	0.02%	0.00%	0.00%	
25	0.19%	0.10%	0.04%	-0.01%	-0.03%	-0.04%	-0.04%	-0.03%	0.00%	-0.01%	0.00%	

TABLE 11.13 Hedge to Roll Over into a 10-Year Option

minimal roll-cost uncertainty in five years. This example only accounts for roll-cost uncertainty due to shifts in volatility level; a more complete treatment would include shifts in the shape of the volatility surface. Expected roll costs and standard deviations of roll costs must now be computed relative to the weighted average of implied volatilities of the hedge package.

11.7 SUMMARY

By way of summary, let us see how the paradigm for managing vanilla options risk deals with the criticisms of the Black-Scholes analysis that have been offered. Compare the analysis here to Taleb (1997, 110–113).

- Black-Scholes unrealistically assumes a constant risk-free interest rate and drift rate of the forward. The way we have set up our Black-Scholes model, directly incorporating rate and drift volatility into the volatility of the forward, shows that this criticism is not a serious one.
- Black-Scholes assumes that asset prices are lognormally distributed. This has long ceased to be true in trading practice. With traders valuing positions at each strike at different market-observed volatilities, any probability distribution believed by the marketplace can be accommodated. In part 2 of Exercise 11.2 you are asked to examine the success of hedging options at one strike with those at another strike, using a Monte Carlo simulation that does not assume asset prices to be lognormally distributed. You will find relatively small uncertainty of hedging results.
- Black-Scholes assumes that hedging in the underlying asset can take place continuously and without transaction costs. These assumptions are closely linked since the presence of transaction costs will certainly force hedging to be less frequent, even if more frequent hedging is theoretically possible. Our Monte Carlo simulations have shown that, with the use of options to hedge other options, the resulting positions can be delta hedged at discrete times, resulting in relatively small uncertainty of hedging results and relatively low transaction costs. Any uncertainty and transaction costs that remain will contribute to wider bid-ask spreads for options.
- Black-Scholes assumes that underlying asset prices will follow a Brownian motion with no sudden jumps. In practice, sudden jumps do occur and these are unhedgeable other than by offsetting options positions. The price-vol matrix reports exposure to price jumps. In part 1 of Exercise 11.2, you are asked to examine the success of hedging options at one strike with those at another strike, using a Monte Carlo simulation that assumes price jumps will take place. You will find relatively small uncertainty of hedging results.

- Black-Scholes assumes that volatility is constant. This is obviously false. The implications of stochastic volatility for the standard deviation of hedging results have been noted. The price-vol matrix reports exposure to changes in volatility, and positions that have small exposure as measured by the price-vol matrix have been shown, using Monte Carlo simulation, to have a relatively small uncertainty of hedging results.
- Black-Scholes assumes that volatility is known. This is also obviously false. Our Monte Carlo simulations were carried out under the assumption that actual volatility was not known when setting hedge ratios, and the resulting uncertainty of hedging results is small.

EXERCISES

11.1 Options Portfolio Risk Measures

Start with a portfolio consisting of less liquid options as follows:

Volume	1	-1	-1
Call/put	Call	Call	Call
Price	100	100	100
Strike	84	107	114
Time to expiry	1.3	0.7	1.7
Implied volatility	25.9%	24.4%	23.0%

1. Calculate the risk exposure of this portfolio.
2. Use the Solver to minimize risk using more liquid options. The liquid options available are as follows:

Call/Put	Call	Call	Call	Call	Call	Call
Price	100	100	100	100	100	100
Strike	100	90	110	100	90	110
Time to expiry	1	1	1	2	2	2
Implied volatility	25.5%	26.0%	24.0%	24.5%	25.5%	23.5%

3. Compare the risk exposure of the risk-minimized portfolio to that of the original portfolio. How much has the risk been reduced? How would you characterize the exposures that remain?

4. Is this a static hedge or will it need to be rehedged through time?
5. Create your own portfolio of less liquid options and go through the same exercise.

11.2 Monte Carlo Simulation of Options Hedging

Program a Monte Carlo simulation to compare the results of dynamic hedging on a single option position and on an option hedged by other options. To begin, try to match the results in Table 11.2. Start with eight simulations ($2 \times 2 \times 2$), corresponding to a pure dynamic hedge/two-sided options hedge, 0 percent standard deviation of volatility/33 percent standard deviation of volatility, and 100 rebalancings/500 rebalancings. Use 1,000 paths for each simulation. When using a standard deviation for volatility, apply it at the point that a volatility is assigned to a path (if you let the volatility vary at each rebalancing along the path, the volatilities will average out along the path and little difference will exist between the results of your 0 percent standard deviation and 33 percent standard deviation cases).

For all eight cases, initial price = strike = 100, time = 1 year, average volatility = 20 percent, rate = dividend = 0 percent, and transaction costs are based on one-fourth point per \$100 bid-asked spread, so any transaction, either buy or sell, incurs a cost of \$0.125 per \$100 bought or sold (but don't charge any transaction cost for establishing the initial delta hedge).

Use the **OptionMC** and **OptionMCHedged** spreadsheets to check your simulation programs. To do this, run your simulation for just 20 time steps. You can then check a particular path by taking the random numbers drawn for that path and substituting them for the random numbers selected in the spreadsheets. You can then compare results.

Once you match the results from Table 11.2, you should try to expand the runs in the following ways:

1. Four runs with 100 rebalancings/500 rebalancings for the pure dynamic hedge/two-sided options hedge, a 33 percent standard deviation of volatility, and a jump process. Jumps should occur on average once on each path, there should be a 50–50 chance that a jump is up or down, and the average absolute jump size should be 10 percent of the current price with a 33 percent standard deviation around this 10 percent. So a one-standard-deviation range

would be from $10\% \times \exp(-0.33) = 7.2\%$ to $10\% \times \exp(0.33) = 13.9\%$. The volatility of the underlying should be adjusted down from 20 percent to whatever level will leave the average pure dynamic hedging cost equal to what it had been without the jump (you will need to try out a few different volatility levels to determine this).

2. A similar set of runs to test the impact of a volatility skew with a standard deviation of 20 percent.
3. For the case of a pure dynamic hedge, 500 rebalancings, and 33 percent standard deviation of volatility, check the impact of imposing different threshold levels for rehedging on the trade-off between the expected transaction cost and standard deviation of P&L.
4. Examine a sample of 50 individual paths and observe the relationship between the final price of the underlying and the total hedge P&L. Does the observed relationship support the claim in Section 11.3 that “despite wild gyrations in underlying prices, [the simulation paths] produce almost identical hedging results”?
5. What pattern do you observe of hedge ratios along the individual paths? For example, how quickly does the hedge ratio go to 100 percent for paths whose final price is above the strike and to 0 percent for paths whose final price is below the strike?

For parts 4 and 5 of this exercise, you need to examine individual paths of the Monte Carlo simulation. Use paths taken from the simulation with 0 percent standard deviation of volatility and 500 rebalancings. If you do not have the time or programming background to create your own Monte Carlo simulation, then carry out these parts of the exercise using the **OptionMC1000** and **OptionMCHedged1000** spreadsheets. Use the following input settings: price = \$100, strike = 100, time to expiry = 1, implied volatility = 20 percent, volatility = 20 percent, skew = 0 percent, and jump probability = 0 percent.

Managing Exotic Options Risk

We need to first determine what we mean by an *exotic option*. Some articles on options emphasize complex formulas and difficult mathematical derivations as the hallmarks that distinguish exotics from vanillas. The criterion I am using in this book emphasizes market liquidity. If you can readily obtain prices at which the option can be bought and sold, then it counts as a vanilla option; if not, then it is an exotic option.

To understand why I favor this definition, consider a forward-start option as an illustrative example. This is an option priced now, but its strike is not set until some future date. Generally, it is set to be at-the-money on that future date. There is certainly no complexity about the formula or mathematical derivation of the formula for this product. It is the standard Black-Scholes formula with the strike and underlying price set equal. However, this product has no liquid market, and relating its valuation and hedging to the valuation and hedging of ordinary European options is not straightforward. Equivalently, we can say that no clear relationship exists between the volatility that is needed as input to the Black-Scholes formula for the forward-start option and the volatilities implied by the prices of standard European options.

The two preceding chapters, on managing forward risk and vanilla options risk, emphasized the use of methods that maximize the degree to which all transactions can be viewed as being managed within a common risk-measurement framework—a single discount curve for forwards and the price-vol matrix for vanilla options. This common framework increases the chance that exposures on different transactions can be netted against one another and offset by transactions involving the forwards and vanilla options with the greatest liquidity. This paradigm does not work for exotic options since none of them have enough liquidity to provide confidence that risks can be offset at publicly available prices.

Therefore, the emphasis throughout this chapter is on methodology that enables, as much as possible, the risks in an exotic option to be represented

as an equivalent vanilla option and forwards position. The vanilla option and forwards position is the liquid proxy representation of the exotic discussed in Section 6.1.2 and in Section 8.4. My primary arguments for anchoring exotic option risk management to a liquid proxy are presented there. Additional reasons were discussed in the arguments favoring a more detailed limit approach in Section 6.2, in the broader setting of general risk decomposition. As applied specifically to exotic options, these reasons are:

- It permits the separation of exotic options risk into a part that can be managed with vanilla options and a residual that cannot. It is important to identify and quantify this residual risk so that adequate reserves can be held against it, and to facilitate the management recognition of pricing that is inadequate to support actual hedging costs. Without separating out the part of the risk that can be hedged with liquid vanilla options, it is quite possible that gains from ordinary vanilla risk positions will obscure losses from the truly illiquid residual.
- It encourages as much of the risk as possible to be managed as part of the far more liquid vanilla options position.
- It reduces the risk of having exotic options positions valued with a methodology that is inconsistent with that used for valuing the vanilla options positions.
- It consolidates exotic options positions into already well-developed reporting mechanisms for vanilla options—price-vol matrices, volatility surface exposures, deltas, and other “greeks.” This has the advantage of building on well-understood reports, thus clarifying the explanation to senior managers. It also guards against large positions accumulating without being recognized by a common reporting mechanism, since all exotics for a given underlying will be consolidated into the same set of vanilla options risk reports.

The use of liquid proxies in risk management closely parallels the use of the *control variate technique* in modeling. The control variate technique uses the best available model to value a particular exotic option, but it also uses the same model to value a related vanilla option (or basket of vanilla options). Since the vanilla options are liquid, they can then be valued directly from the market (or interpolated from direct market prices). The model is only used to value the difference between the exotic and related vanilla. Risk reporting and risk management are similarly divided between reporting and managing the risk of the related vanilla option as we would any other vanilla and creating separate risk reporting and management for the difference between the exotic and related vanilla, thereby reducing the model dependence of valuation. See Hull (2012, Section 20.3) for a discussion that

emphasizes the computational efficiency of this technique, which is strongly analogous to the risk management advantages stressed here.

If the underlying assumptions of the Black-Scholes framework were true—in particular, if volatility was known and constant—the choice of models for exotics would generally be easy. Most exotics can be valued using formulas derived from market assumptions similar to those used in the Black-Scholes analysis of European options. However, when volatility is unknown and variable, there is seldom a direct way of translating a volatility surface used for valuing European options into a single volatility to be used in valuing an exotic. Usually, we will need to rely on more complex formulations tailored to a particular exotic to establish this relationship. Much of this chapter is devoted to developing these formulations for specific exotics.

A distinction that will prove very important when analyzing these models can be made between those where the relationship between the exotic and the vanilla is *static* and those where the relationship between the exotic and the vanilla is *dynamic*. Static relationships mean that the same vanilla (or package of vanillas) can be used to represent the exotic in vanilla option risk reports throughout the life of the exotic. Dynamic relationships mean that the package of vanillas used may need to change in composition over the life of the exotic. Dynamic relationships correspond to the full simulation approach recommended by Derman (2001) that was discussed in Section 8.4. Static relationships, and the quasistatic relationships I will discuss in a moment, correspond to the simulation of a limited hedging strategy I propose as an alternative to Derman's full simulation approach in Section 8.4.3.

Static representations have obvious operational advantages. Once it is booked at the inception of a trade, the representation does not need to be updated. Even more important is the simplicity introduced when the potential cost of differences between the actual exotic and its vanilla option representation over the life of the transaction is estimated. As emphasized in Section 11.3, dynamic representation requires simulation to evaluate potential costs. However, the simulation of dynamic changes in vanilla option hedges can be far more computationally difficult than the simulation of dynamic changes in underlying forwards hedges studied in Section 11.3. The reasons for this are discussed in Section 12.3.2.

The ideal of a static representation cannot always be achieved. It will be possible in Sections 12.1 and 12.4 when we are discussing options whose payout depends on prices at a single future time. However, when discussing options whose payout is a function of prices at different times, as is true in Sections 12.2, 12.3, and 12.5, static representation will not be possible. Our alternatives will be either dynamic representation or *quasistatic* representation, in which changes in the representation are minimized, often to only a single change, to simplify calculations of potential cost. We will use the

simpler term *static* for the remainder of this chapter, but this is shorthand for *quasistatic* representation, and will pay due attention to the estimation of the cost of the hedge changes.

Table 12.1, which was taken from Smithson (2000), shows the principal forms of exotic products and how widely they are used in different markets.

The study of exotic options in this chapter is divided into five sections, following the categories used in Table 12.1.

- **Section 12.1—single-payout options.** These are options whose payoffs are the function solely of the price of an underlying asset at a single future time. We will show how to replicate these options exactly using a basket of forwards and vanilla options. The resulting replication can be used both to value the exotic and represent it in risk reports. The only residual risk will be the liquidity of the resulting basket, particularly in the replication of binary options. A particular example of an important single-payout exotic is a log contract, which makes payments based on the logarithm of the underlying price. Its importance is mostly due to its close linkage to a variance swap, an exotic product not in Table 12.1 but that shows increasing use. We also discuss the volatility swap in this section, a close cousin of the variance swap.
- **Section 12.2—time-dependent options.** These are options whose payoffs are the function of the price of a vanilla option at a single future time. As in Section 12.1, we will show how to eliminate all risk of underlying price movement for these exotics by replication using forwards and vanilla options. The residual risk exposure to implied volatility at a future time can be quasistatically hedged with vanilla options. These exotics include forward-start options, cliquet options, chooser options, and compound options.
- **Section 12.3—path-dependent options.** These are options whose payoffs depend on the price of a single underlying asset at several future times. We will focus on barrier options, but also use the lessons learned to apply to ladder, lookback, double barrier, and partial-time barrier options. We will examine and contrast replication approaches that utilize dynamic hedging with vanilla options and approaches that permit quasistatic hedging with vanilla options.
- **Section 12.4—correlation-dependent options.** These are options whose payoffs depend on the prices of several underlying asset securities and that therefore must be priced based on assumptions about correlations. We will examine several important cases: basket forwards and options, quanto forwards and options, diff swaps, mortgage-backed securities, collateralized debt obligations (CDOs), and convertible bonds.
- **Section 12.5—correlation-dependent interest rate options.** A particular subset of correlation-dependent options are options whose payoffs

TABLE 12.1 Intensity of Use of Option Structures in Various Markets

	Interest Rate Options			FX Options			Equity Options			Commodity Options								
	Options			Exchanges			OTC			Exchanges			OTC			Exchanges		
	OTC	Exchanges	OTC	Exchanges	OTC	Exchanges	OTC	Exchanges	OTC	Exchanges	OTC	Exchanges	OTC	Exchanges	OTC	Exchanges		
First-generation options																		
European style	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	
American style	A	A	A	A	A	A	R	A	R	A	A	A	A	A	A	A	A	
Bermuda style	A	A	A	A	A	A	R	R	R	R	R	R	R	R	R	R	R	
Second-generation options																		
<i>Path-dependent options</i>																		
Average price (rate)	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	
Barrier options	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	
Capped	O	O	O	O	O	O	R	O	R	O	R	O	R	O	R	O	R	
Lookback	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	
Ladder	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	
Ratchet	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	
Shout	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	
<i>Correlation-dependent options</i>																		
Rainbow options	R	R	R	R	R	R	O	O	R	O	R	O	R	O	R	O	R	
Quanto options	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	
Basket options	R	R	R	R	R	R	A	A	A	A	A	A	A	A	A	A	A	
<i>Time-dependent options</i>																		
Chooser options	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	
Forward-start options	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	
Cliquet options	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	
<i>Single-payoff options</i>																		
Binary options	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	
Contingent premium options	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	

FX = foreign exchange; OTC = over the counter; A = actively used; O = occasionally used; R = rarely used; blank = not used. Source: Smithson (2000).

depend on multiple future interest rates. This includes the important special case of American and Bermudan swaptions.

12.1 SINGLE-PAYOUT OPTIONS

In continuous time finance, the Breeden-Litzenberger theorem states that any option whose payout is a smooth function of a terminal forward price can be perfectly replicated by an infinite package of forwards and plain-vanilla calls and puts (see Carr and Madan 2002, Section II.A). The discrete time version states that any option whose payout is a smooth function of a terminal forward price can be replicated as closely as desired by a finite package of forwards and plain-vanilla calls and puts, with the tightness of fit of the replication dependent on the number of vanilla calls and puts in the package. In both cases, replication is static, meaning the forwards and vanilla calls and puts are purchased at the deal inception and then no further hedging is needed. The terminal payout on the replicating package will match the terminal payout of the exotic option.

The discrete time result can be established in two stages:

1. Any smooth function can be approximated as closely as desired by a piecewise-linear function. The tightness of fit depends on the number of pieces of the replication.
2. Each piece of a piecewise-linear function can be replicated by adding another vanilla option to a package of options that replicates all of the pieces up to that point. This can be easily seen from an example.

Consider a function that pays out nothing at prices \$100 or below, pays out \$2 for every \$1 gain in price up to \$102, pays out \$3.5 for every \$1 gain in price from \$102 to \$105, and pays out \$2.3 for every \$1 gain in price above \$105. This payout can be replicated by buying 2 calls at \$100 and 1.5 calls at \$102, and selling 1.2 calls at \$105, as shown in Table 12.2.

The **BasketHedge** spreadsheet on the website for this book enables you to calculate the vanilla option hedges and the associated valuations based on this discrete time approach. The impact of smiles and skews in the volatility surface of the vanilla options on the valuation of the exotic options can be readily calculated using this spreadsheet.

Even if this is not selected as a desirable hedge from a trading viewpoint, it still makes sense as a way to represent the trade from a risk management viewpoint for the following reasons:

- It permits realistic valuation based on liquid, public prices. Alternative valuation procedures would utilize an analytic pricing model, which is

TABLE 12.2 Vanilla Options Replication of a Piecewise-Linear Payout

Price	Payout	+2 Calls at \$100	+1.5 Calls at \$102	-1.2 Calls at \$105
\$100	0.0	0.0	0.0	0
\$102	4.0	4.0	0.0	0
\$105	14.5	10.0	4.5	0
\$110	26.0	20.0	12.0	-6.0

usually easily derivable, but a level of volatility needs to be assumed and no straightforward procedure is available for deriving this volatility from observed market volatilities of vanilla options at different strikes. The hedge package method will converge to this analytic solution as you increase the number of vanilla option hedges used, provided all vanilla options are priced at a flat volatility (it is recommended that this comparison always be made as a check on the accuracy of the implementation of the hedging package method). However, the hedging package method has the flexibility to price the exotic option based on any observed volatility surface (in fact, instead of using the volatility surface, the directly observed vanilla option prices are used, so the pricing is not dependent on any option model).

- The hedge package method gives an easy means of integrating exotic options into standard risk reports, such as price-vol matrices and vega exposure by strike and maturity. Placing as much risk as possible within a single context also increases the chances that risks from one position may offset risks in another position. Only the net risks need to be managed. (See the arguments for requiring internal hedging in Section 6.2.)
- Although the representation will be incomplete due to the use of a finite package of vanilla options, the residual risk can be easily calculated by Monte Carlo simulation based on an assumed probability distribution of final forward prices multiplied by the amount of mishedge. This is an easier calculation of remaining risk than the analytic method, which requires a Monte Carlo simulation of dynamic hedging.

One objection that is sometimes raised to the static hedging strategy for exotic options is that the required basket of vanilla options is unrealistic, in terms of using options at strikes that have little market liquidity, in terms of the number of different options in the basket, or in terms of the required odd lots of individual options.

Although this objection may have validity in the context of a proposed actual hedge to be placed against a particular deal, it does not carry much force in the context of risk management, in which hedging strategies are utilized as devices for representing risk in standard reports through liquid proxies. The tools for managing vanilla European options within a portfolio framework are well established. As was pointed out when discussing dynamic hedging in Section 11.3, good empirical evidence exists that vanilla options at less liquid strikes when statically hedged with vanilla options at more liquid strikes result in dynamic hedging strategies that achieve far greater stability than pure dynamic hedging strategies. As a result, we would argue that risk managers should not hesitate to represent exotic option trades as baskets of vanilla options in a vanilla options portfolio risk report. The advantages are parallel to those cited at the beginning of Section 10.2 for representing an illiquid forward as a static combination of liquid swaps: unified risk reporting increases risk transparency, maximizing liquidity and minimizing transaction costs.

The one point of legitimate concern would be if the resulting representation would be a position too large to be managed with the existing market liquidity. This would be an argument against representing a binary option as a very large position in a very narrow call spread. Instead, liquidity considerations should limit the size of the call spread position that is used as a representation, which in turn limits the narrowness of the call spread used. The resulting residual risks must be managed by the exotics desk through a combination of limits and reserves. We discuss this approach in more detail in Section 12.1.4. Another example would be if the representation revealed heavy reliance on very high- or low-strike vanilla options outside the range at which the firm's vanilla option traders would be comfortable managing the residual risk against more liquid strikes. Note that in both cases, the method of representing exotics exposure as a basket of vanilla options has the advantage of highlighting the regions of illiquidity impacting the exotic, a focus that many analytic pricing methods lack.

These points hold generally for the replication of exotic derivatives with vanilla options. By representing the exotic derivative as closely as possible with a hedge package of vanilla options, you can minimize the remaining basis risk that needs to be managed using techniques specific to the exotic derivative and maximize the amount of risk that can be combined with and managed as part of the vanilla options book, utilizing established risk management tools such as the price-vol matrix.

Examples of options that can be risk managed in this way are calls on the square, cube, square root, or other power of the excess above a strike, or the corresponding puts. Other mathematical functions, such as the logarithm

of the excess above or below a strike, are also possible. This style of option, sometimes collectively known as *power options*, has largely fallen out of favor following the Bankers Trust (BT)/Procter & Gamble (P&G)/Gibson Greetings blowup of 1994, which is discussed in Section 4.3.1. The lawsuits and allegations prompted by large losses on contracts with complex payoff formulas with no discernible tie to any of the end user's economic motives led to a distrust of such derivatives. Currently, most market makers' client appropriateness rules permit such contracts only in very limited circumstances.

Nonetheless, some power options remain in active use. The most prominent are *log contracts*, which are of particular interest because of their link to valuing and hedging variance swaps, and a type of quanto option that is utilized in the foreign exchange (FX) and bullion markets. In addition, the convexity adjustments needed for valuing and hedging certain types of forward risk, which we discuss in Section 10.2.4, can usefully be viewed as a type of power option and managed by this technique. After examining each of these three cases, we will follow with an examination of the important case of binary options, which illustrates the issue of how to handle liquidity risk arising from static replication. Finally, we will show how binary options can be combined with vanilla options to create other exotics—a contingent premium option and an accrual swap.

12.1.1 Log Contracts and Variance Swaps

A *variance swap* is a forward contract on annualized variance whose payout at expiry is:

$$(\sigma_p^2 - K_{\text{VAR}}) \times N \quad (12.1)$$

where σ_p^2 is the realized stock variance (quoted in annualized terms) over the life of the contract, K_{VAR} is the delivery price for variance, and N is the notional amount of the swap in dollars per annualized volatility point squared. The holder of a variance swap at expiry receives N dollars for every point by which the stock's realized variance, σ_p^2 , has exceeded the variance delivery price, K_{VAR} , and pays N dollars for every point by which the stock's realized variance, σ_p^2 , falls short of the variance delivery price, K_{VAR} . This contract can be generalized to assets other than stocks and to amounts other than dollars.

Variance swaps give their holders a vega exposure similar to what they would have by purchasing a vanilla option. However, variance swaps differ from vanilla options in that their vega exposure remains constant over time,

whereas vanilla options may go into or out of the money, reducing their vega exposure. This can be a significant advantage to a position taker whose main concern is to find an investment that expresses her economic view of future volatility. It also has the advantage of enabling her to avoid maintaining delta and gamma hedges, which will be seen as a distraction to the real intention, which is just to express a volatility view. The downside is the relative illiquidity of variance swaps versus vanilla options, leading to their being priced with wider bid-ask spreads. The log contract offers a means to link the hedging and valuation of the illiquid variance swap to that of liquid vanilla options, using the basket hedge methodology.

The link between the variance swap and the log contract comes from the following analytic formula for the value of a log contract:

$$\ln F - \frac{1}{2} \int_0^T \sigma^2 dT \quad (12.2)$$

where \ln is the natural logarithm function, F is the current price of the underlying forward to contract expiry T , and σ^2 is actual realized variance over that time period. This formula is a direct consequence of Equations 10 and 11 in Demeterfli et al. (1999). A derivation can also be found in Neuberger (1996). Under the Black-Scholes assumptions of known constant volatility, this implies that the log contract should be valued at $\ln F - \frac{1}{2} \sigma^2 T$, an analytic formula used in the **BasketHedge** spreadsheet to check the value derived for the log contract when the volatility surface is flat.

Since we can use the spreadsheet to find a set of vanilla options to replicate the log contract, we now have a hedging strategy for a variance swap. Buy a replicating set of vanilla options for twice the volume of log contracts as the volume of variance swaps sold (twice the volume in order to counteract the $\frac{1}{2}$ in front of the integral in the formula). Delta hedge these vanilla options. Since the log contract is losing value at exactly the rate of $\frac{1}{2} \int_0^T \sigma^2 dT$, the delta hedging should be producing profits at exactly the rate needed to cover payments on the variance swap.

In practice, this will not work exactly, due to jumps in underlying prices, as explained in Demeterfli et al. (1999, "Hedging Risks"). Monte Carlo simulation would be necessary to quantify the risk of this tracking error. However, the replication of the log contract still offers a good first-order hedge and valuation for the variance swap.

The section "The Difficulty with Volatility Contracts" in the same article discusses why this approach will not work for *volatility swaps*, which differ from variance swaps by having a payout of $(\sigma_p - K_{VOL}) \times N$ rather than $(\sigma_p^2 - K_{VAR}) \times N$. No static hedge for the volatility contract exists. In the

categorization we are using in this chapter, it is path dependent and needs to be risk managed using the techniques of Section 12.3, utilizing local volatility or stochastic volatility models to determine dynamic hedges. However, its close relationship to the variance swap, and thus to the log contract, suggests the use of a liquid proxy approach: use dynamic hedging just for the difference between the volatility swap and log contract while static hedging the log contract.

For further reading on the modeling and risk management of variance and volatility swaps, I highly recommend Demeterfli et al. (1999) and Gatheral (2006, Chapter 11).

Exercise 12.1 asks you to utilize the **BasketHedge** spreadsheet to look at the impact of changes in the volatility surface on the valuation of log contracts and hence on variance swaps. Demeterfli et al. (1999) also has an instructive section on the “Effects of the Volatility Skew” on variance swaps. Log contracts and variance swaps require hedges over a very wide range of strikes and should therefore show valuation sensitivity across the whole volatility surface. This seems reasonable from an intuitive standpoint since changes in volatility impact variance swaps even when the underlying forward price has moved very far away from the current price, leaving a currently at-the-money option very insensitive to vega. So high- and low-strike vanilla options are needed to retain the vega sensitivity of the package.

12.1.2 Single-Asset Quanto Options

In Section 12.4.5, we discuss dual-currency quanto derivatives in which the percentage change of an asset denominated in one currency is paid out in another currency. For example, a 10 percent increase in the yen price of a Japanese stock will be reflected by a 10 percent increase in a dollar payment at a fixed-in-advance dollar/yen exchange rate. We will see that the forward price of a quanto is the standard forward multiplied by $\exp(\rho\sigma_S\sigma_F)$, where \exp is the exponential to the base e , σ_S is the standard deviation of the asset price, σ_F is the standard deviation of the FX rate, and ρ is the correlation between them.

A related product is a single-currency quanto derivative in which the asset whose percentage change is to be calculated is also the asset whose exchange rate is fixed. Here are two examples:

1. A dollar/yen FX option, which, if the yen rises in value by 10 percent relative to the dollar, will be reflected by a 10 percent payout in yen. Since the yen has gone up in value by 10 percent versus the dollar, the payout in dollar terms is $110\% \times 10\% = 11\%$. In general, for a p percent increase, the payout is $(1 + p\%) \times p\% = p\% + p^2\%$.

2. A dollar/gold option struck at \$300 per ounce. If gold rises in value by 10 percent to \$330 per ounce, the payment is 1 ounce \times 10 percent = 0.1 ounces of gold. The payout in dollars is therefore $0.1 \times \$330 = \33 , which is $\$300 \times 11\%$. In general, for a p percent increase in gold prices, the payout is $p\% + p^2\%$.

Since just a single asset is involved, the σ_S and σ_F in the quanto formula are the same and ρ is equal to 1, so the standard forward is multiplied by $\exp(\sigma_S^2)$. The **BasketHedge** spreadsheet has a worksheet called **Quanto** that calculates the value of a single-asset quanto using a static hedge basket of vanilla options. As you can see from the spreadsheet, the hedge consists of 101 percent of a standard call at the quanto strike plus calls of 2 percent of the notional at all strike levels above the quanto strike. This gives a payoff, if the asset rises by p percent, of:

$$101\% \times p\% + 2\% \times \sum_{i=1}^{p-1} i = 101\% \times p\% + 2\% \times \frac{p^2 - p}{2} = p\% + p^2\%$$

When the static hedge cost is computed from a flat volatility surface, the results agree exactly with an analytic formula derived from the forward multiplied by $\exp(\sigma_S^2)$. If higher volatilities are assumed for higher strikes, the cost of the basket hedge will exceed the cost derived from the analytic formula. If lower volatilities are assumed for higher strikes, the cost of the basket hedge will be less than the cost derived from the analytic formula.

12.1.3 Convexity

In Section 10.2.4, on applying mathematical models of forward risk to indexed flows, we raised the issue of convexity or nonlinearity of some index flows and the complications this can entail for valuing and hedging these flows. We pointed out the availability of analytic formulas that approximate the convexity adjustments needed to account for the impact on valuation of the nonlinearity of these flows. These approximation formulas (see formulas 6.3, 29.1, 29.2, and 29.4 in Hull 2012) all require an interest rate volatility as a key input. However, in a world of nonflat volatility surfaces, which implied volatility should be used? Equivalently, what are the strikes of the options contracts that should be used to hedge this exposure?

The basket-hedging methodology we have developed in this section provides a more precise valuation for convexity adjustments, one that is sensitive to the shape of the volatility surface, and also provides details of the required hedge that can be used to represent the exposure in conventional vanilla option position reports, as shown in the **Convexity** worksheet within the **BasketHedge** spreadsheet.

TABLE 12.3 Payouts of a Binary Option

Final Price = S	Vanilla Call, Strike = K	Cash-or-Nothing Call, Strike = K , Payout = K	Asset-or-Nothing Call, Strike = K
$S \leq K$	0	0	0
$S > K$	$S - K$	K	S

12.1.4 Binary Options

European *binary options* (also known as *digital options* or *bet options*) have highly discontinuous payoffs. The basic form, the *cash-or-nothing option*, which we will focus on in this section, pays either zero if the price finishes below the strike or a set amount if the price finishes above the strike. A variant, the *asset-or-nothing option*, pays zero if the price finishes below the strike or the ending price if the price finishes above the strike. An asset-or-nothing option is simply the sum of a standard vanilla option and a cash-or-nothing option at the same strike that pays the strike price. Table 12.3 illustrates the payouts.

European binary options fulfill the condition of having a payout that is a function of the price of an asset at one definite time. Therefore, it can be treated by the methodology just stated, using a basket of vanilla options to hedge it and using this hedge package to calculate valuation, including skew impact, to calculate remaining risk, and to be incorporated into standard risk reports. However, the discontinuous nature of the payment at the strike leads either to unrealistically large hedge positions in vanilla calls (liquidity risk, since market prices would be impacted by an attempt to transact so many calls) or to significant hedge slippage (basis risk) between the binary option and its hedge.

For example, let's say a customer approaches a trading desk wanting to buy a one-year binary call that will pay \$10 million if the Standard & Poor's (S&P) index is above the current one-year forward level at the end of one year and nothing otherwise. The vanilla option decomposition of a barrier option is particularly simple. It can be represented as a call spread between two vanilla options of equal notional size. Assume you buy a vanilla call at a strike just below the current forward level and sell a vanilla option at a strike just above this level with a spread of 0.01 percent of the price between the two options. You will need to receive \$10 million if the index rises by 0.01 percent above the first strike, since for any index move above the second strike, you are paying as much on the second option as you are receiving on the first. So the notional amount of the call to be bought and sold is $\$10 \text{ million} / 0.01\% = \100 billion .

Let us start by assuming that all vanilla calls are priced at a 20 percent flat implied volatility. The straight analytical formula for the value of the binary option is the amount to be paid $\times N(d_2)$, the term in the Black-Scholes equation for a vanilla option that gives the risk-neutral probability that the price will finish above the strike. In this case, we have:

$$\begin{aligned} & \$10 \text{ million} \times N(d_2), \text{ where } d_2 = \left(\ln(\text{price} / \text{strike}) - \frac{1}{2}\sigma^2 t \right) / \sigma\sqrt{t} \\ & = \left(\ln(1) - \frac{1}{2}20\%^2 \right) / 20\% = 10\%N(-0.1) = 0.46017, \text{ giving a price of the} \\ & \text{binary of } \$4,601,700 \end{aligned} \tag{12.3}$$

Replicating the binary option using a vanilla call spread, the exact choice of vanilla calls to be used makes virtually no difference to the price (as long as we assume a flat implied volatility), but it does make a significant difference to the mix between liquidity risk and basis risk. For example:

- Buy a vanilla call on \$100 billion at a strike of 99.995% of the forward level at a price of $BS(99.995\%, 1, 20\%) = 7.9678802\%$ for \$100 billion $\times 7.9678802\% = \$7,967,880,200$ and sell a vanilla call on \$100 billion at a strike of 100.005% of the forward level at a price of $BS(100.005\%, 1, 20\%) = 7.9632785\%$ for \$7,963,278,500, for a net cost of $\$7,967,880,200 - \$7,963,278,500 = \$4,601,700$.
- Buy a vanilla call on \$2 billion at a strike of 99.75% of the forward level at a price of $BS(99.75\%, 1, 20\%) = 8.0812430\%$ for \$161,624,900 and sell a vanilla call on \$2 billion at a strike of 100.25% of the forward level at a price of $BS(100.25\%, 1, 20\%) = 7.8511554\%$ for \$157,023,100, for a net cost of \$4,601,800.
- Buy a vanilla call on \$500 million at a strike of 99% of the forward level at a price of $BS(99\%, 1, 20\%) = 8.4357198\%$ for \$42,178,600 and sell a vanilla call on \$500 million at a strike of 101% of the forward level at a price of $BS(101\%, 1, 20\%) = 7.5152765\%$ for \$37,576,400, for a net cost of \$4,602,200.

Note the inverse relationship between the width of the call spread (0.01 percent, 0.50 percent, and 2 percent, respectively) and the size of the legs of the call spread (\$100 billion, \$2 billion, and \$500 million, respectively).

The first combination offers the smallest basis risk. It will replicate the binary option exactly as long as the S&P index at the end of one year is outside the range 99.995% to 100.005%—that is, as long as the S&P index does not finish within about one-half basis point of its current forward level. However, liquidity risk is heavy; purchases and sales in the size of

\$100 billion would be certain to move market prices if they could be accomplished at all. (Even if the trading desk does not expect to actually buy this call spread, its use in representing the risk profile of the trade will lead to illiquid dynamic hedging requirements.) At the other end of the spectrum, the third combination is of a size that could possibly be transacted without major market movement, but basis risk is now much larger. Exact replication of the binary option takes place only in a range outside 99% to 101% of the current forward, so there are about 100 basis points of market movement on either side of the current forward level in which replication would be inexact. And replication could be *very* inexact. If the index ended at 100.1% of the forward, for example, the customer would be owed \$10 million, but the vanilla call at 99% would pay only $\$500 \text{ million} \times 1.1\% = \5.5 million , a net loss of \$4.5 million.

Of course, the basis risk can be dynamically hedged with purchases and sales of S&P futures. However, the large payment discontinuity of the binary option can lead to unmanageable hedging situations. For example, suppose you are close to expiration and the S&P is 1 basis point below the forward level. If no further movement occurs, you will make about \$4.95 million $[(99.99\% - 99\%) \times \$500 \text{ million}]$ on the vanilla call and owe nothing on the binary, but an uptick of just 2 basis points will lead to a loss of about \$5 million. Should you put on a delta hedge of a size that will make \$5 million for a 2-basis-point uptick? The problem is that a position of this size will cost you \$10 million for a 4-basis-point downtick, and you do not gain anything from option payouts to offset this loss. In theory, in a world of complete liquidity and no transaction costs, you could put on this hedge only at the exact moment you approach the binary strike and take it off as soon as you move away from that strike; but in practice, such strategies are wholly implausible. The actual experience of trading desks caught needing to delta hedge a sizable binary position that happens to be near the strike as expiration approaches is excruciatingly painful. Traders have their choice of gambles, but they must decide on a large bet in one direction or another.

In light of this, risk managers will always seek to place some sort of controls on binary positions. These controls, which may be complementary, come in the form of both limits and reserves. Limits are placed on the size of the loss that can occur for a certain size price move, the maximum delta position that can be required for a hedge, or the maximum gamma (the change in delta) that can be required for a given price move. Delta and gamma limits are based on the anticipated liquidity and transaction costs of the underlying market in which hedging is being done. Limits on loss size are designed to enable traders to take a purely insurance approach to binaries, hoping to come out ahead in the long run. This requires that

no one binary be too large. Such an approach needs to be combined with eliminating binaries close to a strike and expiration from delta and gamma reports, so that delta hedging is not attempted. It also requires decisions about how binaries should be combined for limit purposes.

To operate like insurance, binaries need to be widely scattered as to maturity date and strike level, and limits need to bucket strikes and maturities in a manner that forces this scattering. However, bucketing should combine binaries in only one direction (bought or sold); it is dangerous to permit the netting of one binary with another except when date and strike (and any other contract terms, such as a precise definition of the index) exactly match.

A valuation and reserve policy should also be consistent with the insurance approach to binaries—profit and loss (P&L) should be recognized only to the extent it can come close to being locked in. Gains that have great uncertainty attached to them should only be recognized when realized. This can be accomplished with several methods. I will provide a detailed example of a method that I consider particularly elegant in its capability to balance liquidity and basis risks, its maximal use of static hedge information, and its good fit with dynamic hedging risk reporting. In this approach, every binary has an internal liquid proxy representation assigned to it that is designed to be as close as possible to the binary in its payouts while still being capable of liquid hedging and conservative relative to the binary in that the internal representation will always produce a lower P&L for the firm than the binary. All risk reports for the firm are based on the internal representation, not the true representation of the binary. No special rules are required for eliminating binaries close to a strike and expiration from the firm's delta and gamma reports, since the internal representation has been designed to be small enough not to require unreasonable hedges. The valuation difference between the true and internal representation, which by design must always be a positive value to the firm, is booked to a reserve account. Since the reserve is always positive, this policy sometimes results in the firm recognizing windfall profits, but never windfall losses.

Let's see how this policy would work in the case we have been considering. A call spread is selected as the internal representation of the binary by choosing the smallest spread that results in a position size that is considered to be small enough to be liquid, either by representing a real possibility for purchase in the market or by being representable in the firm's risk reports by delta positions that can be achieved with reasonable liquidity. However, rather than choosing a call spread that straddles the binary, and therefore has payouts greater than the binary in some scenarios, we choose a call spread that is on one side of the binary and therefore always has payouts greater than the binary. If 2 percent is the width of the call spread

we select as the smallest consistent with a liquid position, then we use as an internal representation a call spread consisting of a sale of \$500 million at a strike of 98 percent and a purchase of \$500 million at a strike of 100 percent (notice that the internal representation has the opposite sign from the hedge that would extinguish it). The resulting valuation would be $\$500 \text{ million} \times BS(98\%, 1, 20\%) - \$500 \text{ million} \times BS(100\%, 1, 20\%) = \$500 \text{ million} \times 8.9259724\% - \$500 \text{ million} \times 7.9655791\% = \$44,629,900 - \$39,827,900 = \$4,802,000$. This is the valuation of the internal representation. The actual binary continues to be valued at \$4,601,700; the difference of \$200,300 is placed into a reserve. If the actual sale price of the binary to a customer is \$5 million, then only \$200,000 of the profit from the difference between the price and valuation goes into immediate P&L recognition; the other \$200,000 goes into a reserve against anticipated liquidity costs of managing the binary risk.

What happens to this reserve? There are several possibilities:

- The firm might decide to actually buy the static overhedge, which costs \$4,802,000. The internal hedge reports of the firm will not show the net position between the internal representation of the binary and the actual call spread hedge. If the S&P index ends up below 98 percent or above 100 percent, no difference will appear between the eventual payout under the binary and the pay-in due to the call spread, and the reserve will end up at zero. If the S&P index ends up between 98 and 100 percent, the call spread will have a pay-in while the binary has no payout. For example, if the S&P index ends at 99 percent, the call spread will pay \$5 million, which will be the final value of the reserve. At expiry of the options, this \$5 million will be recognized in P&L as a windfall gain.
- The firm might not do any static hedging and might just delta hedge based on the internal representation of the static overhedge. Since the static overhedge was selected to be of a size that enables liquid delta hedging, the results in this case should be close to the results in the case that the static overhedge is actually purchased, but with some relatively small variance. As an example, suppose that we are very close to expiry and the S&P index forward is at 99 percent. Based on the internal representation of the call spread overhedge, the appropriate delta will be a full \$500 million long in the S&P index forward, and roughly \$5 million in dynamic hedging profits should already have been realized but held in reserve. If the index ends at 99 percent, the \$5 million in dynamic hedging profits will be taken from the reserve and recognized in P&L as a windfall gain. If the index ends just above 100 percent, the \$5 million in dynamic hedging profits realized to date plus the

\$5 million gain from the 1 percent increase on the \$500 million long in the S&P index will be exactly enough to pay the \$10 million owed on the binary. Note that keeping the \$5 million in dynamic hedging profits realized to date in reserve is necessary to avoid having to reverse a previously recognized gain in order to pay off on the binary.

- Other combinations are possible, such as static hedges that are not overhedges, but all produce similar results.

In Exercise 12.2 you will run a Monte Carlo simulation of the potential differences between final payout on a portfolio of binary options and the overhedge liquid proxy, utilizing the spreadsheet **BinaryMC**. It will allow you to see a practical example of how a well-diversified portfolio of binaries requires lower reserves than a more concentrated portfolio of binaries.

This technique of representing a binary internally as a static overhedge is sometimes objected to by front-office personnel as trading off a very probable gain in order to achieve security. In this view, the \$400,000 that was originally realized on the transaction was real P&L, and \$200,000 was sacrificed in order to achieve security in the very small minority of cases in which the index finishes very close to the strike. The idea that \$200,000 has been thrown away is, in fact, an optical illusion caused by focusing only on those cases in which the index finishes outside the 99 to 101 percent range. The trade still has a \$400,000 expected value—it just consists of a sure \$200,000 in the vast majority of cases in which the index finishes outside 99% to 101% and a set of windfall profits up to \$10 million when the index finishes within this range. The front-office view would be correct if some means were available, such as dynamic hedging, of being almost sure of achieving this \$400,000 result in all cases. But it was exactly the lack of such means—the fact that the use of dynamic hedging to try coming close to achieving \$400,000 in all cases results in some cases with disastrous losses—that caused us to seek an alternative approach. This reserve methodology can be seen as being consistent with moving the front office away from viewing these trades as normal derivatives trades that can be approached in an isolated manner and toward viewing them as necessarily being part of a widely diversified portfolio of binaries. In this context, over a long enough time period, the sum of occasional windfall gains can become a steady source of income. If limits can ensure a wide enough diversification, then reserves may not be necessary.

So far in the example we have assumed a lack of volatility skew. In the presence of skew, the binary will price quite differently. Let's see the impact of using a 20.25 percent implied volatility for a strike of 99 percent and a 20 percent volatility for a strike of 101 percent. The cost of the 99 percent vanilla call is now $BS(99\%, 1, 20.25\%) = 8.534331\%$, resulting in a net cost

of \$5,095,274. Just as with the cases previously discussed, the reduction to a package of vanilla options lets us pick up the impact of volatility skew. We can see that binary options are highly sensitive to skew.

Taleb (1997, Chapter 17) gives a lucid discussion of the practical aspects of hedging binary options. On page 286, Taleb says that “the best replication for a binary is a wide risk reversal (that would include any protection against skew). There will be a trade-off between transaction costs and optimal hedges. The trader needs to shrink the difference between the strikes as time progresses until expiration, at a gradual pace. As such an optimal approach consumes transaction costs, there is a need for infrequent hedging.” Using a call spread (also known as a *risk reversal*) that is wide reduces the size of the vanilla options that are needed, reducing transaction costs and liquidity concerns, and also capturing the volatility skew more accurately, since a wide spread could utilize more liquid strikes. As we have seen, the width of the spread should not materially impact the total hedge cost.

In many cases, the underlying price will finish nowhere near the strike and no further transactions are needed. However, in those cases where the underlying is threatening to finish close to the strike, the basis risk will get too large and the trader will need to roll from the original call spread into a tighter call spread, incurring transaction costs due to the need to purchase and sell options and because the sizes of the option transactions are growing as the spread narrows. Factoring this potential transaction cost into the valuation of binary options is an alternative method for establishing a valuation reserve on a binary. As Taleb (1997, 286) states, “when the bet option is away from expiration, the real risks are the skew. As it nears expiration, the risks transfer to the pin. In practice, the skew is hedgeable, the pin is not.” (We have been using *basis risk* for what Taleb terms the *pin risk*.)

Gatheral (2006, Chapter 8) also has a good discussion of digital options, with a very clear demonstration of the dependence of digital option valuation on the skew of the volatility surface.

12.1.5 Contingent Premium Options

A contingent premium option entails no initial payment by the option buyer, who pays only at option termination under the circumstances that the option finishes in-the-money. This type of option is popular with some clients because of the deferral of cash payment and because the client will not need to pay for an option that turns out to be useless, although it should be noted that an option that finishes just slightly in-the-money will still require a net payment by the option buyer, since the payment due from the option seller will be less than the option’s cost. It is easy to see that a contingent premium option is just a standard vanilla option plus a forward to defer payment of

the option premium plus a binary option to offset the option premium due in the event the price finishes below the strike of the vanilla option.

12.1.6 Accrual Swaps

Accrual swaps are swaps where interest on one side accrues only when the reference rate is within a given range (see Hull 2012, Section 32.6). An accrual swap can be represented as a package of binary caps and floors since interest accruing is an all-or-nothing event. Being above the floor rate requires the payment and being above the cap rate cancels the payment, which can be represented by a payment with the opposite sign.

12.2 TIME-DEPENDENT OPTIONS

Now that we have provided a methodology for hedging and valuing the price of a linear underlying instrument at a single future point, we will extend that approach to exotic options whose payoffs depend on the price of a vanilla option at a single future point. This dependence on a vanilla option's future price can be decomposed into dependence on the price of the underlying of the vanilla option and dependence on its implied volatility. We can hedge the first element of this decomposition by a direct application of the methodology of the preceding section, leaving only dependence on implied volatility. To see how to hedge this piece, let us first look at an exotic that is dependent only on implied volatility and has no dependence on the underlying price.

12.2.1 Forward-Starting and Cliquet Options

A *forward-start option* is specifically constructed to have its price depend entirely on the at-the-money implied volatility of a vanilla option at a specified time. For example, a forward-start option could be sold on April 1, 2013, for a one-year at-the-money option to buy 1,000 shares of IBM that starts on November 1, 2013. The strike of the option will be set on November 1, 2013, at the then underlying price. Hence, no underlying price exposure exists prior to November 1, 2013, and the only exposure prior to that time is to what implied volatility the at-the-money option will sell at on November 1, 2013.

A *cliquet option* is a package of forward-start options, usually with one starting just as the previous one expires. For example, a cliquet might consist of three-month forward-start options beginning March 10, June 10, September 10, and December 10, 2013. Since the payoff on each option in

the package is determined independently of any other option in the package, a cliquet can be valued by valuing each forward-start option separately and then summing.

A natural approach would be to consider valuing a forward-start option with an extension of the method we used to roll into a longer-term option (see Section 11.6.3). The only difference is that we need to set up the target price-vol profile that we want to achieve as that of an at-the-money option, regardless of the underlying price level. The **ForwardStart** spreadsheet on the website for this book shows the details. The essential point is that the difference in the price of the at-the-money option at two different implied volatility levels, σ_1 and σ_2 , can just be represented as $BS(100\%, T, \sigma_1) - BS(100\%, T, \sigma_2)$, where T is the tenor of the at-the-money option to be created. Optimal fitting can then find the combination of current options that has close to the desired profile of volatility exposure at the time the forward-start option expires and the at-the-money option begins.

If you look at the example given in Table 12.4, you will find that the package of current options that creates the desired profile has a significant weighting at many different strike levels, so it will vary in valuation based on both the current smile and the current skew. This is not surprising, given that we are creating an option that has flat exposure to future implied volatility levels at all strikes. The situation parallels that of the log contract, which has flat exposure to variance.

12.2.2 Compound Options

It is now quite straightforward to extend this approach to exotics that depend on both underlying and implied volatility. A call-on-a-call option is one example of a *compound option*, which gives the purchaser of the compound option the right to buy (or sell) a particular vanilla option at a given strike price. It is also known as a *split-fee option* because a major selling point is that a customer who may want an option but is not willing to invest that much in one can put up a smaller down payment to defer the decision. Analytical formulas for compound options, assuming flat volatility surfaces and constant volatility, are well known (see Hull 2012, Section 25.6). We will make use of these formulas to work through an illustrative example.

Let's say that a customer wants to buy a one-year at-the-money call on 100 million euros on April 1, 2013, expiring on April 1, 2014. Assuming 20 percent implied volatility, the cost would be 7.97 percent of the principal amount. We'll assume the at-the-money euro exchange rate is \$0.90. The customer might prefer to pay 4.45 percent to get an option that can be exercised on November 1, 2013. On that date, the customer can either pay 5 percent to get a call on 100 million euros at a strike of \$0.90 expiring on

Discount Spacing Price	5.00%	Years		2	6.9829	-5.9652	2.7291	0.3272	-0.9075	2.2714	-5.8165	4.9924	-2.3706
		call	put										
Volume	-1	call	100	100	100	100	100	100	100	100	100	100	100
Price	5	100	80	100	90	100	110	120	80	90	100	110	120
Volatility	2%	1	2	2	2	2	2	2	1	1	1	1	1
Portfolio		20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%	20.0%
Implied vol	0.12%	-7.58%	-37.71%	103.69%	-60.70%	18.44%	1.43%	-18.29%	29.36%	29.36%	-44.07%	20.38%	-4.84%
BS price	30.3%	0.0%	-148.7%	486.3%	-331.8%	115.3%	10.1%	-80.6%	166.9%	166.9%	-314.0%	176.4%	-49.4%
Delta	0.12%	-0.38%	-0.61%	3.13%	-3.01%	1.37%	0.15%	-0.17%	0.71%	0.71%	-2.20%	1.77%	-0.66%
Vega	0.0%	0.0%	-1.7%	8.6%	-8.3%	3.8%	0.4%	-0.9%	3.7%	3.7%	-11.5%	9.3%	-3.4%
Gamma	0.001%	0.000%	0.011%	-0.060%	0.058%	-0.026%	-0.003%	0.006%	-0.027%	-0.027%	0.085%	-0.068%	0.025%
Theta	0.105%	-6.86%	-34.12%	93.82%	-54.93%	16.69%	1.29%	-16.55%	26.57%	26.57%	-39.88%	18.44%	-4.38%
Current													
Spot-Vol Matrix													
Price		-8%	-6%	-4%	-2%	0%	2%	4%	6%	8%	Vega	Vega	Convexity
	-25	0.52%	0.39%	0.28%	0.18%	0.07%	-0.03%	-0.12%	-0.22%	-0.31%	-0.05%	0.00%	0.00%
	-20	-0.14%	-0.06%	-0.04%	-0.04%	-0.07%	-0.10%	-0.14%	-0.18%	-0.22%	-0.01%	0.01%	0.01%
	-15	-0.30%	-0.18%	-0.13%	-0.10%	-0.09%	-0.08%	-0.08%	-0.09%	-0.10%	0.00%	0.01%	0.01%
	-10	-0.12%	-0.10%	-0.09%	-0.07%	-0.05%	-0.03%	-0.01%	0.01%	0.02%	0.01%	0.00%	0.00%
	-5	0.13%	0.02%	-0.03%	-0.03%	-0.01%	0.01%	0.05%	0.08%	0.11%	0.01%	-0.01%	-0.01%
	0	0.23%	0.06%	-0.01%	-0.02%	0.00%	0.04%	0.08%	0.12%	0.16%	0.01%	-0.02%	-0.02%
	5	0.16%	0.02%	-0.03%	-0.04%	-0.01%	0.03%	0.08%	0.13%	0.18%	0.02%	-0.01%	-0.01%
	10	0.02%	-0.04%	-0.07%	-0.07%	-0.04%	0.00%	0.05%	0.11%	0.16%	0.02%	-0.01%	-0.01%
	15	-0.03%	-0.06%	-0.09%	-0.09%	-0.07%	-0.03%	0.01%	0.06%	0.12%	0.01%	0.00%	0.00%
	20	0.07%	0.00%	-0.05%	-0.08%	-0.08%	-0.07%	-0.03%	0.01%	0.06%	0.00%	0.00%	0.00%
	25	0.32%	0.17%	0.05%	-0.02%	-0.07%	-0.08%	-0.08%	-0.05%	-0.01%	-0.01%	-0.01%	-0.01%

TABLE 12.4 Hedge at Rollover of a One-Year Option with a Forward Start in Two Years

April 1, 2014, or choose to let the option expire. The attraction to the customer is that if the euro declines in value by November 1, 2013, the option will seem unattractive and he will have saved money by having paid only 4.45 percent rather than 7.97 percent for the original option. The customer will pay more than 4.45 percent only if the option turns out to be valuable. Of course, the downside is that if he does want the option, he will have paid a total of $4.45\% + 5.00\% = 9.45\%$ for it rather than 7.97 percent.

When the call-on-a-call option expires on November 1, 2013, the value of the call option that the customer must now decide to purchase or let expire is determined by both the price of the underlying euro exchange rate (forward to April 4, 2014) and the implied volatility for a six-month option on the euro struck at \$0.90. The basket hedging procedure used in Section 12.1 can find a set of vanilla option hedges that eliminate the risk of the uncertainty of the underlying euro exchange rate. However, exposure to the uncertainty of the six-month implied volatility on November 1, 2013, will remain. This implied volatility exposure can be hedged by the same option roll approach as used in Section 12.2.1. The **Compound** worksheet of the **BasketHedge** spreadsheet calculates the vanilla option hedge against the underlying price and also calculates the price-vol matrix exposure of the resulting hedged position. This price-vol matrix can then be used as input to the **ForwardStartOption** spreadsheet to compute a hedge on the residual forward-starting volatility risk.

Exercise 12.1 takes you through pricing this call-on-a-call option in the **BasketHedge** spreadsheet. For a flat volatility surface, the basket hedge reproduces the analytical value, but different valuations are produced in the presence of smile and/or skew. Further steps in the exercise have you utilize the spreadsheet to calculate hedges and valuations for other compound options and choose options in which the decision on whether an option should be a call or a put can be deferred.

12.3 PATH-DEPENDENT OPTIONS

So far we've dealt strictly with exotic options whose payment is based on the price of an asset at a single time period—that is, European-style options. Now we want to look at how an option that is based on the prices of a single asset at many time periods can be handled. Barrier options are a good example to focus on for the following reasons:

- They illustrate dependence on the entire volatility surface, in terms of both time and strike level.
- They have a large range of variants.

- They are overwhelmingly the most traded exotic options among FX options and are also used with equities, commodities, and interest rates.
- They can be used as building blocks to form static hedges for other exotic options, such as lookback and ladder options.

A *barrier option* is one whose payoff is equal to that of a standard call or put, but that pays off only under the condition that some price level (called the *barrier*) has been breached (or not) at some time period prior to the time the call or put payoff is determined. Options that pay only if a barrier has been breached are called *knock-in* (*down and in* if the barrier is below the asset's price at the time the option is written, and *up and in* otherwise). For example, a one-year down-and-out call on the S&P index with a barrier of 1,050 will have no payout if the S&P index goes below 1,050 at any time during the year. Options that pay only if a barrier has not been breached are called *knock-out* (either *down and out* or *up and out*). Variations include *double barrier options* that either knock out if either a down-and-out or an up-and-out condition has been reached or knock in if either a down-and-in or up-and-in condition has been reached. Another variation is a *partial-time barrier*, where the barrier condition can be activated only during a specified time period that begins after the option start date and/or ends before the option termination date. A variation that can be combined with all of these options is a fixed rebate to be paid if an option is knocked out.

We will first show that standard analytic models for barrier options are inadequate, both for valuation and for risk representation, in the presence of nonflat volatility surfaces for vanilla options. We will therefore need to turn our attention to two alternative approaches to valuing and hedging barriers: dynamic hedging utilizing both vanilla options and the underlying and quasistatic hedging with vanilla options. One particular quasistatic hedging approach, developed by Peter Carr, is particularly useful for developing an intuitive understanding of the risk profile of barrier options. We will then demonstrate how to statically hedge lookback and ladder options with barrier options and how to handle rebates. Finally, we will briefly discuss how the methods developed for standard barrier options can be applied to the broader class of single-asset exotic options, including double barriers and partial-time barriers.

One noticeable difference between this section and all of our previous discussions of options is that we are concerned with the *drift*, which can be thought of either as the difference between the risk-free rate and the dividend rate, or more generally as the discount rate between forward prices at different expiries. Up until now, we didn't need to worry about drift because we were considering only options whose value would be determined by the

asset price at a single point in time; hence, all hedges could be based on a forward with a single expiry date. Since we are now considering options that depend on price behavior at several points in time, hedges may need to involve forwards for different expiry dates and the relationship between forward prices can no longer be ignored.

12.3.1 Standard Analytic Models for Barriers

Good analytic models based on partial differential equations (PDEs) have been developed for barrier options; see Hull (2012, Section 25.8) for the equations. Analytic models have great advantages in terms of computational speed relative to Monte Carlo and tree-based models. The ease of calculating a valuation by just plugging input variables into a formula explains much of the success of the Black-Scholes equation. The formulas for barrier options require a bit more computation than Black-Scholes, but they are still quite manageable. However, the analytic models for barriers have the drawback that they need to assume a single level of volatility, and there are no good rules for translating a volatility surface observed for European options into a single volatility to be used for a particular barrier option. In fact, cases can be shown where no single volatility assumption can be utilized with the standard analytic approach to give a reasonable price for the barrier option. We will illustrate this point with the following example. Consider an at-the-money three-month up-and-out call that knocks out at a barrier 20 percent above the strike. Its valuation at different volatility levels, using the standard analytic formula shown in Hull (2012, Section 25.8) is shown in Table 12.5.

Note that the analytic result has option values that first increase as the volatility level rises, since rising volatility causes the call value to increase. At higher volatility levels, the option values decrease as the volatility level rises, since rising volatility increases the probability of a knock-out. Since the barrier level starts far away from the current price, it is only at high volatilities that the impact of rising volatility on the probability of a knock-out dominates the impact of rising volatility on the value of the call.

The methods for utilizing the full volatility surface, which we will discuss shortly, would agree with these analytical results for flat volatility surfaces. However, if we assume a nonflat volatility surface, with an implied volatility of 20 percent for a European call struck at 100 and 18 percent for a European call struck at 120, approaches that utilize the full volatility surface (either the Derman-Kani dynamic hedging approach or the Carr static hedging approach) would price the barrier option at 3.10, which is 10 percent higher than the 2.81 maximum value the barrier option reaches at any volatility level using the analytic approach. The reason for this is

TABLE 12.5 Value of a Barrier Based on Analytic Formula

Volatility	Value of Up-and-Out Call
1.00%	0.1995
2.00%	0.3989
3.00%	0.5984
4.00%	0.7979
5.00%	0.9973
6.00%	1.1968
7.00%	1.3962
8.00%	1.5956
9.00%	1.7942
10.00%	1.9897
11.00%	2.1772
12.00%	2.3499
13.00%	2.5008
14.00%	2.6242
15.00%	2.7166
16.00%	2.7771
17.00%	2.8070
18.00%	2.8087
19.00%	2.7858
20.00%	2.7421
21.00%	2.6816
22.00%	2.6080
23.00%	2.5245
24.00%	2.4340
25.00%	2.3390
26.00%	2.2415
27.00%	2.1432
28.00%	2.0455
29.00%	1.9492
30.00%	1.8552

that the lower volatility as you approach the barrier decreases the chance of penetrating the barrier without simultaneously lowering the value of the call.

This example also shows why the analytic method is inadequate for representing the risk in standard option reports. The analytic method does not give any breakdown of how much of the risk should be represented as sensitive to changes in the at-the-money vanilla options versus how much should be represented as sensitive to changes in the out-of-the-money vanilla options.

12.3.2 Dynamic Hedging Models for Barriers

Dynamic hedging models price barrier options (or any other exotic option whose payoff is a function of a single underlying asset) based on the cost of dynamically hedging the exotic with a portfolio of the underlying asset and vanilla European options. This is analogous to the Black-Scholes model pricing of vanilla European options based on the cost of dynamically hedging with the underlying asset. These models utilize the full set of the current prices of vanilla European options, so they make use of the full volatility surface along with a theory of how these vanilla option prices can evolve with time. If you utilize an actual dynamic hedging strategy consistent with the model, you will be successful in replicating the model's price for the exotic to the extent that the model's theory about the evolution of the vanilla options prices is correct and that transaction costs are manageable.

Two principal types of dynamic hedging models are used for exotics:

1. **Local volatility models that assume that volatility is a known and unvarying function of time and the underlying price level.** These models are natural extensions of the Black-Scholes model, which assumes that volatility is known and unvarying, but which also assumes it is the same at all times and underlying price levels. Based on the assumption of the local volatility model, you can derive a definite price at any future time and the underlying price level of any vanilla or exotic option. The cost of the dynamic hedge therefore differs from the originally derived price only to the extent that future volatilities prove to follow a varying function of time and underlying price level (or that transaction costs are significant).
2. **Stochastic volatility models that assume that volatilities will vary over time and that might include price jumps, based on some assumed model.** The cost of the dynamic hedge differs from the derived price to the extent that the process of actual volatility variation differs from that assumed by the model (or to the extent that transaction costs are significant).

A relatively straightforward implementation for a local volatility model is the trinomial tree approach of Derman and Kani (1994), which builds the unique trinomial tree for modeling the price diffusion of the underlying asset that meets the following two criteria:

1. Volatility is a known and unvarying function of time and the underlying price level.
2. The tree correctly prices *all* European calls and puts on the underlying asset at different strike levels and times to expiry.

A thorough discussion of the Derman-Kani approach and its application to barrier pricing can be found in Chriss (1997, Chapters 9 and 11). If any reader wants to implement this model, I strongly recommend reading Chapter 5 of Clewlow and Strickland (1998), which provides wonderfully detailed instructions and examples.

A general introduction to stochastic models can be found in Derman and Kani (1998). A frequently used computationally tractable stochastic volatility model is that found in Heston (1993). A model that is attracting current interest is the *variance gamma model*, which is explained in Madan, Carr, and Chang (1998). Gatherall (2006) and Lee (2001) contain insightful analysis on the differences between local volatility and stochastic volatility models in the pricing of exotic options. Matytsin (1999) suggests that a combination of stochastic volatility and jump processes is needed to explain observed volatility surfaces implied by the vanilla option prices. The jump processes are needed to explain the steepness of smile and skew observed at shorter-term maturities, whereas stochastic volatility is needed to explain the steepness of smile and skew at longer-term maturities.

Dynamic hedging utilizes the full volatility surface in pricing barrier options. It can be readily employed for representing the barrier option in risk reports through its vanilla option hedges. Dynamic hedging can also be applied to any derivative based on a single underlying. Its drawback is its vulnerability to incorrect assumptions about volatility evolution and possible instability of the hedge representation.

The most thorough discussion of the vulnerability of dynamic hedging models to incorrect assumptions about volatility evolution that I know of is in Gatherall (2006), a relatively short book that is long on elegance and insight. At the close of Chapter 4, Gatherall states that “From the results of our computation, we can see that the local volatility model and the stochastic volatility model price European options almost identically” and that “to value an option, it’s not enough just to fit all the European option prices, we also need to assume some specific dynamics for the underlying.” In Chapter 8, Gatherall then analyzes the difference in evolution of the volatility surface implied by local volatility models versus stochastic volatility models. He states, “If the payoff we are hedging depends (directly or indirectly) on the volatility skew, and our assumption [which is implied by a local volatility model] is that the . . . skew is independent of the volatility level, we could end up losing a lot of money if that’s not how the market actually behaves.”

Once an exotic has been priced by a given model, the exotic can be hedged by a set of vanilla options that have the same sensitivity to the model’s input parameters as the exotic. As long as the model’s input parameters remain unchanged, the hedge does not require changing. However, changes

in observed vanilla option prices may require changes to input parameters to fit current prices, and once parameters change, the hedge may need adjustment.

How stable is the resulting representation? To what degree does it require frequent and sizable adjustments in the options hedges that can result in hedge slippage as a result of both transaction costs (generally considerably higher for options than for the underlying) and the instability of the hedge against parameter changes? The more the price of a product is dependent on assumptions about volatility evolution, the greater the instability of hedges. Although trading desks may gain experience with the stability of particular models in particular markets through time, it is difficult to obtain a risk measure in advance. The projection of hedge changes through Monte Carlo simulation (as recommended by Derman [2001] as discussed in Section 8.2.6.2), which has proved very useful in establishing results for the hedging of vanilla options with other vanilla options, is orders of magnitude more difficult to achieve for exotics. This is because each step on each path of the Monte Carlo simulation requires recomputation of the hedge. When the only hedge change is in the underlying, this is a very simple calculation of the $N(d_1)$ in the Black-Scholes formula. When the hedge change is in an option, a complete recalculation of the model being used to link the vanilla options and the exotic option together is required.

12.3.3 Static Hedging Models for Barriers

The uncertainty surrounding the hedging costs of using dynamic hedging for barriers provides the motivation to search for static or near-static hedging alternatives. Static hedging models price barrier options based on the cost of a replication strategy that calls for an almost unvarying hedge portfolio (at least of the vanilla options; it would be possible to use a dynamic hedge of the underlying, although the particular static hedging models we discuss only utilize vanilla options in the hedge portfolio). These models utilize nearly static hedge portfolios both as a way to reduce transaction costs and as a way to reduce dependence on assumptions about the evolution of volatility. Chapter 9 of Gatheral (2006) analyzes these nearly static hedges of barrier options from a different vantage point than mine, but with broadly similar conclusions.

Three approaches to the static hedging of barriers can be distinguished:

1. The approach of Derman, Ergener, and Kani, which is broadly applicable to all exotic options whose payoff is a function of a single underlying asset, but has considerable exposure to being wrong about future volatility levels.

2. The approach of Carr, which is more specifically tailored to barrier options, utilizing an analysis of the Black-Scholes formula to form a hedge portfolio that is immune to changes in overall volatility level and volatility smile. However, the Carr approach is still vulnerable to changes in the volatility skew. It is easier to implement than the Derman-Ergener-Kani approach for barriers in the absence of drift (that is, forward equal to spot) and produces a very simple hedging portfolio that helps develop intuitive understanding of the risk profile of the barrier.
3. Approaches that utilize optimal fitting give solutions close to those provided by the Carr approach for single barriers in the absence of drift, but are more flexible in handling drift and are less vulnerable to changes in volatility skew. Optimal fitting can be generalized to broader classes of exotics, but with less ease than the Derman-Ergener-Kani approach.

All three approaches are based on the idea of finding a basket of vanilla options that statically replicate the differences between the barrier option and a closely related vanilla option. To facilitate the discussion, we will confine ourselves to the case of a knock-out call, since a knock-in call can be handled as a vanilla call less a knock-out call, and all options can be treated as call options to exchange one asset for another (refer back to the introductory section of Chapter 11). The idea is to purchase a vanilla call with the same strike and expiration date as the knock-out being sold and then reduce the cost of creating the knock-out by selling a basket of vanilla options (this basket may have purchases as well as sales, but the net initial cash flow on the basket is positive to the barrier option seller). The basket of vanilla options must be constructed so that:

- It has no payoff if the barrier is never hit. In this case, the payout on the barrier option, which has not been knocked out, is exactly offset by the pay-in from the vanilla call that was purchased, so nothing is left over to make payments on the basket.
- Its value when the barrier is hit is an exact offset to the value of the vanilla call. When the barrier is hit, you know you will not need to make any payments on the barrier option, so you can afford now to sell the vanilla call you purchased. You do not want to later be vulnerable to payouts on the basket of vanilla options you sold, so you must purchase this basket. In order for cash flows to be zero, the basket purchase price must equal the vanilla call sale price.

You can guarantee the first condition by only using calls struck at or above the barrier in the case of a barrier higher than the current price and by only using puts struck at or below the barrier in the case of a barrier lower

than the current price. If the barrier is never hit, then you certainly won't be above the up barrier at expiration, so you won't owe anything on a call, and you certainly won't be below the down barrier at expiration, so you won't owe anything on a put.

All three static hedging techniques take advantage of knowing that at the time you are reversing your position in these vanilla options, the underlying must be at the barrier. A useful analogy can be made between these approaches to static hedging and the one we examined for forward-start options in Section 12.2. For forward-start options, we purchased an initial set of vanilla options and then had a fixed date on which we would make a single switch of selling our initial package of vanilla options and buying a new vanilla option. For barrier options, we cannot know in advance what the time of the switch will be, but we can know what the forward price of the underlying will be at the time of the switch. As with forward starts, we confine ourselves to one single switch out of the initial vanilla option hedge package. All of these approaches therefore share many of the advantages we saw for the static hedge technique for forward starts:

- A clear distinction between the portion of expected cost that can be locked in at current market prices of vanilla options (including current volatility surface shape) versus the portion that requires projections of what the volatility surface shape will be at the time of the switch.
- An estimate of uncertainty for establishing limits and reserves can be based on readily observable historical market data for possible volatility surface shapes. The impact of uncertainty is easy to calculate since it only needs to be computed at one particular point.
- Future liquidity costs, such as the potential payment of bid-ask spread, are confined to a single switch.
- Although it is to be expected that trading desks will, in practice, adjust the static hedge as market circumstances evolve, it remains useful as a risk management technique to evaluate the consequences of an unadjusted hedge.

The three approaches differ in how they attempt to ensure that the option package will be equal in value to the vanilla call at the time the barrier is hit. The Derman-Ergener-Kani approach (see Derman, Ergener, and Kani 1995) uses a package of vanilla options that expire at different times. The algorithm works backward, starting at a time close to the expiration of the barrier option. If the barrier is hit at this time, the only vanilla options still outstanding will be the vanilla call and the very last option to expire in the package. Since both the underlying price is known (namely, the barrier) and the time to expiry is known, the only remaining factor in determining the

values of the vanilla options is the implied volatility, which can be derived from a local or stochastic volatility model (if it is derived from a stochastic volatility model, it will be based on expected values over the probability distribution). Thus, the Derman-Ergener-Kani approach can be viewed as the static hedging analog of the dynamic hedging approaches we have been considering.

Once the prices of the vanilla options at the time the barrier is hit are calculated, you can easily determine the amount of the option that is part of the basket that needs to be sold in order to exactly offset the sale of the vanilla call with the purchase of the option in the basket. You then work backward time period by time period, calculating the values of all vanilla options if the barrier is hit at this time period and calculating the volume of the new option in the basket that is needed to set the price of the entire basket equal to the price of the vanilla call. At each stage, you only need to consider unexpired options, so you only need to consider options for which you have already computed the volumes held.

The following points about the Derman-Ergener-Kani approach should be noted:

- If the barrier is hit in between two time periods for which vanilla options have been included in the package, the results are approximated by the nearest prior time period. The inaccuracy of this approximation can be reduced as much as you want by increasing the number of time periods used.
- The approach can easily accommodate the existence of drift (dividend rate unequal to risk-free rate), since a separate computation is made for each time the barrier could potentially be hit.
- Since the approach relies on the results of a local or stochastic volatility model to forecast future volatility surface levels and shapes, it is vulnerable to the same issue as when these models are used for dynamic hedging—the hedge works only to the extent that the assumptions underlying the model prove to be true. As Derman, Ergener, and Kani state, “The hedge is only truly static if the yield curve, the dividend, and the volatility structures remain unchanged over time. Otherwise, the hedge must be readjusted.” This is illustrated in Table 12.6, which shows the potential mismatch in unwind cost at a period close to expiry based on differences between model-assumed volatilities and actual volatilities at the time the barrier is hit.

Note that the Derman-Ergener-Kani approach is vulnerable to model errors relating to both the level of volatility surface and the shape of volatility surface.

TABLE 12.6 Unwind Costs of Derman-Ergener-Kani Hedge of Barrier Option

Strike at-the-money, barrier at 95 percent of forward, and three months to expiry.
 Down-and-out call value at initial 20 percent volatility is 3.1955.
 Unwind with one month to expiry.

Volatility at Unwind	Unwind Gain or Loss
10.00%	0.4479
15.00%	0.2928
20.00%	0.0000
25.00%	-0.3595
30.00%	-0.7549

The Carr approach (see Carr, Ellis, and Gupta 1998) avoids this dependence on projecting future volatility surfaces and is much simpler to implement, but at a price—it cannot handle volatility skews (though it can handle volatility smiles) and its simplicity depends on the absence of drift (dividend rate equals risk-free rate).

The Carr approach achieves a degree of model independence by using a framework that corresponds directly with the Black-Scholes equation and determining a hedge package that will work, providing no drift or volatility skew is present. In these circumstances, one can calculate exactly a single vanilla put that will be selling at the same price as the vanilla call in the case that a down barrier is hit. It is based on the principle of put-call symmetry. In the boxes, we first explain how the principle of put-call symmetry can be derived from the Black-Scholes equation and then show how the exact Carr hedges can be derived from put-call symmetry.

PUT-CALL SYMMETRY

The principle of put-call symmetry says that if you have two strikes, K_1 and K_2 , whose geometric average is the forward price, that is, $\sqrt{K_1 K_2} = F$, then the current price of a call strike at K_1 for expiry T , $C(K_1, T)$, and the current price of a put struck at K_2 for the same expiry T , $P(K_2, T)$, are related by the equation:

$$C(K_1, T) / \sqrt{K_1} = P(K_2, T) / \sqrt{K_2}$$

This formula is a direct and easy consequence of the Black-Scholes formula. From Hull (2012, Section 17.8), the Black-Scholes formula for the price of a call and put based on the forward price is:

$$C(K_1, T) = e^{-rt} (FN((\ln(F/K_1) + \sigma^2 T / 2) / \sigma \sqrt{T}) - K_1 N((\ln(F/K_1) - \sigma^2 T / 2) / \sigma \sqrt{T}))$$

$$P(K_2, T) = e^{-rt} (K_2 N((\ln(K_2 / F) + \sigma^2 T / 2) / \sigma \sqrt{T}) - FN((\ln(K_2 / F) - \sigma^2 T / 2) / \sigma \sqrt{T}))$$

But since $F = \sqrt{K_1 K_2}$,

$$K_2 / F = K_2 / \sqrt{K_1 K_2} = \sqrt{K_2} / \sqrt{K_1} = \sqrt{K_1 K_2} / K_1 = F / K_1$$

So,

$$C(K_1, T) / \sqrt{K_1} = e^{-rt} (\sqrt{K_2} N((\ln(F/K_1) + \sigma^2 T / 2) / \sigma \sqrt{T}) - \sqrt{K_1} N((\ln(F/K_1) - \sigma^2 T / 2) / \sigma \sqrt{T}))$$

And substituting F/K_1 for K_2/F ,

$$\begin{aligned} P(K_2, T) / \sqrt{K_2} &= e^{-rt} (\sqrt{K_2} N((\ln(F/K_1) + \sigma^2 T / 2) / \sigma \sqrt{T}) \\ &\quad - \sqrt{K_1} N((\ln(F/K_1) - \sigma^2 T / 2) / \sigma \sqrt{T})) \\ &= C(K_1, T) / \sqrt{K_1} \end{aligned}$$

Since we have utilized the Black-Scholes formula in our derivation, this result holds only under the Black-Scholes assumption of a flat volatility surface for the expiry time T or if the deviation from flat volatility surface is exactly the same at strike K_1 and K_2 . However, since the forward is the geometric average of these two strikes, this is equivalent to saying that one strike is the same percentage above the forward as the percentage the other strike is below the forward. For their volatilities to be equal, the volatility surface must have a smile shape, not a skew shape, using the terminology of Section 11.6.2.

DERIVING THE CARR HEDGE

Since no drift is present, the forward price is equal to the spot price, which is the barrier level, H . Since the call is struck at K , we can find a reflection strike, R , such that $\sqrt{KR} = H$ and, by put-call symmetry, $\sqrt{R}Call(k) = \sqrt{K}Put(R)$. Since $\sqrt{KR} = H$, $R = H^2 / K$, $\sqrt{R} = H / \sqrt{K}$, you need to purchase $\frac{\sqrt{K}}{\sqrt{R}} = \frac{K}{H}$ puts struck at H^2 / K .

For an up barrier, one must separately hedge the intrinsic value and the time value of the vanilla call at the time the barrier is hit. The intrinsic value can almost be perfectly offset by selling binary options that pay $2 \times I$, where I is the intrinsic value. Any time the barrier is hit, there will be nearly a 50–50 chance that the binary will finish in-the-money, so its value is close to $50\% \times 2 \times I = I$. In fact, the standard lognormal pricing of a binary results in assuming slightly less than a 50 percent chance of finishing above the barrier, so we need to supplement the binary with I / H of a plain-vanilla call struck at the barrier. The exact value of the binary is $2 \times I \times N\left(-\frac{\sigma\sqrt{T}}{2}\right)$, and the value of the vanilla call struck at the barrier, and hence exactly at-the-money when the barrier is hit, is

$$\begin{aligned} & (I / H) \times H \times \left(N\left(\frac{\sigma\sqrt{T}}{2}\right) - N\left(-\frac{\sigma\sqrt{T}}{2}\right) \right) \\ & = I \times \left(1 - 2 \times N\left(-\frac{\sigma\sqrt{T}}{2}\right) \right) \end{aligned}$$

The sum of these two terms is then I .

The Carr approach has several advantages:

- It shows that it is at least plausible to price the barrier based on options with tenor equal to the final tenor of the barrier, indicating that this is probably where most of the barrier's risk exposure is coming from.
- Having a large binary component of the hedge is an excellent means of highlighting and isolating the pin risk contained in this barrier that dies in-the-money. Techniques we have already developed for managing pin risk on binaries can now easily be brought into play. For example, we could establish a reserve against the pin risk of the binary

(see Section 12.1.4). This approach is quite independent of whether the trading desk actually sells a binary as a part of the hedge—the risk of the binary is present in any case.

- Because the Carr approach uses a small number of options in the hedge package, it is very well suited for developing intuition about how changes in the shape of the volatility surface impact barrier prices.
- Even if you choose to hedge and price using a dynamic hedging approach, the Carr methodology can still be useful as a liquid proxy. Dynamic hedging can be employed for the difference between the barrier and the static hedge determined by the Carr approach. By choosing an initial hedge that, on theoretical grounds, we expect to be close to a good static hedge, we expect to minimize the degree to which changes in option hedges are required. However, by using dynamic hedging, we allow for as much protection as the accuracy of the model provides against uncertainty in skew and drift.
- Neither the presence of volatility smiles nor the uncertainty of future volatility smiles impacts the Carr approach. Since it deals with options that are symmetrically placed relative to the at-the-money strike, all smile effects cancel out.

The simplicity of the Carr approach is lost in the presence of drift or volatility skew. See the appendix to Carr and Chou (1996) for a method of using a large number of vanilla options to create a volatility-independent static hedge of barrier options in the presence of drift. See Carr (2001) for a method of handling volatility skew.

To appreciate how the Carr model performs and to gain the benefit of its insight into the risk structure of barriers, you should study the **CarrBarrier** spreadsheet provided on the website for this book. The spreadsheet shows the hedge structure for all eight possible simple barrier structures and the result of the barrier unwind for a specified scenario. Exercise 12.3 guides you through some sample runs. Here are some of the points you should be looking for:

- The one common element in all eight variants is the use of the reflection option—the one that utilizes the principle of put-call symmetry. It captures the time value of the barrier option at the point the barrier is hit.
- The sample run displayed in Table 12.7 shows that on unwind, for the down call and up put cases, the reflection option exactly offsets the value of the option that needs to be purchased for the in cases and needs to be sold for the out cases. For the up call and down put cases, a binary piece also needs to be offset, but the reflection option offsets the entire time value. In Table 12.8, in which the only change from Table 12.7 is

that the volatility at unwind has been raised, the binary piece (the sum of the binary and binary correction) is unchanged from Table 12.7, but the time value has increased exactly equally for the vanilla option and the reflection option.

- The time value when the barrier is hit depends on how far the barrier is from the strike. In the Table 12.7 example, the up barrier of 110 is further from the 100 strike than the 95 down barrier is, so the up reflection options have far less value than the down reflection options. You can think of the reflection option as taking value away from the out option and transferring it to the in option.
- The up call and down put cases are ones with binary components, since these in options will begin life already in-the-money and these out options cause an in-the-money component to be extinguished. The size of the binary component at the time the barrier is hit is the exact difference between the strike and barrier. It is divided into two pieces: the principal piece is the binary option and the secondary piece is the vanilla option used to supplement the binary. The total value of these two components at initiation will be less than the potential value on hitting the barrier, precisely reflecting the (risk-neutral) probability that the barrier will be hit.
- By trying different values for barrier-hitting scenarios, you will see that as long as volatility skew and drift are both equal to zero, the total impact of buys and sells in all eight cases is always zero. That is, the hedge works perfectly regardless of the assumptions made as to the time remaining when the barrier is hit, the at-the-money volatility, the volatility smile, or the risk-free rate. However, if either drift or volatility skew differs from zero, gains and losses will occur when the barrier is hit, varying by case. Examples are shown in Tables 12.9 and 12.10. It would clearly be a relatively easy task to calculate the size of potential losses based on assumptions about how adverse drift and skew could be at different possible times the barrier is hit. This could serve as input for the determination of reserves and limits.
- When the initial volatility skew, volatility smile, and drift are set equal to zero, pricing given by the standard analytic formula for barriers (shown on the top line in each column) exactly equals the total creation cost of the Carr hedges, as can be seen from the zero on the line labeled “difference.” When any of these values is different from zero, the Carr hedge gives a different value than the analytic formula. For example, Table 12.11 shows a case that corresponds to the one analyzed in Table 12.5, showing a 3.104 value for the up-and-out call in the presence of a volatility skew compared with a 2.7421 value using the analytic formula. Note that the presence of volatility skew (or drift) in the initial conditions does not imply that the Carr hedge will

not work. Only conditions at the time the barrier is hit determine the efficiency of the hedge.

In Exercise 12.4 you will run a Monte Carlo simulation of the cost of a hedging strategy that hedges a barrier option with the Carr hedge, utilizing the spreadsheet **CarrBarrierMC**.

A more general approach to static hedging that can handle all drift and volatility shape conditions is optimization, in which a set of vanilla options is chosen that fits as closely as possible the unwind of the barrier option at different possible times, drifts, volatility levels, and volatility surface shapes that may prevail when the barrier is hit. The optimization approach is discussed in Dembo (1994). Often no perfect static hedge can be found, but in these cases the optimization produces information on the distribution of possible hedge errors that can be useful input for determining a reasonable reserve. A similar approach can be taken to many different types of exotic structures.

The **OptBarrier** spreadsheet illustrates how optimization can be used to find a static hedge for a barrier option. If the possible conditions when the barrier is hit are restricted to zero drift and volatility smile but no skew, then the Excel Solver will find a set of vanilla options that almost exactly matches the barrier unwind for all volatility levels and times to expiry (although the particular set of hedges chosen may lack the clarity of insight that the Carr hedges offer). Of course, this is not a surprise since we know from the Carr approach that a perfect static hedge is possible under these circumstances. When different nonzero drift and volatility skew conditions are allowed, the match of the barrier unwind is no longer as exact.

The spreadsheet determines how much this slippage can be across all the specified cases of hitting time, skew, and drift. As with the Carr approach, this information can then be used to set reserves and limits. The difference from the Carr approach is the objective to find a hedge that minimizes the amount of this slippage. Exercise 12.3 guides you through some sample runs.

As a concluding note, observe that there is a lower limit on the uncertainty of unwind costs for any static hedging approach. Any dynamic hedging model can be used to compute the unwind cost of a selected static hedging strategy. So any difference in the pricing of barrier options between different dynamic hedging models translates into uncertainty of unwind costs. Practical experience with dynamic hedging models shows that differences in assumptions (for example, stochastic volatility versus local volatility and the frequency of jumps) give rise to substantial differences in barrier options prices utilizing the same input for current vanilla options prices. So you can search for static hedges that minimize the uncertainty of unwind costs, but an irreducible uncertainty will always remain that can be controlled only through limits and reserves. Static hedging greatly simplifies the calculations needed for limits and reserves.

Price	100.00	CDO	CDI	CUO	CUI	PDO	PDI	PUO	PUI
Strike	100.00	3.1955	0.7923	0.6343	3.3535	0.0778	3.9100	3.8791	0.1087
Up barrier	110.00	-3.9878	0	-3.9878	0	-3.9878	0	-3.9878	0
Down barrier	95.00	0	0	3.1581	-3.1581	3.2171	-3.2171	0	0
Time to expiry	0.25	0	0	0.0867	-0.0867	-0.0994	0.0994	0	0
Rate	0.00%	0.7923	-0.7923	0.1087	-0.1087	0.7923	-0.7923	0.1087	-0.1087
Drift	0.00%	-3.1955	-0.7923	-0.6343	-3.3535	-0.0778	-3.9100	-3.8791	-0.1087
ATM volatility	20.00%	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Vol smile	0.00%	90.25	90.25	121	121	90.25	90.25	121	121
Vol skew	0.00%								

		At barrier							
Time to expiry	0.25	-1.8881	1.8881	-10.9539	10.9539	-6.8881	6.8881	-0.9539	0.9539
Rate	0.00%	0	0	9.6012	-9.6012	5.1994	-5.1994	0	0
Drift	0.00%	0	0	0.3988	-0.3988	-0.1994	0.1994	0	0
ATM volatility	20.00%	1.8881	-1.8881	0.9539	-0.9539	1.8881	-1.8881	0.9539	-0.9539
Vol smile	0.00%	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Vol skew	0.00%								

Forward	100
Up forward	110
Down forward	95

TABLE 12.7 Carr Static Hedge

	Price	100.00	CDO	CDI	CUO	CUI	PDO	PDI	PUO	PUI
Strike	100.00		3.1955	0.7923	0.6343	3.3535	0.0778	3.9100	3.8791	0.1087
Up barrier	110.00		-3.9878	0	-3.9878	0	-3.9878	0	-3.9878	0
Down barrier	95.00		0	0	3.1581	-3.1581	3.2171	-3.2171	0	0
Time to expiry	0.25		0	0	0.0867	-0.0867	-0.0994	0.0994	0	0
Rate	0.00%		0.7923	-0.7923	0.1087	-0.1087	0.7923	-0.7923	0.1087	-0.1087
Drift	0.00%		-3.1955	-0.7923	-0.6343	-3.3535	-0.0778	-3.9100	-3.8791	-0.1087
ATM volatility	20.00%		0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Vol smile	0.00%		90.25	90.25	121	121	90.25	90.25	121	121
Vol skew	0.00%									
At barrier										
Time to expiry	0.25		-5.5195	5.5195	-14.2920	14.2920	-10.5195	10.5195	-4.2920	4.2920
Rate	0.00%		0	0	9.2034	-9.2034	5.3983	-5.3983	0	0
Drift	0.00%		0	0	0.7966	-0.7966	-0.3983	0.3983	0	0
ATM volatility	40.00%		5.5195	-5.5195	4.2920	-4.2920	5.5195	-5.5195	4.2920	-4.2920
Vol smile	0.00%		0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Vol skew	0.00%									

Forward	100
Up forward	110
Down forward	95

TABLE 12.8 Carr Static Hedge with Higher Volatility at Unwind

Price	100.00	CDO	CDI	CUO	CUI	PDO	PDI	PUO	PUI
Strike	100.00	3.1955	0.7923	0.6343	3.3535	0.0778	3.9100	3.8791	0.1087
Up barrier	110.00	-3.9878	0	-3.9878	0	-3.9878	0	-3.9878	0
Down barrier	95.00	0	0	3.1581	-3.1581	3.2171	-3.2171	0	0
Time to expiry	0.25	0	0	0.0867	-0.0867	-0.0994	0.0994	0	0
Rate	0.00%	0.7923	-0.7923	0.1087	-0.1087	0.7923	-0.7923	0.1087	-0.1087
Drift	0.00%	-3.1955	-0.7923	-0.6343	-3.3535	-0.0778	-3.9100	-3.8791	-0.1087
ATM volatility	20.00%	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Vol smile	0.00%	90.25	90.25	121	121	90.25	90.25	121	121
Vol skew	0.00%								

		At barrier								
Time to expiry	0.25	-1.9757	1.9757	-10.8303	10.8303	-6.9757	6.9757	-0.8303	0.8303	
Rate	0.00%	0	0	9.2028	-9.2028	5.0002	-5.0002	0	0	
Drift	0.00%	0	0	0.3988	-0.3988	-0.1994	0.1994	0	0	
ATM volatility	20.00%	1.8010	-1.8010	1.0830	-1.0830	1.8010	-1.8010	1.0830	-1.0830	
Vol smile	0.00%	-0.1746	0.1746	-0.1457	0.1457	-0.3739	0.3739	0.2527	-0.2527	
Vol skew	10.00%									

Forward	100
Up forward	110
Down forward	95

TABLE 12.9 Carr Static Hedge with Nonzero Skew at Unwind

	CDO	CDI	CUO	CUI	PDO	PDI	PUO	PUI
Price	100.00							
Strike	100.00							
Up barrier	110.00	0.7923	0.6343	3.3535	0.0778	3.9100	3.8791	0.1087
Down barrier	95.00	0	-3.9878	0	-3.9878	0	-3.9878	0
Time to expiry	0.25	0	3.1581	-3.1581	3.2171	-3.2171	0	0
Rate	0.00%	0	0.0867	-0.0867	-0.0994	0.0994	0	0
Drift	0.00%	0.7923	0.1087	-0.1087	0.7923	-0.7923	0.1087	-0.1087
ATM volatility	20.00%	-3.1955	-0.6343	-3.3535	-0.0778	-3.9100	-3.8791	-0.1087
Vol smile	0.00%	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Vol skew	0.00%	90.25	90.25	121	90.25	90.25	121	121
At barrier								
Time to expiry	0.25	1.6691	-10.2694	10.2694	-7.3790	7.3790	-1.0913	1.0913
Rate	0.00%	0	9.0052	-9.0052	5.4974	-5.4974	0	0
Drift	-3.00%	0	0.3610	-0.3610	-0.2179	0.2179	0	0
ATM volatility	20.00%	2.1120	0.8243	-0.8243	2.1120	-2.1120	0.8243	-0.8243
Vol smile	0.00%	0.4429	-0.4429	0.0789	0.0125	-0.0125	-0.2671	0.2671
Vol skew	0.00%							

Forward 100
Up forward 109.178086
Down forward 94.2901652

TABLE 12.10 Carr Static Hedge with Nonzero Drift at Unwind

	CDO	CDI	CUO	CUI	PDO	PDI	PUO	PUI
Price	100.00							
Strike	3.9244	0.0633	2.7421	1.2457	0.8479	3.1399	3.9874	0.0003
Up barrier	-3.9878	0	-3.9878	0	-3.9878	0	-3.9878	0
Down barrier	0	0	0.8708	-0.8708	3.7355	-3.7355	0	0
Time to expiry	0	0	0.0130	-0.0130	-0.0935	0.0935	0	0
Rate	0.1266	-0.1266	0.0000	0.0000	0.1266	-0.1266	0.0000	0.0000
Drift	-3.8611	-0.1266	-3.1040	-0.8838	-0.2192	-3.7686	-3.9878	0.0000
ATM volatility	0.0633	-0.0633	-0.3619	0.3619	0.6287	-0.6287	-0.0003	0.0003
Vol smile	81	81	144	144	81	81	144	144
Vol skew	-10.95%							
At barrier								
Time to expiry	-0.7124	0.7124	-20.1473	20.1473	-10.7124	10.7124	-0.1473	0.1473
Rate	0	0	19.2024	-19.2024	10.3988	-10.3988	0	0
Drift	0	0	0.7976	-0.7976	-0.3988	0.3988	0	0
ATM volatility	0.7124	-0.7124	0.1473	-0.1473	0.7124	-0.7124	0.1473	-0.1473
Vol smile	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Vol skew								

Forward	100
Up forward	120
Down forward	90

TABLE 12.11 Carr Static Hedge with Nonzero Skew at Initiation

12.3.4 Barrier Options with Rebates, Lookback, and Ladder Options

We will show how to use barrier options to create a static hedge for barrier options with rebates, lookback, and ladder options. Thus, we can transfer the techniques we have studied for using vanilla options to represent and hedge barrier option positions to create vanilla option representations and hedges of barrier options with rebates, lookback, and ladder options.

The use of a rebate feature in a barrier option can be regarded as a binary option triggered by a barrier. For example, suppose you have a down-and-out call that pays a rebate of \$2 million if the down barrier is hit and the call is canceled. This can be viewed as the sum of a down-and-out call with no rebate and a down-and-in binary option that pays \$2 million if the barrier is hit. However, since a binary option can be represented by being long one vanilla call and short another vanilla call, as discussed in Section 12.1.4, a down-and-in binary can also be treated as being long one down-and-in call and short another down-and-in call. So the rebate can be hedged and valued through the methodology we have already developed for barriers without rebates.

Lookback options come in two varieties: those that pay the difference between the maximum price that an asset achieves during a selected period and the closing price and those that pay the difference between the maximum price that an asset achieves during a selected period and a fixed strike. Symbolically, the lookback pays either $S_{\max} - S_T$ or $\max(0, S_{\max} - K)$. We can reproduce the payoffs of a lookback of the first type exactly by buying a lookback of the second type with a strike equal to the current price of the asset (S_0), selling the asset forward to time T , and buying a forward delivery of S_0 dollars at time T . Since S_{\max} is certainly $\geq S_0$, $\max(0, S_{\max} - S_0) = S_{\max} - S_0$, the total payoff of this combination at time T is:

$$\max(0, S_{\max} - S_0) - S_T + S_0 = (S_{\max} - S_0) - S_T + S_0 = S_{\max} - S_T \quad (12.4)$$

So if we can hedge the second type of lookback option by static hedging with barriers, we can create the first type of lookback option by static hedging with barriers as well.

Lookback options have a closely related product called *ladder options* that pay $\max(0, S_{\max} - K)$ rounded down by a specified increment. For example, if $K = 100$ and $S_{\max} = 117.3$, the lookback call of the second type would pay 17.3, a ladder with increments of 1 would pay 17, a ladder with increments of 5 would pay 15, and a ladder with increments of 10 would pay 10. Since a lookback call can be approximated as closely as we want by a ladder with a small enough increment, it is sufficient to show how to statically hedge a ladder with barriers.

It is easy to create a static hedge for a ladder option using up-and-in binary options. For each ladder rung, you buy an up-and-in binary option of the same tenor that pays the increment conditional on the rung being breached at some point during the life of the option. For example, if $K = 100$ and we have a ladder with increments of 5, we buy an up-and-in binary option having a payoff of 5 and a barrier of 105, another with a payoff of 5 and a barrier of 110, and so on. If the highest level the underlying reaches during the life of the ladder option is 12, then 10 will be owed on the ladder option, but the binary up-and-ins with barriers of 105 and 110 will both have been triggered for a payment of $5 + 5 = 10$.

12.3.5 Broader Classes of Path-Dependent Exotics

Now that we have looked at several dynamic hedging and static hedging alternatives for managing risk on standard barrier options, we want to examine how these approaches can be generalized to the full universe of single-asset exotics. We will focus most of our attention on double barriers and partial-time barriers, since these are reasonably popular products and since any techniques that are flexible enough to handle these variants would be flexible enough to handle any product.

Double barriers knock out (or knock in) if either a higher or a lower barrier is crossed. An example would be a one-year call option struck at 100 that knocks out if the price during the year is ever either above 120 or below 80. Partial-time barriers have a restricted time period during which the barrier provision applies. An example would be a one-year call option struck at 100 that knocks out if the price is below 90 any time between the end of month 3 and the end of month 9. If the price goes below 90 prior to month 3 but then goes back above 90 by the end of month 3, no knock-out occurs. Similarly, if the first time the price goes below 90 is after month 9, no knock-out occurs.

The greatest flexibility is offered by dynamic hedging, using either local volatility or stochastic volatility models, and by the Derman-Ergener-Kani approach to static hedging. Both can be easily generalized to double barriers and partial-time barriers. Local volatility models that solve for the exotic option values on a tree constructed to fit vanilla option prices can be easily adapted to solve for virtually any set of payoffs. Stochastic volatility models, which may require Monte Carlo simulation solutions, can easily handle any deterministic payout. The Derman-Ergener-Kani static hedging algorithm can solve for hedge packages that give zero unwind costs for double barriers and partial-time barriers just as easily as for standard barriers. The `DermanErgenerKaniDoubleBarrier` and `DermanErgenerKaniPartialBarrier` spreadsheets illustrate this computation. An interested reader could use

these spreadsheets as a guide to program a general calculator for applying the Derman-Ergener-Kani method to more complex barriers.

The drawbacks of dynamic hedging and Derman-Ergener-Kani static hedging that we analyzed for standard barriers apply in a more general setting as well. It will still be difficult to project the potential effects of hedge slippage for dynamic hedging. This is a heightened concern for double barriers since they have a reputation among exotics traders as particularly treacherous to dynamically hedge since they are almost always threatening to cross one barrier or the other. The dependence of Derman-Ergener-Kani on the model used to calculate the hedge ratios, and hence its vulnerability to being wrong about future volatility levels, remains true for the expanded product set.

Peter Carr and his collaborators have done a lot to expand the applicability of his static hedging approach beyond standard barriers. In particular, Carr, Ellis, and Gupta (1998, Section 3.1) have developed a static hedge for double barriers, and Carr and Chou (1997) have developed a static hedge for partial-time barriers. Similar results are presented in Andersen and Andreasen (2000). These hedges offer one of the major advantages of the Carr hedge for standard barriers—protection against shifts in volatility levels. However, they do not offer another major advantage of the Carr hedge for standard barriers: They are not simple to compute and do not provide much intuitive insight into the risk structure of the exotic being hedged. The specialized nature of each construction does not offer significant guidance as to how to build hedges for other exotics.

Optimal fitting would seem to offer the best hope for an easy-to-generalize static hedge that will minimize sensitivity to model assumptions. However, unlike the Derman-Ergener-Kani method, which automates the selection of the vanilla options to be used in hedging a particular exotic, the optimal fitting approach relies on practitioner insight to generate a good set of hedge candidates. A poor choice of possible hedges results in a poorly performing static hedge. A possible solution is to try to generalize the Derman-Ergener-Kani approach to fit to a range of volatility surfaces rather than to a single one. Some promising results along these lines have been obtained by Allen and Padovani (2002, Section 6). A copy of this paper is on the book website.

12.4 CORRELATION-DEPENDENT OPTIONS

Valuation and hedging strategies for derivatives whose payoff is a function of more than one underlying asset are critically dependent on assumptions about correlation between the underlying assets. With only a few exceptions

(which are discussed in Section 12.4.3), there is an absence of sufficiently liquid market prices to enable implied correlations to be inferred in the way implied volatilities can be derived from reasonably liquid prices of vanilla options. So much of the focus of risk management for these derivatives revolves around controlling the degree of exposure to correlation assumptions and building reserves and limits against the differences between actual realized and estimated correlations.

An important distinction within derivatives with multiasset payoffs should be made between those whose payoff is based on a linear combination of asset prices (for example, the average of a set of prices or the difference between two prices) and those whose payoff is based on a nonlinear combination of asset prices (for example, the maximum of a set of prices or the product of two prices). When the payoff is based on a linear combination of asset prices, risk management is considerably simpler, even if the payoff itself is a nonlinear function of the linear combination of asset prices, such as an option on the average of a set of prices. We therefore discuss these two types of derivatives in separate sections. A final section discusses options that depend on a different type of correlation—the correlation between underlying asset value and the probability of option exercise.

12.4.1 Linear Combinations of Asset Prices

Derivatives whose payoff depends on a linear combination of asset prices share several important characteristics that simplify their risk management:

- If the payoff function is a linear function of the linear combination of asset prices, then the derivative does not have any option characteristics and can be perfectly hedged with a static portfolio of the underlying assets. In such cases, the valuation of the derivative is independent of correlation assumptions. This is not true of derivatives whose payoff function is a linear function of a nonlinear combination of asset prices, such as a forward based on the product of an asset price and an FX rate (a so-called quanto) that requires dynamic hedging.
- Even when the payoff function is a nonlinear function of the linear combination of asset prices, such as an option on the average of a set of prices, and therefore requires dynamic hedging, the rules for dynamic hedging are particularly simple to calculate.
- Even when dynamic hedging is required, it is often possible to make very good approximations of valuation and the risk of incorrect correlation assumptions using a standard Black-Scholes model.

We will examine each of these characteristics more closely. We will then make use of the approximation technique discussed previously to answer questions about how the risk of these derivatives should be managed.

12.4.1.1 Derivatives Whose Payoffs Are Linear Functions of Linear Combinations of Asset Prices

In principle, any derivative whose payoff is a linear function of a linear combination of asset prices, such as a forward on the average price of a basket of assets, can be statically hedged by buying the properly weighted basket of forwards. In practice, this could be operationally difficult for a basket composed of a very large number of assets, and a market maker may choose to hedge with a differently weighted basket selected to statistically track the derivative payoff closely, with a resulting possibility of tracking error. However, in either case, the performance of this hedging strategy will not be influenced by the level of correlations of assets within the basket. In particular, the valuation of a basket should not be influenced by whether the assets in the basket are well diversified or highly concentrated. Both well-diversified and highly concentrated baskets should be valued as the weighted average of the valuations of the individual components.

At first, this may seem to violate intuition, since firms devote considerable resources to calculations such as value at risk (VaR) that rate highly concentrated baskets as riskier than well-diversified baskets. Shouldn't some penalty in valuation be applied for an asset basket that carries more risk? The answer from capital market theory is that only systemic risk, which is not capable of being diversified away, should be penalized and that the role of tools such as VaR is to make certain that a firm has considered the proper hedges against risk that can be diversified away. So a trader entering into a forward on the average price of a basket will be charged a higher risk premium by his firm's risk systems for running an open position (that is, not putting in place the basket hedge) in a highly concentrated basket than in a well-diversified basket. But in either case, he has the ability to put on the hedge closing out the position, so concentration should only play a role in the evaluation of the risk of running an open position, not in the valuation of the derivative. A particularly clear discussion of this point can be found in Varian (1987, "Value Additivity Theorem").

As Varian emphasizes, this principle only applies as long as payoffs are linear and ceases to apply when payoffs are nonlinear. This is true both for nonlinearity of the payoff function, such as an option on the average price of a basket of stocks, and the nonlinearity of a combination of asset prices, such as a forward on the maximum price of a set of stocks. As soon as nonlinearity is introduced, considerations that only play a minor role in the risk assessment of linear products begin to play a role in valuation. For example, the probability of extreme tail events based on the correlation of

default probabilities plays no role in the valuation of a CDO based on a basket of loans and/or bonds so long as the CDO divides ownership of the basket proportionally. (A CDO is an example of an asset-backed security; see Section 10.1.8.) However, CDOs often divide the ownership of the basket into tranches, with some tranches paying all credit losses up to a certain level and other tranches paying only losses above that level. This enables the investor market to be segregated more efficiently by creating some bonds that are tailored to investors seeking lower credit risk and other bonds that are tailored to investors willing to take on more credit risk in return for adequate compensation. Tranching CDOs introduces nonlinearity of payoffs. As a result, valuation is dependent on the probability of extreme tail events based on the correlation of default probabilities. For further discussion of this point, see Section 13.4.1.

A second point to note is that the arbitrage principle only applies if the assets comprising the basket are sufficiently liquid. If not, investors who would have a hard time acquiring a diversified basket of assets may be willing to pay a premium to receive a payment on an index based on the average price of such a basket. This offers a profit opportunity to market makers who can efficiently acquire diverse baskets that other market participants would find difficult to replicate. The market maker can then offer to pay an index based on its earnings on the basket and build a premium into the index. This diversification premium has definitely been observed in the default swaps market.

12.4.1.2 Rules for Dynamic Hedging The required dynamic hedges for an option on a linear combination of asset prices are very easy to determine. Standard deltas can be derived from option pricing models, and the delta hedge can then be formed by multiplying this delta times the linear weights of each asset in the basket. This simplifies ongoing hedging calculations and the calculation of required hedges in Monte Carlo simulations of hedging strategies.

Consider an at-the-money one-year option on a 5,000-share stock basket consisting of 20 percent IBM, 45 percent General Electric (GE), and 35 percent Merck. If the volatility of the basket is assumed to be 25 percent, the delta, using the Black-Scholes formula, is 55 percent. The hedge should be:

$$\begin{aligned} 5,000 \times 55\% \times 20\% &= 550 \text{ shares of IBM} \\ 5,000 \times 55\% \times 45\% &= 1,237.5 \text{ shares of GE} \\ 5,000 \times 55\% \times 35\% &= 962.5 \text{ shares of Merck} \end{aligned} \quad (12.5)$$

12.4.1.3 Approximation of Option Values The calculation of the value of an option on a linear combination of asset prices can be reasonably approximated by calculating the volatility of the underlying basket based on the weights

of each asset in the basket, the implied volatilities of each asset, and the assumed correlations between assets. This calculated volatility can then be used as input to the Black-Scholes formula for the basket option.

Continuing the previous example, assume that the volatility of IBM stock is 30 percent, the volatility of GE stock is 33 percent, and the volatility of Merck stock is 28 percent, with correlations between IBM and GE of 60 percent, between IBM and Merck of 50 percent, and between GE and Merck of 40 percent. Then the volatility of the basket can be estimated as:

$$\begin{aligned} & \text{SquareRoot} [(20\% \times 30\%)^2 + (45\% \times 33\%)^2 + (35\% \times 28\%)^2 \\ & + 2 \times (20\% \times 30\% \times 45\% \times 33\% \times 60\% + 20\% \times 30\% \\ & \times 35\% \times 28\% \times 50\% + 45\% \times 33\% \times 35\% \times 28\% \times 40\%)] \\ & = 25.2\% \end{aligned} \quad (12.6)$$

This is only an approximation for two reasons. The first reason is that the representation of an asset's distribution by a single implied volatility is only accurate if the implied volatility surface for that option is flat, that is, the same at all strike prices. However, as discussed in Section 11.6.2, this is rarely the case. The second reason is that even if we had an example in which the implied volatility surfaces of the options on all the individual assets were flat, meaning that the market was pricing them all as if they were lognormally distributed, a linear combination of lognormal distributions is not lognormal, so the implied volatility surface for the basket option would not be flat and thus could not be represented by a single volatility.

For assets with reasonably flat implied volatility surfaces, this approximation technique will give accurate enough results to be useful as a way of building intuition about the degree to which basket option prices depend on the implied volatilities of the individual assets and on the assumed correlations between them. This is how we will make use of this approximation in the remainder of this section.

Actual valuations require more accurate numerical techniques. In practice, two are generally used. One technique is a Monte Carlo simulation in which each asset process is specified by a full distribution that corresponds to the implied volatility surface for that asset, following the approach discussed in Section 12.3.2. Assumed correlations between assets can be enforced by the technique discussed in Hull (2012, Section 26.7). This technique is flexible enough to support more complex assumptions, such as correlations that vary based on the price level or price movement of the component assets. Finally, the value of the basket can be computed along each sample path and the resulting value of the option can be calculated.

The flexibility to have correlation vary with price level or price movement can be important since large downward price moves tend to be

accompanied by higher correlation than ordinary price moves. This can result in baskets being priced at higher volatility skews than individual components of the basket since it increases correlation and hence increases volatility at lower price levels. For further discussion of this point, see Derman and Zou (2001).

The Monte Carlo approach affords great flexibility, including the incorporation of stochastic volatility and price jump assumptions. Its drawback is difficulty in valuing American-style options that require the determination of optimal early exercise strategies. Further developments in Monte Carlo modeling do allow approximations of American option valuation; see, for example, Broadie, Glasserman, and Jain (1997).

The alternative approach for American-style options on baskets is the three-dimensional tree approach described in Hull and White (1994). This approach enables the combination of two trinomial trees that have been fitted to the full implied volatility surface, using the techniques discussed in Section 12.3.2, to be combined into a single tree based on assumed correlations, which can vary by node. Basket values can then be computed on the combined tree and option values determined by working backwards on the tree. This approach has the advantage of greater precision in determining early exercise strategies. The disadvantages are that it is only computationally feasible for baskets involving two assets and it is restricted to using local volatility models to replicate the implied volatility surface, which lacks the flexibility to incorporate stochastic volatility or price jumps. A possible combination of the two methods for more than two assets would be to determine the option price for the final exercise using the more precise Monte Carlo method and estimating the extra value due to possible early exercise using the three-dimensional tree technique using the first two principal components of the assets as the two variables to be modeled on the tree.

12.4.2 Risk Management of Options on Linear Combinations

We will now take advantage of the simple formula available to approximate the value of an option on a linear combination of assets to examine how risks arising from positions in these options should be managed.

One possible risk management technique is pure dynamic hedging of options positions in a particular linear combination. This is operationally straightforward, as discussed in Section 12.4.1.2. However, it encounters the same deficiencies of reliance on the delta-hedging strategy that we discussed in Section 11.1. The same arguments favoring the use of other options in hedging that were given in Section 11.1 apply, but it is unusual to find any liquidity in options on asset combinations. This suggests the use of options on individual assets comprising the basket as part of the hedge.

TABLE 12.12 The Impact of Hedging Basket Options with Single-Stock Options

	Standard Deviation	Transaction Costs
Dynamically hedge with underlying stocks only	28.7%	2.3%
Purchase at-the-money options on stocks A and B and then dynamically hedge	14.0%	1.9%

Consider the following simple example. An option has been written on the average of two assets, A and B. Compare the simulation results of a pure dynamic hedge with the underlying stocks with the simulation results of a hedge that involves first purchasing options on assets A and B and then dynamically hedging the resulting position with the underlying stocks.

Suppose a one-year at-the-money option has been written on the average of the prices of two stocks, A and B. Assume that both A and B have 20 percent volatility on average with a 33 percent standard deviation of volatility and that correlation between the two assets averages 0 percent with a 33 percent standard deviation. We will simulate two hedging strategies: Use a pure dynamic hedge with the underlying stocks, or first purchase an at-the-money option on A and an at-the-money option on B and then dynamically hedge the resulting position with the underlying stocks. The ratio of the notional of purchased options on individual stocks to the notional of the sold basket option we will use is 70 percent, split equally between the option on A and the option on B. This 70 percent ratio is suggested by the average volatility of the basket option being $\sqrt{(50\% \times 20\%)^2 + (50\% \times 20\%)^2} = 14.14\%$, which is just a little bit more than 70 percent of the 20 percent average volatility of the individual stocks. Simulation starting with different ratios of individual stock options to the basket options confirms that 70 percent is the ratio that results in the lowest standard deviation of the dynamic hedging results. Table 12.12 compares the results between the two hedging strategies.

Although a substantial reduction in uncertainty and transaction costs results from utilizing an option in the constituent stocks as a hedge, it is not as large a reduction as was shown for hedging vanilla options with vanilla options at other strikes in Table 11.2. Even if we were certain of the correlation, the static hedge utilizing the purchase of at-the-money options on stocks A and B can only reduce the standard deviation to 12.2 percent. The intuitive reason for this is that the relationship of one strike being located midway between two other strikes is obviously stable, whereas the underlying stock options can move into or out of the money without a

similar move on the part of the basket option. For example, if stock A's price rises by 20 percent and stock B's price falls by 20 percent, the previously at-the-money call options on stock A and B will now be substantially in-the-money and out-of-the-money, respectively. In both cases, their sensitivity to volatility will be considerably reduced from the time of initiation. This is not true for the basket option, which will still have its same initial sensitivity to volatility since it is still at-the-money relative to the average price of A and B.

A possible remedy would be to dynamically change the amount of single stock options being used to hedge in response to changes in relative volatility sensitivity of the basket option and single stock options. This has many similar virtues and drawbacks with the proposal to dynamically hedge barrier options with vanilla options that was considered in Section 12.3.2. One advantage in this case is that it is considerably easier to calculate the required option hedges in the Monte Carlo simulation, provided you are willing to accept the degree of approximation of the simple formula.

Whether employing static hedging or dynamic hedging with single-asset options, the following rules should apply:

- Any residual exposure to the uncertainty of correlation should be reflected in reserve policies and limits, since this is an exposure that cannot be hedged with liquid instruments.
- Residual unhedgeable exposure to the uncertainty of single-asset volatility should be quantified, as shown in the Monte Carlo example in Table 12.12, and reflected in reserve policies and limits.
- Valuation procedures and risk measurement should be in agreement. If implied volatilities of individual assets are used as an input to the valuation of a basket option, then the exposure to changes in each constituent asset's implied volatility should be reflected, either statically or dynamically, in price-vol matrix reports and other volatility exposure measures computed for the individual asset. Similarly, delta exposure should be reflected in individual underlying asset position reports. If this principle is not followed, valuation exposure to changes in the price or volatility of an asset can grow without control by being included in more and more basket products.
- In some cases, individual asset volatility may be so slight a contribution to the risk of a basket option that it is not worth the effort of utilizing the implied volatility as an input to valuation or reflecting exposure to volatility changes in individual asset risk reports. The basket option will then effectively be managed as if it was an option on a separate underlying unrelated to the single-asset options. Note that this does not change the use of the individual underlying to perform delta hedging.

TABLE 12.13 Sensitivities of Option on Basket

Correlation Level	1% Shift in Volatilities	10% Shift in Correlation
90%	0.97%	0.51%
75%	0.94%	0.53%
50%	0.87%	0.57%
25%	0.79%	0.62%
0%	0.71%	0.69%
-25%	0.61%	0.79%
-50%	0.50%	0.95%
-75%	0.35%	1.30%
-90%	0.22%	1.85%
-95%	0.16%	2.31%
-98%	0.10%	2.90%

The **BasketOption** spreadsheet on the website for this book shows the calculation of basket option exposures to changes in correlation and individual asset volatility under the approximation of the simple formula. Table 12.13 shows some sample results for an equally weighted two-asset basket with both assets having a 20 percent volatility. The impacts shown are for a 1 percent shift in the volatilities of both assets (for example, 20% + 1% = 21%) and a 10 percent shift in correlation (for example, 75% + 10% = 85%).

Note how the relative contribution of individual stock volatility relative to correlation declines sharply as correlation levels become negative. This is very relevant for options on the spread between two asset prices, since the hedge basket then consists of a positive position in one asset and a negative position in the other. If the assets are strongly correlated, their positions in the basket will show high negative correlation. In these cases, hedging the individual option volatilities is questionable.

One reporting issue for all multiasset derivatives is whether to take correlation into account when reporting delta and vega exposure of the derivative. As a concrete example, consider a forward on the average of two stocks, A and B, whose prices are 90 percent correlated. If the overall basket position has an exposure of \$1 million for a 10 percent rise in the average price, should you show the exposure to A as \$500,000 or as something closer to \$1 million to reflect the probability that a rise in the price of A will be accompanied by a rise in the price of B? Clearly, for purposes of the firm's consolidated risk-management reports, \$500,000 is the right figure since the consolidated reports will also be showing a \$500,000 exposure

to B and these two positions will contribute to the consolidated reporting of total exposure to a 10 percent increase in stock prices. If you used a position closer to \$1 million for the A exposure, it would have the absurd result, when combined with exposure to B, of showing an exposure greater than \$1 million to a 10 percent increase in stock prices. However, including a correlation may be appropriate for specially tailored reports for traders who want a quick rule of thumb about how much the basket price will move when stock A's price moves (perhaps because A's price is more liquid than B's). A particular example that has attracted industry attention is the sensitivity of convertible bond prices to changes in the underlying stock price, which we discuss further in Section 12.4.4.

A particular example of a basket option is an Asian option on a single asset. An Asian option is an option on the average price of the asset over a specified set of observations. This is equivalent to an option on a basket of forwards where all the forwards are for the same underlying asset. Obviously, one would expect correlations on such forwards to be quite high. In fact, the conventional Asian option pricing formula assumes a correlation of 100 percent (see Hull 2012, Section 25.12), which is equivalent to assuming constant interest rates, which is slightly inaccurate. Note that the time period over which each forward will contribute volatility to the basket is different, which is a key element to be taken into account in the pricing of the option.

12.4.3 Index Options

As a generalization, we have stated that most multiasset derivatives are illiquid. But this rule has clear exceptions—most prominently, options on interest rate swaps and options on equity indexes. Options on interest rate swaps, also known as *swaptions*, are mathematically and financially equivalent to options on a basket of forwards so they reflect an implied correlation. This special case is treated at length in Section 12.5. Options on stock indexes, such as the S&P 500, NASDAQ, FTSE, and Nikkei, are among the most widely traded of all options. Comparing implied volatilities of stock index options with implied volatilities of options on single stocks that are constituents of the index will therefore yield implied correlation levels. We look at the risk management consequences, which can also be applied to other liquid index options such as options on commodity baskets and FX baskets.

The first principle is that the valuation of a reasonably liquid index option should always be directly based on market prices for the index option and not derived from prices for options on individual stocks in the index and a correlation assumption. Correlation assumptions, no matter how

well based in historical analysis and economic reasoning, should never be allowed to replace a market-derived implied correlation to assess the price at which risk can be exited. This is just an application of the same reasoning that says that reasonably liquid options need to be valued using implied volatilities, not volatility assumptions based on history.

This does not mean that room is not available for models that analyze the index option price in terms of its constituent parts. Traders frequently employ trading strategies based on how rich or cheap the implied correlation is relative to correlations based on historical and economic analysis. When they conclude that implied correlations are too low, they buy the index option and sell options on individual stocks in the index, hoping to gain if realized correlation is higher than implied. This is called a *convergence position*. When they conclude that implied correlations are too high, they buy options on individual stocks and sell the index option. This is called a *divergence position*. Corporate risk managers need to make a judgment about how high or low realized correlation can go in measuring the riskiness of these positions.

Index options are also potentially useful in hedging illiquid basket options. For example, if a market maker has written an option on an average of 50 stocks, all of which are components of the S&P index, hedging the volatility risk of the basket option by buying an option on the S&P 500 index is likely to leave less residual risk than buying options on the 50 individual stocks and it will certainly be far more efficient from an operational risk viewpoint (an error is more likely tracking 50 options positions in single stocks than 1 options position in the index). Also in favor of the index option hedge is that index options are almost always more liquid than single stock options.

However, if the option written was on the average of two stocks that are components of the S&P 500 index, hedging the volatility risk of the basket option by buying options on the two single stocks is likely to leave less residual risk than buying an option on the S&P 500 index. At some point between two and 50 stocks, the index hedge is less uncertain than the individual stock hedge, but it needs to be found empirically through simulation. Simulation is also necessary to measure the residual uncertainty of the index stock hedge for purposes of calculating reserves and limits. The most accurate means of simulation is a Monte Carlo with dynamic hedging in an underlying asset package for which the deltas on individual stocks are computed as the net of the delta on the basket option and the delta on the index option. An approximation that is much easier to compute and reasonably accurate for large baskets is to assume no delta hedging and just compute the tracking error between the two options that occurs at the final payoff.

12.4.4 Options to Exchange One Asset for Another

At the beginning of Chapter 11, we stated that all vanilla options could be viewed as the option to exchange one asset for another. It is equally true, following a result of Margrabe, that every option to exchange one asset for another can be evaluated by the Black-Scholes option formula used for vanilla options (see Hull 2012, Section 25.13). So why should we try to view these as multiasset options? Because by bringing in a third asset that plays no role in the original contract, we can in some cases increase the liquidity of the option's valuation. This can most easily be seen by a concrete example.

Consider an option to exchange 10,000 ounces of gold for £4.5 million. Clearly, this option will be exercised if and only if an ounce of gold at the expiration of the option is worth more than £450. Equally clearly, this contrast has absolutely no reference or relationship to dollars. However, it can be viewed, as a mathematical equivalence, as a spread option on the difference between the dollar price of 10,000 ounces of gold and the dollar price of £4.5 million. To see this equivalence, consider the following:

- The option will be exercised if and only if an ounce of gold is worth more than £450. This is equivalent to saying it will be exercised if and only if the dollar price of an ounce of gold is worth more than the dollar price of £450, which is equivalent to saying it will be exercised if and only if the dollar price of an ounce of gold minus the dollar price of £450 is greater than 0. Multiplying by 10,000, this is equivalent to saying it will be exercised if and only if the dollar price of 10,000 ounces of gold minus the dollar price of £4.5 million is greater than 0.
- If the option is exercised, it can be exercised by buying 10,000 ounces of gold for its then current market price in dollars, exchanging the gold under the options contract for £4.5 million, and selling the £4.5 million for its then current market price in dollars. The (necessarily positive) difference between the dollar sale price and the dollar purchase price represents the payoff of the option.

What has been gained by introducing dollars into the picture? If sterling options on gold have no liquid market, but dollar options on gold and dollar-sterling options have a liquid market, then the gold-sterling spread option can be valued and risk managed based on the implied volatilities of dollar-gold and dollar-sterling vanilla option hedges. Some residual uncertainty will still exist due to the assumed correlation level, but this residual uncertainty may be less than the uncertainty of an illiquid gold-sterling exchange option. As we saw in Table 12.13, this will depend on the gold and sterling-dollar prices not being too highly correlated with one another. If they are highly

correlated, implying a very negative correlation for the long and short positions in the spread basket, then little can be gained from being able to hedge the sensitivity to implied volatilities of dollar-gold and dollar-sterling.

A particular case of an option to exchange one asset for another that draws considerable attention is the large market in convertible bonds; see Hull (2012, Section 26.4) and Tsiveriotis and Fernandes (1998). Convertible bonds offer the bondholder an option to exchange the bond for a fixed number of shares of the firm issuing the convertible bond. Convertible bonds generally have reasonably liquid markets, so there is rarely a valuation advantage to viewing them as spread options. However, when determining trading strategies and evaluating risk exposures, it is often convenient to assess the dependence of convertible bond valuations on the implied volatility of the equity option (more precisely, the equity-cash option), the assumed volatility of the option on a straight (nonconvertible) bond issued by the firm, and the assumed correlation between the bond and the stock.

As discussed at the end of Section 8.3, one trading strategy often pursued is to try to take advantage of the implied volatility for an equity option on the stock of a particular firm being higher than the equity volatility implied by the price of a convertible bond issued by that firm. A trader may decide that buying a convertible is an inexpensive way of buying volatility on the firm's equity price. Or a trader might choose to run a basis position long the convertible bond and short the equity option. Risk analysis of such positions should be sensitive to the reasonableness of assumptions about the volatility of the bond option and the correlation between the bond and stock that have been used to conclude that the convertible bond's equity volatility is cheap. The valuation of a convertible should always be based on observed market prices, not on assumptions about correlation.

Another issue that frequently arises in the management of convertible positions is determining the correct delta to use in hedging a convertible position with stock. It has often been observed that when stock prices are so low that the convertible is far from its exercise price, the actual response of the convertible price to changes in the stock price is far larger than would be expected from a delta derived from a model that only accounts for volatility of the stock price. The explanation of this observation can be found in the correlation between the bond and stock. When stock prices are far below its exercise prices, a convertible bond ought to behave very much like a straight bond, but both the bond and stock price will be impacted in similar ways by changes in the outlook for the firm's earnings (this is discussed in more detail in Section 13.2.4).

If a convertible bond behaves more like a straight bond than a stock, then a straight bond would seem like a better hedge. However, there might be reasons for using the stock as a hedge, such as greater liquidity or ease in

borrowing the stock relative to the straight bond. In such instances, hedging ratios should certainly reflect the assumed correlation between stock and bond prices. But you must be careful to remember that the correlation assumption drives this delta. For example, if the firm's risk reports show a sensitivity to credit spread for the convertible, also showing a high sensitivity to stock price for the convertible in the firm's risk reports would involve a double count of the sensitivity to the bond price—once directly and once through the bond-stock correlation.

12.4.5 Nonlinear Combinations of Asset Prices

When a derivative's payoff is the function of a nonlinear combination of a set of asset prices, none of the three simplifying characteristics that hold for a linear combination can be assumed to be in force. This can be illustrated by a single concrete example: a *quanto* forward whose payoff is calculated by the product of an asset price and FX rate.

On January 25, 2002, stock in the Sony Corporation was trading at 6,080 yen per share and the yen was trading at 134.79 yen per dollar. So the then current dollar price of a share of Sony stock was $6,080/134.79 = \$45.11$. The six-month forward price for Sony stock on that date was also roughly 6,080 yen per share and the six-month forward exchange rate was 133.51 yen per dollar. Suppose a customer comes to a market maker looking to purchase 1,000,000 shares of Sony stock for six-month forward delivery at a dollar price. Possible contracts (see Reiner 1992 for a full discussion) could be:

- Make the purchases at a dollar price fixed in advance. The market maker has a static hedge available (it is an exchange of assets, as discussed in Section 12.4.4). She can purchase 1,000,000 shares for six-month forward delivery at $1,000,000 \times 6,080 = 6,080,000$ yen and purchase 6,080,000 yen for six-month forward delivery at $6,080,000/133.51 = \$45,539,660$, which is the price, without profit margin, she should charge the customer.
- Make the purchase at a dollar price based on the exchange rate, which will be in effect in six months. The market maker has a static hedge available. She can purchase 1,000,000 shares for six-month forward delivery at $1,000,000 \times 6,080 = 6,080,000$ yen. The dollar price will be determined in six months based on the then prevailing exchange.
- Agree that the dollar price per share will differ from the current six-month forward price of $6,080/133.51 = \$45.54$ per share by the percentage change in the yen price per share. So if the yen price in six months is $6,080 \times 110\% = 6,688$, the price per share to be paid will be $\$45.54 \times 110\% = \50.094 . This is a quanto.

No static hedge is available for a quanto. The market maker can begin with a purchase of 1,000,000 shares for six-month forward delivery for 6,080,000 yen and a 6-month forward exchange of 6,080,000 yen for \$45,539,660. However, if the forward share price rises by 10 percent, she now has FX risk on an additional $1,000,000 \times 6,080 \times 10\% = 608,000$ yen and must enter into a forward exchange of these yen for dollars. If the forward FX rate rises by 10% to $133.51 \times 110\% = 146.86$ yen per dollar, she now has stock price risk of an additional 10 percent, since her stock price hedge is for a fixed amount of yen and what she needs is a hedge for a fixed amount of dollars. As the yen weakens against the dollar, she needs to increase the amount of hedge denominated in yen to maintain the dollar amount of the hedge. This pattern, a change in one asset price requiring a dynamic change of the hedge amount of the other asset, is typical of derivatives with payoffs based on the product of two asset prices.

The formula for valuation of a quantoed forward, under the assumption of a bivariate lognormal distribution, is the price of a standard forward multiplied by $\exp(\rho\sigma_S\sigma_F)$, where σ_S is the volatility of the stock price denominated in yen, σ_F is the volatility of the FX rate (that is, the yen price denominated in dollars), and ρ is the correlation between the stock price denominated in yen and the FX rate. A brief explanation of this formula can be found in Hull (2012, Section 29.3), and a more detailed derivation can be found in Baxter and Rennie (1996, Section 4.5). Two important consequences follow from this formula. First, the value of the derivative, even though it is not an option, is dependent on the volatilities of the assets and the correlation. Second, if the correlation is zero, then the valuation formula for a quanto is the same as the valuation formula for a standard forward, so the total impact of the dynamic hedging required must balance out to zero (however, this dynamic hedging could still result in transaction costs).

A derivative with very similar characteristics to a quanto is a difference swap, in which the payoff is based on the future difference between interest rates in different currencies multiplied by a notional principal denominated in one of the currencies. For example, the difference between a dollar interest rate and a yen interest rate may be multiplied by a dollar notional amount. The future dollar interest rate multiplied by the dollar notional amount represents a quantity that can be statically hedged, but a yen interest rate multiplied by a dollar notional amount is a quantoed combination that requires dynamic hedging of both the yen interest rate and the dollar/yen FX rate. For more details, see Hull (2012, Section 32.2) and Baxter and Rennie (1996, Section 6.5).

Once the bivariate lognormal assumption is dropped, more complex valuation algorithms are required. Both the Monte Carlo and trinomial tree approaches discussed in Section 12.3 have the flexibility to be directly

applied to quantos or any other derivative based on a nonlinear combination of asset prices. Both approaches build probability distributions for each asset separately and can incorporate a full volatility surface (and, in the case of Monte Carlo, can incorporate stochastic volatility and price jumps). Both approaches can factor in any desired correlation assumptions between assets. Both approaches can then compute any desired function of the asset prices, no matter how complex, based on the individual asset prices at each node (and Monte Carlo can incorporate full price histories of the assets if they play a role in the function).

Nonlinear functions of multiple asset prices can range from the simplicity of the maximum or minimum price of a basket of assets to the complexity of an involved set of rules for successively dropping high and low prices out of a basket on which an average is being calculated. Some assets in the basket may represent quantoed translations from other currencies. As a further step, options can be written on any of these nonlinear functions, and exotic features such as barriers can be introduced. So long as the Monte Carlo or tree is valuing the nonlinear function correctly, it should also value the option correctly. A general designation for derivatives based on nonlinear functions of multiple asset prices and their derived options is a *rainbow contract*.

Hedging considerations for derivatives on nonlinear combinations are exactly parallel to those for derivatives on linear combinations, so the approach in Section 12.4.2 can be applied. The only difference is that the simple approximation formulas used in that section do not apply. Computations of sensitivities to shifts in asset prices, implied volatilities, and assumed correlations generally need to be evaluated by rerunning the Monte Carlo or trinomial tree valuation model with shifted inputs.

Another interesting example that is similar in structure to the quanto is counterparty credit exposure on a derivative such as an interest rate swap or a FX forward. As discussed in Section 14.3.5, counterparty credit exposure can grow or diminish through time as a function of the interest rate or FX rate driving the value of the derivative. This credit exposure can be hedged by the purchase of credit derivatives or the short sale of bonds issued by the counterparty. The total value of the credit exposure is then the product of the value of the derivative and the credit spread on the counterparty. Similar to a quanto, a dynamic hedge is required. A change in the value of the derivative requires a change in the size of the credit hedge and a change in the size of the credit spread requires a change in the size of the derivative hedge.

In Section 11.3, we examined a case of mean reversion in which there is a narrower dispersion of final underlying price levels than would be implied by a pure random walk and we questioned whether dynamic hedging

costs would be a function of the higher short-term volatility or the lower long-term dispersion. Our answer, based on both Monte Carlo simulation and theory, was that sufficiently frequent rehedging makes dynamic hedging costs depend entirely on short-term volatility, but a trader who wanted to take advantage of anticipated lower long-term dispersion could do so by rehedging less frequently (but with an attendant trade-off of a higher uncertainty of hedging costs).

Let's ask a parallel question for correlation. Suppose you anticipate that two assets will have a strong correlation in terms of long-term trend, but very little correlation in terms of short-term moves. If you are dynamically hedging a position whose valuation depends on correlation, will your dynamic hedging costs be a function of the low short-term correlation or the high long-term correlation?

You shouldn't be surprised to find that the answer is the same for correlation as it is for single-asset volatility. If you re hedge often enough, only the short-term correlation impacts hedging costs. If you want to take advantage of the anticipated long-term trend, you must hedge less frequently and accept a higher uncertainty of hedging costs in exchange for expected hedging costs being influenced by the longer-term correlation.

Many people find this conclusion highly nonintuitive. Consider an example. Suppose you are hedging the counterparty credit risk on an FX forward and that over the life of the forward the exposure continues to grow while the credit rating of the counterparty continuously deteriorates, but the individual moves are uncorrelated. As the exposure grows, you are going to have to buy more credit protection, and it may be hard to believe that you will not have to pay for this increased credit protection at the higher price levels brought on by the deteriorating credit rating.

To help see how this works mechanically, I have provided the **CrossHedge** spreadsheet, which enables you to enter a price history of six prices for each asset and which looks at the hedging of an exotic paying the product of the two asset prices. The spreadsheet shows the hedging and its costs under two assumptions: if the price moves between the two assets are completely uncorrelated and if the price moves between the two assets are perfectly correlated. The complete lack of correlation is implemented by having each price move on the first asset precede in time each price move on the second asset, so there is time to change the hedge quantity before the second asset's price changes. (Remember that for a payoff tied to the product of two asset prices, a change in the price of one asset requires a change in the hedge quantity of the other asset.) Perfect correlation is implemented by simultaneous changes in prices.

Table 12.14 shows the case of deteriorating credit on counterparty credit risk. The first asset is the exposure amount and the second asset is the

discount on the counterparty's bonds. As credit deteriorates, the discount goes all the way to 100 percent, corresponding to the worst possible case of default with no recovery. Despite the fact that the exposure is steadily growing while the discount is steadily increasing, the change in the value of the product is completely hedged in the uncorrelated case. Examining the impact of the individual hedges should impart a better sense of how the hedge works—each change in credit quality and exposure has been hedged by having the right size hedge in place at the time of the change.

12.4.6 Correlation between Price and Exercise

Standard option pricing assumes a correlation of 100 percent between price and exercise; that is, option buyers will exercise their options when, and only when, the price of the underlying asset makes it profitable to exercise. However, in some instances, it can be argued that a correlation of less than 100 percent should be assumed. These arguments rely on a combination of historical experience, showing that a previous correlation has been less than perfect, and on a behavioral analysis of the option buyers, demonstrating that they have motivations that conflict with optimal option exercise. In terms of game theory, standard option analysis, which assumes a correlation of 100 percent, is equivalent to a zero-sum game in which a loss by the option seller is exactly offset by a gain for the option buyer. A correlation of less than 100 percent corresponds to a non-zero-sum game.

Table 12.15 shows the impact of different correlation assumptions, multiplying the payoff based on price and exercise by the probability and summing to get an expected return.

For example, it may be argued that a municipality that has the option to require early repayment of a fixed-rate term deposit without paying any penalty, which is equivalent to a swaption, will only exercise this option in response to a change in its cash needs, which are uncorrelated with interest rate levels. Support for this analysis should certainly include historical studies of how similar municipalities have exercised these options. However, even reasonable explanations of behavior and historical precedent may be questionable evidence. In the absence of any actual legal constraint or internal costs that exercise would entail, it is possible that institutions will become more efficient exercisers of options over time, as they gain financial sophistication or as large economic movements (for example, unusually high interest rates on new deposits) create increased incentives to focus attention.

Such arguments may become more plausible when the option must be exercised by a large group of individuals. Correlation now becomes a question of what proportion of a population will exercise options in a timely fashion, and their diversity of circumstances will argue for less than perfect

TABLE 12.15 Correlation between Price and Exercise

Standard Option						
Price	Exercise	No Exercise		Probability		
110	-10	0	×	50%	0	= -5
90	+10	0		0	50%	
No Correlation						
Price	Exercise	No Exercise		Probability		
110	-10	0	×	25%	25%	= 0
90	+10	0		25%	25%	
Some Correlation						
Price	Exercise	No Exercise		Probability		
110	-10	0	×	35%	15%	= -2
90	+10	0		15%	35%	
Negative Correlation						
Price	Exercise	No Exercise		Probability		
110	-10	0	×	15%	35%	= +2
90	+10	0		35%	15%	

correlation. An example would be a pension plan that guarantees some minimum return on a particular investment strategy. If option exercise were a zero-sum game, the individual investors would withdraw from the plan whenever the investment was below the minimum return in order to collect the guarantee. However, financial institutions that provide these guarantees value them based on behavioral assumptions about the individual participants, whose varied circumstances with regard to age, career, and tax status make the cost of exercising the option different for each subgroup.

An important example of an option exercised by a large group of individuals is the very sizable market in asset-backed bonds, where each bond is backed by a pool of mortgages, automobile loans, or other consumer loans. Although these assets often provide consumers the legal right to prepay the loan without penalty, individual circumstances often get in the way of an economically efficient exercise of this right. First, refinancing a loan often involves substantial personal costs (for example, legal fees, title searches, and the time devoted to the transaction). For an institution on a large loan, these would probably be insignificant relative to gains from exercise, but this may not be true for an individual. Second, some consumers may not be able to refinance due to a deteriorating credit rating or decrease in asset value. Others may have strong personal motives that outweigh the costs of financing, such

as a required move or a divorce forcing a home sale that causes a desirable rate mortgage to be prepaid or the desire to trade a car for a newer model.

Given the enormous size of this asset class and the plausibility of less than perfect correlations, financial firms have invested and continue to invest large amounts of money in research to develop accurate models of this correlation. A good introduction to this asset class is Davidson et al. (2003), with its Chapter 9 introducing the modeling issues involved. For some assets, such as automobile loans, the general conclusion is that correlation tends to be close to zero. For mortgages, correlation is definitely strongly positive; falling mortgage rates trigger massive refinancings, and rising mortgage rates trigger considerably slow refinancings. However, correlation is certainly far from perfect, and the stakes in properly identifying which mortgage bonds represent good investments are sufficiently high to support detailed research trying to predict the relationship between refinancing behavior and prevailing mortgage rates by population subcomponent, such as the geographic region or size of mortgage. The relationships developed are often quite complex. The behavior depends not just on current mortgage rates, but also past mortgage rates and yield curve shape. Consumers are found to be sensitive not only to the current refinancing advantage, but also to beliefs as to whether that advantage will be growing, since the costs of refinancing are high enough to cause consumers to attempt to minimize the number of times they refinance. Another factor, known as *burnout*, indicates that a consumer population that has already experienced a period of low rates will show lower refinancing response (as a proportion of mortgages still outstanding) in a subsequent low rate period. This is presumably due to the proportion of those who did not refinance the first time who cannot afford to refinance. Monte Carlo models of the correlation are used to project consumer behavior under a variety of possible future interest rate movements, and bonds are ranked on the basis of *option-adjusted spread* (OAS)—the spread the bond is earning over a comparable-maturity Treasury after taking into account the cost of the refinancing option based on the assumed correlation.

Why does this spread remain? One reason is certainly that these correlation relationships are only estimates based on past data that could prove to be wrong. When unanticipated shifts in consumer behavior on refinancings are observed, such as a prolonged period of very low rates resulting in greater consumer education about refinancings, leading to refinancing levels that substantially exceed those predicted by models based on past data, OAS can show large rapid increases. To some extent, this will later reduce as Monte Carlo models are updated to accommodate the new experience, but some OAS increase may persist, reflecting an increase in uncertainty over the accuracy of such models.

12.5 CORRELATION-DEPENDENT INTEREST RATE OPTIONS

Throughout Chapter 11 on vanilla options and in Sections 12.1 and 12.2, we have dealt with options whose underlying can be regarded as a forward to a set future date. As we discussed at the beginning of Chapter 11, all uncertainty about discounting rates for these models can be collapsed into the volatility of the forward. However, some options have payoffs that depend on forwards for several different future dates (but with all forwards on the same spot underlying). The primary example would be an American option that gives the option holder freedom to determine the timing of payoff. More complex dependence on different forwards can be seen in the products we examined in Section 12.3, such as barrier options.

Options that depend on forwards for several different future dates can usefully be viewed as options on multiple underlyings with all relationships between these forwards built into the correlation structure assumed between the forwards. Indeed, this is the approach to multifactor interest rate models that has predominated over the past two decades in the form of the Heath-Jarrow-Morton (HJM) models (see Hull 2012, Section 31.1) and the LIBOR market models, also known as Brace-Gatarek-Musiela (BGM) models (see Hull 2012, Section 31.2).

Should we then just view these products as a particular class of options with multiple underlyings and consider their risk management issues as already having been dealt with in Section 12.4? One reason for not availing ourselves of this convenient shortcut is that the large volume of these options that actively trade encourages extra effort to try to find a simpler structure and faster computation time for subsets of this product. Another reason is that this represents the only class of multiasset options where some reasonable liquidity exists in products that require correlation inputs to value, so it is worth studying how much information on correlation can be extracted from observed market prices.

Three levels of models are essentially available, of increasing mathematical and computational complexity. The simplest level includes the binomial and trinomial tree models in which the relationship between different forwards is treated as constant. In Section 12.5.1, we examine risk management using these models and the conditions under which more complex models are required. The second level includes the single-factor interest rate models in which the relationship between different forwards is treated as stochastic. In Section 12.5.2, we examine risk management using these models and the conditions under which the third level of full-blown multifactor HJM or BGM models are required. Finally, in Section 12.5.3, we look at how much correlation information can be extracted from observed market prices.

12.5.1 Models in Which the Relationship between Forwards Is Treated as Constant

We have already encountered binomial and trinomial tree models in which the relationship between forwards is treated as constant—the local volatility models discussed in Section 12.3.2. Recall that this section was devoted to options whose payoff depends on the underlying price of a single asset at several different times. Because values of the asset at several different times are involved, we needed to be concerned with hedging and valuation depending on several different forwards. However, the only way to avoid treating these different forwards as multiple assets is to assume that a constant relation exists between them. This is in effect what is done in the local volatility models of Section 12.3.2, since the only variable changing on the tree is the spot price of the asset and all forward prices are derived based on fixed interest rate relationships between forward and spot prices.

In this section, we study the simplest, most widely traded, and best-known version of a product that depends on the underlying price of a single asset at several times—the *American option*. American options differ from European options by a single added feature: the right of the option buyer to exercise the option at any time. A simple variant restricts the right to exercise to several specified times and is variously known as a *semi-European*, *semi-American*, or (as a geographic middle ground between European and American) *Bermudan option*.

American and Bermudan options have long been valued using binomial trees (the Cox-Ross-Rubinstein model) and more recently using trinomial trees to allow for nonflat volatility surfaces. See Hull (2012, Chapter 12) for the mathematics of the binomial tree. See Clewlow and Strickland (1998, Chapter 5) for the use of trinomial trees to incorporate the volatility surface. The key assumption is that the relationship between the forwards remains fixed. Most typically, this is represented by a constant interest rate and forward drift (or constant dividend rate, with *drift* defined as the interest rate less the dividend rate). However, any constant set of relationships between forwards can be accommodated with no increase in complexity or cost of computation, as discussed in Hull (2012, Section 20.5).

Four factors drive the value of early exercise (all of this discussion is for calls—we are continuing our convention from Chapter 11 of treating all options as calls):

1. **Price.** When prices rise, it increases the probability of price levels high enough to warrant early exercise, so early exercise value increases.
2. **Volatility.** The more volatile the price, the greater the incentive not to exercise early in order to take advantage of the time value of the option.

However, high volatility means a greater percentage of price moves will be large enough to warrant early exercise. So the two impacts of higher volatility run in opposite directions. In practice, the second effect is usually larger, and higher volatility increases early exercise value.

3. **Financing cost.** The higher the net cost of funding the delta hedge of the option, the greater the incentive to exercise early. However, if net financing cost is earning the option buyer money on his delta hedge, it discourages early exercise. An equivalent way of viewing this is through the drift of the forward. If drift is positive, this decreases the incentive to exercise the call early since it is likely the call will be worth more after the upward drift. If drift is negative, this increases the incentive to exercise the call early since it is likely the call will be worth less after the downward drift.
4. **Discount rate.** Early exercise allows earlier receipt of option payoffs. This is more valuable the higher the discount rate, so high discount rates encourage early exercise.

The **AmericanOption** spreadsheet illustrates the computation of American option values using a Cox-Ross-Rubinstein binomial tree. It focuses on the computation of the early exercise value, defined as the excess value the American option possesses over the corresponding European option. Table 12.16 shows some sample results.

Note from Table 12.16 the relatively small impact of discount rates on early exercise relative to drift. Since exchange-traded American options are all options on a fixed forward, they all have zero drift, so early exercise value is quite small. This explains the claim made at the start of Chapter 11 that exchange-traded American options have little valuation difference from European options.

Hedges can be established for the impact on the early exercise value of all four of these factors, as illustrated in Table 12.16. For delta and vega, we calculate the ratio of American option delta and vega to the corresponding European option delta and vega, enabling the American to be represented in delta reports and price-vol matrices for the corresponding European option. For discount and drift, the sensitivity of the early exercise value to a 100 basis point shift is calculated and can be used to establish a hedge. This is a comparable situation to vega hedging an option you are valuing using the Black-Scholes model—the theory behind the model says volatility is constant, but you are going “outside the model” to hedge against volatility uncertainty. Here we are determining the early exercise value using a model that says that discount rate and drift are constant, but we are establishing a hedge against an uncertain discount rate and drift. The liquid proxy for the American option would be a combination of the corresponding European

Strike	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%
Time	1	1	1	1	1	1	1	1	1	1
Volatility	20%	20%	20%	20%	20%	20%	20%	20%	20%	30%
Rate	0%	1%	0%	0%	0%	5%	10%	10%	10%	0%
Drift	0%	0%	-1%	-5%	0%	0%	0%	0%	0%	-1%
European price	7.97%	7.89%	7.44%	5.57%	7.58%	7.21%	7.21%	7.21%	7.21%	11.37%
American price	7.97%	7.90%	7.52%	6.09%	7.66%	7.40%	7.40%	7.40%	7.40%	11.45%
Early exercise	0.00%	0.01%	0.08%	0.51%	0.09%	0.19%	0.19%	0.19%	0.19%	0.08%
Vega	0.40%	0.39%	0.39%	0.38%	0.38%	0.36%	0.36%	0.36%	0.36%	0.39%
Early exercise as % of:										
European price	0.00%	0.17%	1.05%	9.23%	1.13%	2.65%	2.65%	2.65%	2.65%	0.70%
Vega	0.00%	3.46%	19.72%	136.95%	22.74%	53.11%	53.11%	53.11%	53.11%	20.16%
American / European delta	100.02%	100.24%	100.96%	110.39%	101.63%	103.96%	103.96%	103.96%	103.96%	100.58%
American / European vega	100.00%	100.18%	99.99%	99.84%	101.15%	102.68%	102.68%	102.68%	102.68%	100.08%
Rate sensitivity per \$1mm	\$136	\$163	\$202	\$214	\$201	\$219	\$219	\$219	\$219	\$291
Drift sensitivity per \$1mm	\$778	\$844	\$1,044	\$1,267	\$971	\$1,037	\$1,037	\$1,037	\$1,037	\$956

TABLE 12.16 Early Exercise Values and Hedges for American Option

option and the extra hedges needed for the exposure to discount rate and drift. This can easily be converted into a Monte Carlo simulation of the differences in final payout between an American option and this liquid proxy, given a simulation of changes in the underlying price, the discount rates, and drift.

The critical assumption when calculating these hedges is that discount rate and drift risk can be valued and hedged as variables independent of the spot price risk. Equivalently, the assumption is that the level of forward rates is uncorrelated with the shape of the forward rate curve. This assumption is reasonable for most equities, questionable for FX and commodities (refer back to our discussion of mean reversion in Section 11.3), and certainly false for interest rate options, since high correlation will exist between the rate determining the payoff and the rates determining the discount and drift.

Is it possible that the impact of this correlation is small enough to ignore for practical purposes? As shown in Table 12.16, when drift is positive or zero or when it is not too negative, the total size of the early exercise value is not too large so any impact of correlation can probably be ignored. When drift is quite negative, early exercise value becomes significant and it is likely that the impact of correlation between interest rates needs to be taken into account. To do so requires some type of term structure model; the factors influencing choices between these models are discussed in the next section.

This is particularly true for options on bonds or on swaps, where the *pull to par* causes drift to be very negative. Because the duration of a bond or swap gets shorter as time passes, the impact of interest rates on prices is continuously declining. So an option holder faced with an early exercise decision knows that the current price premium is likely to diminish through time—if interest rates don't move further in her favor, current rate levels will translate into a smaller price advantage in the future. This is true both for options that pay on rising interest rates and those that pay on falling interest rates, since both high bond prices based on low interest rates and low bond prices based on high interest rates move in the direction of par if rates stay the same as time to maturity diminishes.

If any substantial reduction in the duration of an underlying bond or swap occurs during the tenor of an option, this negative drift will require a term structure model. If no substantial reduction in duration occurs over the option tenor, then a Cox-Ross-Rubinstein model with the duration held constant can be used as a reasonable approximation. A rule of thumb that is often used is that this approximation is suitable as long as the duration of the underlying at the start of the option life is at least 10 times as great as the option tenor. So this rule of thumb would allow the use of a Cox-Ross-Rubinstein model for a six-month option on a 10-year bond, but would insist on a term structure model for a one-year option on a five-year bond.

12.5.2 Term Structure Models

The most liquid options products based on interest rates are *caps*, *floors*, and *European swaptions*. A European swaption is an option to enter into a swap at some fixed future date at a rate fixed at the time of entering into the option. A one-period swap is a forward rate agreement (FRA) and, by convention, an option on an FRA is called a *caplet* if it is an option to receive floating and pay fixed (i.e., it pays off when rates are high) and is called a *floorlet* if it is an option to pay floating and receive fixed (i.e., it pays off when rates are low). Market practice is to sell caplets and floorlets in bundles, called strips, which are called caps and floors, respectively (so a swaption is an option on a bundle of FRAs, a swap, and a cap or floor is a bundle of options on FRAs). For example, a five-year cap on three-month LIBOR would consist of a strip of nineteen options on three-month FRAs that have starting dates beginning at times starting three months from now and ending four years and nine months from now. Market convention will quote a single volatility for a cap or floor, which is then applied to each of the constituent FRA options—but this is just a convention to make it easy to communicate. Actual pricing of a cap or floor evaluates each individual FRA option at the appropriate volatility, sums the resulting prices to arrive at the price of the cap or floor and then solves for a single volatility, which, applied to each individual FRA option, would result in this summed price.

A European swaption on a one period swap is identical to a caplet or floorlet. In our discussion of term structure models, the models used to price complex interest rate products, in both this section and in Section 12.5.3, we will for convenience sometimes refer to all of the liquid instruments being used for calibration of the models as swaptions, even though those which are options on individual FRAs are more accurately called caplets or floorlets.

Broadly speaking, term structure models come in two varieties: single-factor models that assume that the correlation between all forwards is 100 percent and multifactor models that can accommodate less than perfect correlation structures. Both types of model can handle a correlation between the underlying of the option and drift. Multifactor models are obviously more accurate, but add a considerable cost in computation time and complexity. Since American and Bermudan options on swaps and bonds are by far the most utilized exotic in the interest rate options market, there is a strong incentive to try to use single-factor models for this product as long as accuracy is reasonable.

A critical fact about interest rate options, which any term structure model needs to deal with, is that options of the same tenor for bonds (or swaps) of different maturities tend to have lower interest rate volatilities for the long

TABLE 12.17 Annualized Volatility of Dollar Par Swap Yields

Tenor	1 Year	3 Years	4 Years	5 Years	6 Years	7 Years	8 Years
Annualized volatility	16.95	16.24	15.95	15.81	15.26	15.18	14.88
Tenor	9 Years	10 Years	12 Years	15 Years	20 Years	30 Years	
Annualized volatility	14.72	14.65	14.25	12.50	12.87	11.99	

maturity. This can be confirmed both by observations of implied volatilities from market quotes and from historical volatility observations of par bond or swap yields. (For example, Table 12.17 shows annualized volatilities by tenor based on six years of dollar par swap yields between 1996 and 2001—see the **DataMetricsRatesData** spreadsheet for the underlying data.)

Broadly speaking, this fact can be explained by some combination of the following two theses:

1. Forward rates are less than perfectly correlated with one another and the longer the bond maturity, the more its volatility is dependent on the correlation between forwards.
2. Longer-term forwards have lower volatility than shorter-term forwards.

The latter theory implies that interest rates are mean reverting, since it requires the standard deviation of longer-term forwards to be lower than that produced by a pure random walk driven by the volatility of shorter-term forwards. To see the interaction between the correlation and volatility of longer-term forwards when explaining swaption volatility, refer to Section 12.5.3.

Because it assumes that all correlation between forwards is 100 percent, a single-factor model must utilize the lower volatility of long-term forwards to drive the observed volatility structure of swaptions. To what extent does forcing one of these two levers to bear all of the explanatory weight distort valuation and hedging? In principle, to answer this question, build the best multifactor term structure model you can; calibrate both this multifactor model and the single-factor model that is proposed for production use to the current set of vanilla cap, floor, and European swaption prices; and then compare their output in valuing exotic products.

Although this is too daunting a computational task to attempt here, I will give a flavor of what this analysis is like for one very simple case: a three-year time horizon; three liquid vanilla products—a one-year caplet on a one-year LIBOR, a two-year caplet on a one-year LIBOR, and a one-year

swaption on a two-year swap; and a flat implied volatility surface with respect to strike. We will assume the two-year swap is on a one-year LIBOR. We will take advantage of the equivalence between swaps and packages of forward rate agreements (FRAs), as noted in Section 10.1.6. The notation we will employ is to label a FRA by the time at which its rate is determined and the time at which it settles. So a 2–3 FRA has a rate determined at the end of two years based on what would then be the one-year rate.

The model will be calibrated to the current one-year LIBOR, 1–2 FRA and 2–3 FRA, the one-year volatility of the 1–2 FRA, the first-year volatility of the 2–3 FRA, the second-year volatility of the 2–3 FRA, and the one-year correlation between the 1–2 FRA and the 2–3 FRA. In addition to valuing the liquid vanilla products, we will value four exotics:

1. A two-year Bermudan swaption that can be exercised either at the end of year 1 based on the then prevailing two-year LIBOR or at the end of year 2 based on the then prevailing one-year LIBOR.
2. A two-year caplet on a one-year LIBOR that can knock out depending on the level of a one-year LIBOR in one year.
3. A forward-start caplet on a one-year LIBOR that has a one-year tenor and begins in one year with a strike set to the then one-year LIBOR.
4. A one-year tenor option on the spread between a two-year LIBOR and a one-year LIBOR.

Our full term structure model is in the **TermStructure** spreadsheet. It is a simple Monte Carlo implementation. It takes advantage of the fact that only two exercise points are available for the Bermudan to value it by the following trick. At the end of two years, exercise is a simple decision. If you are in the money at the end of one year, you have a choice between early exercise, which gives you a two-year par swap, or waiting a year, which is equivalent to a one-year caplet on a one-year LIBOR. So you just choose the maximum value between the two-year swap and the one-year caplet on the one-year LIBOR.

Using a flat initial rate curve of one-year LIBOR = 1–2 FRA = 2–3 FRA = 7 percent, two scenarios can be computed as shown in Table 12.18, which can be verified with the spreadsheet.

Notice the following:

- The inputs have been deliberately chosen to calibrate to the same vanilla option prices in both scenarios.
- The higher correlation in scenario 2 must be balanced by the lower volatility of the longer-term 2–3 FRA in the first year in order to match the one-year swaption price. This must be followed by higher volatility in

TABLE 12.18 The Valuation of Interest Rate Volatility Products under Two Scenarios

	Scenario 1	Scenario 2
Inputs		
First-year volatility of 1–2 FRA	20.00%	20.00%
First-year volatility of 2–3 FRA	19.50%	14.00%
Second-year volatility of 2–3 FRA	14.83%	20.00%
First-year correlation of 1–2 FRA and 2–3 FRA	50.00%	100.00%
Valuations		
One-year caplet on one-year LIBOR	0.519	0.519
One-year swaption on two-year swap	0.810	0.810
Two-year caplet on one-year LIBOR	0.559	0.559
Bermudan swaption	0.936	0.949
Knock-out caplet	0.400	0.447
Forward-start option	0.645	0.541
Spread option	0.518	0.153

the second year when its time to maturity is shorter so that the combined first- and second-year volatilities fit the price of the two-year caplet.

- Despite a very large difference in correlations between the two scenarios, the Bermudan swaption values close to equal in both scenarios. This reflects a trade-off between lower volatility of the 2–3 FRA in the first year, which decreases the value of early exercise, and higher volatility of the 2–3 FRA in the second year, which increases the value of the option in those cases in which early exercise does not occur.
- The knock-out caplet also shows values close to equal in both scenarios. Lower correlation increases the chances that a high 2–3 FRA, which leads to a higher caplet value, will be accompanied by a 1–2 FRA that is low enough that the caplet will not knock out. This leads to a higher caplet value but is offset by the lower second-year volatility that accompanies the lower correlation.
- Lower correlation causes the forward-start option to have a higher value by adding volatility in the relation between the strike and forward to the volatility of the forward.
- The largest difference between the two scenario valuations is for the spread option, which is the product most directly tied to yield curve shape rather than level. It values much higher when lower correlation permits greater variability in shape.

This single case is consistent with the intuition of most practitioners in the interest rate options market. For Bermudan swaptions, a one-factor

model can be calibrated to current vanilla prices and give reasonable results, but as you move toward products that are more dependent on the future shape of the yield curve, multifactor models become more of a necessity. Although this demonstration for a two-period case is far from conclusive for longer-term swaptions, see Andersen and Andreasen (2001) for similar conclusions in a more general setting. This spreadsheet can be useful for gaining intuition about the direction and order of magnitude of correlation assumptions on different interest rate exotics.

When multifactor models are utilized, traditionally the primary choice has been between models that assume a normal distribution of the short-term rate, such as Hull-White, and models that assume a lognormal distribution of the short-term rate, such as Black-Derman-Toy or Black-Karasinski. See Hull (2012, Section 30.3) and Rebonato (1998, Chapters 12 and 13) for an exposition of these models.

The discussion on which of these approaches to use has often centered on whether one believes that normal or lognormal distributions of rates give closer correspondence to historical experience. This line of argument is getting to seem rather dated in light of the almost universal adoption of full volatility surfaces that accommodate mixtures of normal and lognormal assumptions in equity, FX, and commodity options models (see Section 11.6.2). As we discussed with barrier options in Section 12.3.1, not getting the shape of the implied volatility surface correct can result in major errors in the valuation of exotics. Bermudans share a key characteristic of barriers in that the strike level that determines the termination of the option can be different than the strike level that determines the value of the option, making the correct fitting of the relative volatility between these two strike levels an important determinant of valuation. A more modern approach to utilizing the full implied volatility surface when creating a single-factor interest rate options model can be found in Khuong-Huu (1999).

Other factors that go into the choice and accuracy of a single-factor model include:

- The Hull-White model offers a strong computational advantage in that the forward value of a bond or swap can be computed by analytic formula for any node of the tree (see Hull 2012, Section 30.3). By contrast, lognormal models of the short rate must extend the tree all the way out to the maturity of the bond or swap and solve backwards on the tree to determine a forward value.
- It is possible for interest rates to become negative in some portion of the tree in normal models of the short rate. If you believe this is economically unrealistic (refer back to the discussion in Section 10.3.2), then you would want to get estimates of the degree of impact this could

have on valuations and hedges; see Rebonato (1998, Section 13.9) for a balanced discussion of this issue and other strong and weak points of the Hull-White model.

- The limitation of having just a single factor to calibrate with leads to conflicts between the desire to correctly fit observed prices of potential hedging instruments and the desire to avoid unrealistic evolutions of the rate curve; see Rebonato (1998, Sections 12.5 and 13.9) for an extended discussion.
- Black-Derman-Toy is a binomial tree model, in contrast to the trinomial tree models of Hull-White and Black-Karasinski, and is far easier to implement and maintain than the trinomial tree models. The price paid for this convenience is that the speed of mean reversion is determined and cannot be set as an input parameter. Overcoming this weakness was the primary motivation for the introduction of Black-Karasinski (see Hull 2012, Section 30.3). As a result, Black-Derman-Toy can only calibrate to a limited subset of vanilla options on any given run. For instance, in our two-period example, it could only calibrate to the one-year swaption on a two-year swap and the two-year caplet, but not to the one-year caplet. This could potentially reduce the number of possible hedging instruments that have been correctly priced by the model; see Rebonato (1998, 12.5) for further discussion.
- All of the single-factor models share the issue that shifts in rate levels will cause shifts in the package of vanilla options that form a good hedge for an American or Bermudan option. Table 12.19 shows an illustrative example. This table is based on a 10-year annually exercisable Bermudan call option on a 10-year swap with a coupon rate of 7 percent and flat volatility surface at 20 percent. As should be expected, falling rates increase the value of the call, making early exercise more likely and thus increasing the impact of early volatility relative to later volatility. Rising rates decrease the value of the call, making early exercise less likely and thus increasing the impact of late volatility relative to earlier volatility. It is then easy to solve for a set of European options with similar exposure to the forward volatility curve. However, a package of vanilla options that matches the distribution of exposure at one rate level will no longer match the exposure at a different rate level.

Rebonato (2002), particularly Chapters 8, 9, and 10, is an excellent source of detailed examples and exposition regarding the subtleties of calibrating term structure models to market prices of caps, floors, and European swaptions. Rebonato's discussion of term structure models is very much consistent with the conclusions of Gatheral (2006) regarding dynamic hedging models discussed in Section 12.3.2—models that correctly price all of the liquid instruments can still differ substantially in how the volatility

TABLE 12.19 Impact of Rate Levels on the Forward Volatility Curve Dependence of a Swaption

Year	Flat Rate Level		
	5%	7%	9%
1	1%	0%	0%
2	12%	4%	0%
3	13%	8%	3%
4	13%	11%	6%
5	12%	12%	8%
6	11%	13%	13%
7	11%	13%	15%
8	10%	13%	17%
9	9%	13%	19%
10	8%	12%	20%

surface evolves. And differences in volatility surface dynamics can translate into substantial differences in the cost of hedging an exotic instrument with more liquid instruments.

Looking back once more to Section 8.4, the risk management approach to this should be Monte Carlo simulation of the P&L resulting from following a hedging strategy implied by a particular model, as recommended by Derman (2001). Once again, the difficulty is the computational burden of needing to compute required rehedging along all the different Monte Carlo paths. In this case, I don't have a static or quasistatic hedging alternative to offer that I have actually had experience with. A suggested approach would be, to take a Bermudan swaption as an example, to start with a liquid proxy of a package of vanilla swaptions as in Table 12.19, based on current rate levels. The idea would be to hold this package fixed as you go forward on the Monte Carlo path.

This approach runs into two problems. The first is that some of the European options will reach expiry and deliver a payoff, leaving the Bermudan option decidedly underhedged. Perhaps a simple rule could be followed, such as every time a European swaption reaches expiry, bring the package of European swaptions back up to 100 percent of the Bermudan swaption by buying new European swaptions in the same proportion as the remaining European swaptions in the original package. The second problem is how to decide when on each path Bermudan options should be exercised without needing repeated reruns of the term structure model. One approach could be to use some rule of thumb to govern exercise. Another approach would be to assume exercise on each path will take place at the time that, looking back at the path from final expiry, would be the least favorable to the trading desk.

12.5.3 Relationship between Swaption and Cap Prices

Since a European option on a swap or bond can be a reasonably liquid instrument, and since we can view it as equivalent for valuation purposes to an option on the baskets of FRAs, which the swap is equivalent to, we can try to extract information on market-implied correlations between FRAs from liquid prices. How much correlation information can we extract? Not that much, unless we are willing to make some additional assumptions.

To see why, let's start by considering a simplified market in which only two FRAs trade a 1–2 year and a 2–3 year. The natural options would be a one-year caplet on the 1–2 year, a two-year caplet on the 2–3 year, and a one-year swaption on the combination of 1–2 year and 2–3 year. To price these three options, we need inputs for the following underlying variables: the volatility of the 1–2 year FRA in year 1, the volatility of the 2–3 year FRA in year 1, the correlation between these two FRAs in year 1, and the volatility of the 2–3 year FRA in year 2. Unfortunately, four underlying variables are present and only three options need to be priced. So it will not be possible to extract a correlation from the prices, as we have seen in the example of the previous section, unless we are willing to place some tight restrictions on the possible structure of FRA volatilities.

When we move to more realistic market assumptions, the situation does not improve. The **Swaptions** spreadsheet can take price inputs for one-year LIBOR caplets from one to 10 years and all possible swaption prices involving an integral number of years less than or equal to 10 (for convenience, the prices are quoted as the equivalent Black-Scholes implied volatility). Based on an assumption as to correlation structure, the spreadsheet uses the Excel Solver to find a structure of underlying FRA volatilities that explains the prices. From your experimentation with the spreadsheet (see Exercise 12.10), you can confirm that a wide range of different correlation assumptions is consistent with a single set of prices. We have assumed zero volatility skew and smile throughout this discussion, but changing this assumption will not improve the situation.

It is possible to come to conclusions about the probability of different underlying FRA volatility structures based on historical observation, and this may result in constraints that would at least give a tight range of possible market-implied correlations. For example, one proposal that has both intuitive appeal and some empirical support is to assume that the volatility of FRAs is a function of how far they are from maturity. So the volatility of a 2–3 year FRA in its second year, when it is in the final year of its life, should be the same as the first-year volatility of a 1–2 year FRA and the third-year volatility of a 3–4 year FRA. The intuition behind this assumption is that new information has its greatest impact on nearby borrowing rates, so we

should expect to see greater volatility in nearby rates and lower volatility as you go farther out in maturity (this is equivalent to assuming mean reversion of interest rates, as we saw in Section 12.5.2). So if the caplet volatility in the market for a 1–2 year FRA is 23 percent, but is 22 percent for a 2–3 year FRA, it is reasonable to assume that this 22 percent can be decomposed into a 21 percent volatility in the first year, when the FRA still has over a year to go, and a 23 percent volatility in the second and last year.

This assumption is powerful enough to enable all FRA correlations to be derived from swaption prices. To see this, consider that if you have N different FRAs for which you provide volatility assumptions, this can provide

pricing for $\frac{N^2 + N}{2}$ different swaptions (N in period 1, $N - 1$ in period 2,

and so on — $\sum_{i=1}^N i = \frac{N^2 + N}{2}$). The total number of correlations that can be

specified between FRAs is $\frac{N^2 - N}{2}$ since the N correlations of a FRA with

itself must be 100 percent and a correlation between FRA $_i$ and FRA $_j$ must equal the correlation between FRA $_j$ and FRA $_i$. If you specify that FRA volatility is completely determined by time to maturity, it reduces the number of volatilities that can be specified to N . The total of specified volatilities

plus specified correlations is then $N + \frac{N^2 - N}{2} = \frac{N^2 + N}{2}$. So if all $\frac{N^2 + N}{2}$

swaption prices are specified, a unique set of FRA volatilities and correlations that can explain them must exist.

However, it is possible that placing severe constraints on the relationship between different FRA volatilities will not leave enough freedom to find implied correlations that fit market swaption prices. It can also be the case that caplet volatilities decline too steeply with time to be consistent with the assumption of FRA volatility being a function only of time to maturity; compare this with the discussion in Rebonato (1998, Section 4.5).

Rebonato (2002, Section 9.1.3) makes a case that swaptions volatilities tend to be persistently higher than caplet volatilities due to supply and demand considerations. This is due to consistently high demand from corporate borrowers for cap protection of borrowing costs, while issuers of puttable bonds and buyers of callable bonds are willing to sell the options they own for a fixed upfront price, creating a supply of swaption protection. In Section 9.1.3, along with Sections 1.2 and 6.1.2, Rebonato warns against trying to fit models of exotic interest rate products to both caplet and swaption volatilities, since the difference in volatility levels due to the imbalance of supply and demand factors may result in unrealistic implications for the evolution of volatilities, which may in turn lead to future trading losses. This point is roughly

similar to one made in Section 10.2.1 of this book, the need to account for the trade-off between basis risk and liquidity risk in considering the degree to which an exact fit to market prices should be attempted in a model designed to infer prices of illiquid instruments from more liquid instrument prices.

EXERCISES

12.1 Using the BasketHedge Spreadsheet

1. For a flat volatility assumption (that is, smile = 0 and skew = 0), check the calculation of the square root option in the **Main** worksheet against another pricing method. The method could be analytic (that is, based on solving a PDE), use Monte Carlo simulation, or use a binomial or trinomial tree. Whatever method you choose, make sure you check its accuracy by pricing ordinary options and comparing the answers to the Black-Scholes formula.
2. Pick another type of nonlinear payoff. Change Column C in the **Main** worksheet to calculate a hedge and pricing. Check the results for a flat volatility assumption against another pricing method, as in part 1 of this exercise.
3. Check the impact of smile and skew on the pricing of each of the following:
 - a. The square root option.
 - b. The option you priced in part 2 of this exercise.
 - c. The single-asset quanto priced in the **Quanto** worksheet.
 - d. The log contract priced in the **Log** worksheet.
 - e. The convexity risk hedge priced in the **Convexity** worksheet.
 - f. The call-on-a-call option priced in the **Compound** worksheet.
4. Change Column C in the **Compound** worksheet to price:
 - a. A put-on-a-call compound option.
 - b. A call-on-a-put compound option.
 - c. A chooser option that as of the first expiry time (B1) turns into whichever is more valuable between a call and a put priced at the same strike (B5) to a second expiry time (B4) (see Hull 2012, Section 25.7).
5. For a call-on-a-call option and all three of the options in part 4 of this exercise, use the **Compound** worksheet to determine how much sensitivity remains to future implied volatility after exposure to the price level has been hedged.

12.2 Using the BinaryMC Spreadsheet

Assume you are long one binary option and short a second binary option of the same size. Create a set of examples to show that there is a lower probability of loss:

- a. The closer the two binary options are in maturity date.
- b. The closer the two binary options are in strike.
- c. The greater the correlation in the underlying instruments of the two binary options.

Also show that the variability of results can be reduced by narrowing the spread between the call options used as liquid proxies for the binary options.

12.3 Using the CarrBarrier Spreadsheet

Using the same price strike, up barrier, down barrier, and original time to expiry as the one used in Table 12.7, perform the following:

1. Test the validity of the claim that unwind P&L is zero whenever drift and skew at unwind are zero. Try different combinations of time to expiry, at-the-money volatility, smile, and rate at the time the barrier is hit. Also try different combinations of drift and skew at the time the option is originated.
2. What conclusions can you draw about the pattern of dependence of unwind P&L on different values of drift?
3. What conclusions can you draw about the pattern of dependence of unwind P&L on different values of skew?

12.4 Using the CarrBarrierMC Spreadsheet

Create a set of examples to show the sensitivity of loss probability to changes in the standard deviation of skew and the standard deviation of drift.

12.5 Using the OptBarrier Spreadsheet

Take a down-and-out call case that you have analyzed using CarrBarrier and analyze it using OptBarrier. Use the optimization criterion of 100 percent of the maximum absolute error:

1. First use OptBarrier with four possible times and four possible at-the-money volatilities, but only one possible smile, skew, and

drift—smile, skew, and drift are all set to zero. Confirm that the values you derive for the option price are close to those that **CarrBarrier** derived.

2. Change skew to a single value of 10 percent and see what option values result.
3. Change drift to a single value of -3 percent and see what option values result.
4. Change skew to have two values—one 0 and one 10 percent—and see what option values result and what the resulting degree of uncertainty of closeout cost is. Compare this uncertainty of forward cost to that of the **CarrBarrier** for the same level of skew and drift.
5. Change drift to have two values—one 0 and one -3 percent—see what option values result and what the resulting degree of uncertainty of closeout cost is. Compare this uncertainty of forward cost to that of the **CarrBarrier** for the same level of skew and drift.

12.6 Using the DermanErgenerKani Spreadsheet

1. Use the spreadsheet to check the results given in Table 12.6. Then examine the impact on unwind P&L of deviations between the assumptions about unwind conditions in C8:C12 and the actual unwind conditions in C17:C21. Create a table to show the impact of changes in rate, drift, smile, and skew.
2. Verify that any changes made in initial conditions in B8:B12 will only change the initial price of setting up the hedge and will not have any impact on unwind P&L.

12.7 Using the BasketOption Spreadsheet

1. Check on the sensitivities shown in Table 12.13.
2. Create some examples to check that the General Case and the 3 Asset Case give the same answers for cases with just two or three assets.
3. Using the General Case, tabulate the rate of change in base case volatility and sensitivity to changes in volatility and correlation as the number of assets increases. How does this differ at base correlation rates of 0, 25, and 50 percent?

12.8 Using the CrossHedge Spreadsheet

Try different price paths for the two assets and confirm that they always show zero P&L for the uncorrelated case. What patterns do you

observe for the P&L in the correlated case? For example, what distinguishes cases that lead to gains from cases that lead to losses? What influences the size of the gains or losses?

12.9 Using the TermStructure Spreadsheet

1. Reproduce the results in Table 12.18, which will verify that two different combinations of volatility and correlation input can produce the same valuations for vanilla products but different valuations for exotic products.
2. Find other combinations of volatility and correlation inputs that produce the same valuations for the vanilla products and determine the sensitivity of the exotic products to these inputs.
3. Create your own exotic product by specifying a different payout structure in column J and determine its sensitivity to different combinations of input volatility and correlation that leave vanilla product pricing fixed.

12.10 Using the Swaptions Spreadsheet

Start with input swaption and FRA rates as follows:

- All FRA rates at 7.0 percent.
- Swaption volatilities from Table 12.20.

These swaption volatilities display the usual pattern observed in the market of declining as swap tenor increases:

1. Input correlations of 90 percent for all combinations and use the Solver to find a set of FRA volatilities that correspond to this case.
2. Replace all the 90 percent correlations with 80 percent correlations and use the Solver to find a set of FRA volatilities that correspond.
3. You now have two different sets of FRA volatilities that can explain the same set of swaption volatilities—one based on higher correlation levels than the other. What are the patterns of difference you see between these two sets of volatilities, and how would you explain the linkage between these patterns and the difference in correlation levels?

TABLE 12.20 Swaption Volatilities Input for Exercise 12.10

Option Expiry	Swap Tenor									
	1	2	3	4	5	6	7	8	9	10
1	16.000%	14.700%	13.700%	12.800%	12.400%	12.000%	11.700%	11.500%	11.300%	11.100%
2	17.700%	15.100%	13.700%	13.100%	12.600%	12.100%	12.000%	11.800%	11.600%	
3	17.100%	15.100%	14.000%	13.200%	12.600%	12.500%	12.400%	12.200%		
4	17.200%	15.000%	13.800%	12.900%	12.800%	12.700%	12.500%			
5	16.200%	14.200%	13.100%	12.700%	12.800%	12.600%				
6	14.800%	13.300%	12.900%	12.900%	12.700%					
7	14.400%	13.400%	13.200%	12.900%						
8	14.900%	13.700%	13.100%							
9	14.000%	12.900%								
10	13.000%									

Credit Risk

The field of credit risk management has undergone major transformations over the past two decades. Traditional commercial bank lenders, whose focus used to be almost exclusively on the analysis of individual borrowers with a small dose of limits to avoid excessive concentration in a region or industry, have increasingly viewed overall portfolio management as a major part of their function. This has opened the door to rapid growth in the use of quantitative risk management techniques. At the same time, the introduction of an array of vehicles for transferring credit risk between creditors—the increased use of loan sales, loan syndication, and short sales of bonds, along with the introduction of many varieties of credit derivatives, asset-backed securities, and collateralized debt obligations (CDOs)—has served as a tool for portfolio management.

Over the same time period, many new players have become active participants in credit risk markets. While there have always been nonbank investors in corporate bonds, such as insurance companies, pension funds, and mutual funds, the variety of new instruments available for investors in credit risk—asset swaps, total return swaps, credit default swaps (CDSs), CDOs—has both introduced new investors, such as hedge funds, and increased the participation of existing investors.

In looking at the principles guiding credit risk management, one sees a genuine dichotomy between the views of traditional commercial bank lenders and the views of many nonbank investors. Investors who focus primarily on liquid corporate bonds and CDSs view risk management on these instruments as no different from market risk management of equity or interest rate positions—the general principles of Section 6.1.1 would apply, with emphasis on stop-loss limits, liquidation of positions, timely marking to market, and use of value at risk (VaR) and stress testing to assess liquidation risk. Traditional commercial bank lenders, with many loans to creditors whose debt has little liquidity and with large positions of illiquid size to creditors whose debt does have liquidity, see little value in such short-term

views of risk and concentrate instead on long-term (multiyear) analysis of portfolio risk.

This dichotomy of views relates back to the discussion in Section 1.2, with the credit risk of commercial bank lenders looking like actuarial risk, requiring an approach more like the one we've outlined in Sections 6.1.2 and 8.4. Caught in the middle are investors who have hybrid exposure to liquid and illiquid names—they need to use a mixture of short-term market risk management techniques for their more liquid risks and long-term portfolio analysis for their less liquid names. Among the players caught in the middle are market makers in over-the-counter derivatives, who almost always have a customer mix of counterparties with both liquid and illiquid debt.

The approach in this chapter is to start with the short-term risk management of liquid positions in Section 13.1, then to turn to long-term portfolio risk management in Section 13.3. In between, Section 13.2 looks at non-market-based methods for the internal analysis of single-name credit instruments. This topic is important as critical input to the portfolio models of Section 13.3, as a vital supplement to the techniques of Section 13.1 for names with good but limited liquidity, and as a fundamental element in trading models even for the most liquid names. Finally, Section 13.4 looks at the risk management of multiname credit derivatives such as CDS indexes and CDOs, which require a challenging mix of the portfolio management techniques of Section 13.3 and the more market-based approach of Section 13.1, a challenge that much of the financial industry badly failed in the 2008 crisis. The important topic of the management of credit risk for derivatives counterparties is placed in a separate chapter, Chapter 14, which will draw heavily on the conclusions of this chapter.

13.1 SHORT-TERM EXPOSURE TO CHANGES IN MARKET PRICES

When dealing with sufficiently liquid debt, credit instrument risk management can be designed to look very similar to interest rate risk management, but there are some important differences. As with interest rate risk management, a good part of the challenge is coming up with a unifying principle for combining the risks of many different types of instruments with a wide variety of terms and conditions. As with interest rate risk management, the key tool will be a focus on cash flows as a unifying principle (refer back to the start of Section 10.2). This principle does not work as cleanly for credit instruments as it does for interest rates, but with some modification it will still be able to serve.

We will model our discussion in this section closely on our interest rate discussion in Chapter 10. Section 13.1.1 looks at the variety of credit instruments, Section 13.1.2 looks at the mathematical models for valuing credit instruments, and Section 13.1.3 examines the design of risk reports.

13.1.1 Credit Instruments

13.1.1.1 Bonds and Asset Swaps The market for corporate bonds has been around for a long time, and these instruments are generally well quoted for certain firms. It has always been advantageous for companies seeking capital to issue bonds, partly because of resulting tax advantages, and also not to dilute the ownership in the company by issuing too much equity. Most of the time, corporate bonds are fixed-rate bonds, because this is what most investors in bonds prefer, even though many companies prefer to borrow at a floating rate, generally indexed to the London Interbank Offered Rate (LIBOR) (companies wishing to exchange floating-rate payments they want to make for the fixed-rate payments required on their bonds are a major source of demand for interest rate swaps). Most investors in corporate bonds share the following three characteristics:

1. They have cash to invest.
2. They are willing to take on credit risk, because they have a favorable view of the credit prospects of a particular firm or set of firms.
3. They are willing to take on rate risk or have a longer-term investment horizon and so view locking into long-term rates desirable.

Some investors are interested only in the first two features because they don't necessarily want to take a position with a view on rate risk. This is why *asset swaps* were created. An asset swap is a combination of a corporate bond and an interest rate swap contract that swaps the bond's coupon into a floating payment. So the purchaser of an asset swap will receive a fixed spread as payment so long as the bond does not default. But an even larger market developed for a purer form of credit-linked instruments, the *credit default swap* (CDS), that isolates credit risk without either of the other two aspects of corporate bonds.

13.1.1.2 Credit Default Swaps Credit default swaps were created in the 1990s. Their definition is very simple. While there is no default on the underlying, the protection provider receives a fixed spread payment on a regular basis (for example, every six months) from the protection buyer. If there ever is a default during the lifetime of the contract, the protection seller will pay the protection buyer the full par value of the bond. Since the protection seller

will then only be able to recover the value of the bond less loss given default (LGD), the seller will have a loss equal to the par value times the loss given default rate. So the protection seller is in exactly the same financial position as the buyer of an asset swap, receiving fixed spread payments if there is no default, losing the par amount times the loss given default rate if there is a default.

A CDS is meant to look like an asset swap, but without the need to invest cash. While this feature makes it very attractive to some investors looking to take on credit risk, it is an even more important product for investors with a negative view of a firm's credit or who are seeking protection against a firm's default. These investors previously could only achieve the position they desired by selling short a corporate bond. But the market for borrowing corporate bonds is extremely thin and expensive. The advent of the CDS, like any new forward market, provides far greater liquidity to those wishing to take short positions. (You might wonder why an investor seeking protection against a firm's default could not just sell the asset causing this exposure. But not all assets exposing an investor to losses when a firm defaults are as easy to sell as a corporate bond. Some may be difficult to sell, such as bank loans and extensions of trade credit; others may be impossible to sell, such as counterparty credit exposure on derivatives.) It also provides far greater liquidity for those wishing to express relative value views that one set of credit spreads will widen relative to another set.

The growth of the CDS market has been explosive, growing at a rate of about 100 percent per year in many years. The most troublesome issue in the creation of the CDS market has been difficulties in deciding on a settlement mechanism in the event of default. First, since payoff by the protection seller only occurs in the event of a default, exact definition of a default event must be agreed upon. Does default mean any delay in a scheduled payment of the borrower or only one of a particular magnitude? Is a formal declaration of bankruptcy a necessity? What happens if the terms of the borrower's debt are voluntarily renegotiated with creditors? (And how can you tell how voluntary it has been? The 2011 and 2012 experience with renegotiation of Greek government bonds has been a particularly worrisome example; see the *Economist* article "Fingers on the Trigger" of June 2, 2011.) Second, how should the amount owed by the protection seller to the protection buyer in the event of default be determined, and should this determination involve physical settlement or cash settlement? Third, what becomes of a CDS when the reference firm ceases to exist through merger or acquisition? Multiple solutions have been proposed to these issues with many different variants incorporated into documentation of individual deals. This is a particular headache for market makers in CDSs, who must be certain that transactions that seem to offset one another in terms of tenor and reference

entity actually do offset one another when contractual details of settlement procedure are considered.

Protection sellers would prefer that the debt instruments used for settlement be as narrow as possible, preferably the single most liquid bond issued by the company. But CDSs with such a narrow class of deliverables have led to severe problems in settlement, with protection buyers having to scramble to purchase a deliverable bond, resulting in driving up the price of that bond so high that it is close to par—the resulting profit between the purchase price and sale to the protection seller at par has not been nearly enough to compensate for actual default losses on which protection was sought. (For more details, see the articles from the *Economist*: “Is There Money in Misfortune?” July 16, 1998, and “Of Devils, Details and Default,” December 3, 1998.) It has now become much more common to define a broad class of deliverables, even including much less liquid credit instruments such as bank loans and trade credit. This makes it far easier for the protection buyers, since they can often deliver the actual credit instrument on which they were seeking protection and, in any case, have a wide choice of instruments to deliver. But this has made settlement more difficult for the protection seller, both because of lower liquidity of the instrument being delivered and because the protection buyer’s choice of deliverable instrument gives the buyer a cheapest to deliver option, comparable to the cheapest to deliver option into the Treasury bond future, referenced in Section 10.1.4. Both these effects cause protection sellers to demand higher credit spreads than they would otherwise. We discuss issues of relative pricing between bonds and CDSs in Section 13.1.2.3.

Some of the impact on market prices of credit protection buyers scrambling to acquire deliverable instruments can be eased by a cash settlement provision defined in terms of quoted prices for a specified bond. But the illiquidity of corporate bond markets, particularly in conditions following the default of the bond issuer, makes quoted prices suspect. This problem has been greatly exacerbated by the growth of multiname credit derivatives that have resulted in the notional value of CDS contracts referenced to a firm exceeding the total value of the firm’s debt. This has led to the establishment of auction procedures for establishing prices at which cash settlement can take place; see Helwege et al. (2009) for details concerning the auction mechanism and its implications for CDS market participants.

Another possible solution to this problem is to have a default swap with a fixed payment in the event of default, known as *binary credit default swaps*. This resolves the issue of how to determine payment, but may not be a good fit to the risk needs of a holder of a bond or loan. Suppose I am holding a \$100 million bond issued by ABC. I can buy a standard default swap on \$100 million notional. If the loss in the event of default turns out to

be \$20 million, it should pay roughly \$20 million. If it turns out to be \$80 million, it should pay roughly \$80 million. However, if I buy a default swap with a fixed dollar payout, I must make a guess as to the loss in the event of default and run the risk that I have either purchased too little protection or paid for too much protection. Default swaps with fixed payoffs are also harder to value since this requires an estimate of the probability of default, whereas a standard bond price is based on the product of the probability of default and loss given default. See Section 13.1.2.1 for further discussion of this point.

Default swaps, more than any other derivative instrument, have led to the concept of *legal basis risk* (see Section 3.2.1). A market maker may believe its risk on a default swap is matched exactly by the protection purchased through another default swap, only to find it has to make a payment under the contractual language of the first swap but receives nothing under the slightly different language of the second swap.

The International Swaps and Derivatives Association (ISDA), the industry group that sets standards for derivatives contracts, has made several valiant attempts to remedy the situation by standardizing contract wording. The resulting checklist of possible contract terms is a daunting document. Even so, new disputes continue to arise. ISDA has also established determination committees that rule on disputed issues such as the impact of mergers and whether a renegotiation is voluntary or forced. And ISDA has standardized the auction mechanism for determining prices at which cash settlement can take place. Any firm participating in this market needs to be thoroughly aware of all the relevant history of the disputes, of past actions of ISDA determination committees, and of ISDA auction procedures, and needs to be certain it fully understands the terms of the risk it has taken on. A good synopsis of the ISDA standards and the motivation behind them can be found in Gregory (2010, Section 6.3). For further background on the issues, see Henderson (1998), Falloon (1998), Cass (2000), Bennett (2001), Helwege et al. (2009), and the following articles from the *Economist*: “Is There Money in Misfortune?” (July 16, 1998); “Of Devils, Details and Default” (December 3, 1998); “Fixing the Holes” (August 12, 1999); “The Swaps Emperor’s New Clothes” (February 8, 2001); “The Tender Age” (April 20, 2006); and “Fingers on the Trigger” (June 2, 2010). The “Legal and Documentation” section of the ISDA website, found under the “Functional Areas” heading, provides many documents relating to contractual disputes and the findings of determination committees (www2.isda.org/functional-areas/legal-and-documentation).

The total return swap, which we encountered in Section 10.1.7, is another derivative instrument that can be structured for investors who want to take on credit risk without putting up cash. Unlike the CDS, which is designed to look like an asset swap, the total return swap is designed to look

like a straight investment in a corporate bond. The mechanics are that the investor enters into a swap in which he receives all of the coupon payments from the bond and any change in the bond price (positive or negative, so he may owe payments) and pays an amount equal to LIBOR times the par amount of the bond. So cash flows are very similar to borrowing at LIBOR and investing in the bond, but with the advantage that the counterparty to the total return swap can use it to create a short position in the bond and thereby express a negative view on the credit or protect a credit exposure. Total return swaps have proved to be far less popular instruments than the CDS, perhaps because asset swap positions are more sought after than fixed-rate corporate bond positions (for those investors not willing to put up cash) and perhaps because the reliance on a single bond raises settlement issues unfavorable to the investor similar to a CDS with a single deliverable.

13.1.2 Models of Short-Term Credit Exposure

In Section 10.2, we were able to base all modeling of interest rate risk on a single principle, that the value of each individual cash flow that is bundled together in an interest rate contract can be determined independently of the value of any other cash flow bundled in that contract. We would like to use a similar principle for credit instruments, but run into three roadblocks, one having to do with the treatment of credit instruments in bankruptcy proceedings, the second due to the large convexity risks of credit instruments, and the third due to basis risk between bonds and CDSs.

Before we can address these issues, we first need a fundamental framework within which we can discuss credit risk. Ultimately, the cost of credit risk must be based on expectations and uncertainty concerning loss from default. Without the possibility of default, credit instruments would just be priced based on the risk-free discount curve. Default loss can be analyzed into three components as follows:

$$D(I) = P_D(B) \times L_D(I) \times A_D(I) \quad (13.1)$$

- where
- I = a credit instrument
 - B = the borrower on the instrument
 - $D(I)$ = the default loss on the instrument
 - $P_D(B)$ = the probability that the borrower will default
 - $L_D(I)$ = the percentage loss on the instrument conditioned on default
 - $A_D(I)$ = the amount that will be owed on the instrument conditional on default

We use $P_D(B)$ instead of $P_D(I)$ because cross-default legal provisions come close to guaranteeing that a borrower will default on either all or none of its debt.

For liquid instruments like bonds and CDSs, $A_D(I)$ is a fixed amount—the amount of currency borrowed—so we need only concern ourselves in this section with the $P_D(B) \times L_D(I)$ term. In Section 13.2, when we look at illiquid instruments, we will encounter cases (lines of credit and counterparty credit risk) for which $A_D(I)$ can vary.

Market prices of credit instruments cannot distinguish the effects of default probability and loss given default—in other words, you can extract information from market prices on $P_D(B) \times L_D(I)$, but cannot distinguish between $P_D(B)$ and $L_D(I)$. To get a clear view of the entanglement of default probability and loss given default in market prices for credit instruments, let's consider the simplest possible case. Suppose company XYZ has a two-year zero coupon bond that is trading at \$85.50 per \$100.00 par amount, while a two-year zero coupon government bond is trading at \$90.00 per \$100.00 par amount. The \$4.50 haircut on the corporate bond implies that the market is pricing the bond as if the expected loss from default over a two-year period would be 5 percent ($\$90 \times 95\% = \85.50). However, this loss could consist of $P_D(B) = 5\%$, $L_D(I) = 100\%$; $P_D(B) = 10\%$, $L_D(I) = 50\%$; or any other combination that results in $P_D(B) \times L_D(I) = 5\%$. This inability of splitting probability of default from expected loss given default will need to be kept in mind when we discuss utilizing market data in internal models of credit risk in Section 13.2, and in models of portfolio credit risk in Section 13.3.

13.1.2.1 Impact of Bankruptcy Law The ability to value all cash flows received on the same date using the same discount factor is a vital assumption in the methodology used to maximize liquidity in the forwards markets, as discussed in Section 10.2. The reason this assumption breaks down for credit instruments relates to provisions of bankruptcy law. In almost all jurisdictions, the claim for a two-year coupon due on a five-year bond is not the same in bankruptcy as the claim for the same amount of principal on a two-year bond. The common rule for bankruptcy is that the holder of a bond or loan can make a claim on the principal, but not on any coupon interest. Offsetting this loss of interest that can be claimed is the ability to call for immediate payment of principal, regardless of maturity.

For bonds or loans trading close to par—that is, the coupon on the bond is close to the current par coupon—the advantage and disadvantage almost cancel out. A five-year bond loses five years' worth of coupons, but can accelerate principal due by five years, while a two-year bond loses only two years' worth of coupons, but can accelerate principal due by only two years. The par coupon can be thought of as the rate of interest that exactly compensates an

investor, at current market discount factors, for deferral of receiving principal; therefore, foregoing coupons on the par coupon bond will precisely offset the acceleration of principal. For similar reasons, a floating-rate bond or loan, whose coupon resets to current market levels, should have the advantage of principal acceleration closely balance out the loss of coupon payment.

However, for bonds or loans selling at a premium, either because of a fixed coupon higher than the current par coupon or a floating rate at a positive spread to current market levels, the bankruptcy rules will cause more of a loss on default than that felt by a par bond or loan. Conversely, a bond or loan selling at a discount will experience less of a loss on default than that experienced by a par bond or loan. As a result, the rule that all cash flows on the same date are equivalent, regardless of what package they are part of, breaks down. A coupon payment is worth more in default if it is packaged as part of a discount bond than a coupon payment for the same date that is packaged as part of a premium bond.

Exercise 13.1 familiarizes you with the mathematics needed to deal with this situation. The **CreditPricer** spreadsheet used in the exercise takes as input the current risk-free zero coupon curve, an assumed set of annual default rates, and an assumed loss given default rate, and computes the resulting par curve for a corporate bond and resulting spreads to the risk-free par curve. The calculation looks at the value of payments received if no default occurs plus the accelerated principal payments received if default occurs. You will also find this calculation explained and illustrated in Hull (2012, Section 23.4). The exercise demonstrates that spreads to the risk-free par curve will differ for differing assumptions of loss given default. This shows that it is not just the product $P_D(B) \times L_D(I)$ that matters in this case, but also the individual components, since the value of the principal acceleration depends on the loss given default assumption. The exercise also shows you how to use the same spreadsheet to solve for market-implied default rates based on an observed par curve and an assumed loss given default rate. It further shows that if prices are available for several coupons with the same maturity, then information about the split between $P_D(B)$ and $L_D(I)$ can be extracted.

One issue in which the difficulty in splitting market quotes into probability of default and loss given default components involves binary credit default swaps. The price of a binary CDS should just be the cost of a standard CDS divided by $1 - L_D(I)$, since the standard CDS will pay $1 - L_D(I)$ dollars for each dollar of principal of the CDS in the event of default, but the binary CDS pays the full principal in the event of default.

13.1.2.2 Convexity of Credit Instruments To illustrate the difficulty that convexity poses for credit risk management based on short-term exposure to

market prices, consider the following simple example. (By contrast, convexity has little impact on interest rate instruments; see Section 10.4.) Consider two obligations of company XYZ: a two-year zero coupon bond and a 10-year zero coupon bond. Assume that a risk-free two-year zero is trading at \$90 per \$100 par value and a risk-free 10-year zero is trading at \$60 per \$100 par value. If the expected loss from default for XYZ is roughly 1 percent a year, we would expect to see a haircut for the two-year zero of $\$90.00 \times 2\% = \1.80 and a haircut for the 10-year zero of $\$60.00 \times 10\% = \6.00 . If market confidence in XYZ worsened slightly, expected loss from default might rise from about 1 percent a year to about 1.1 percent a year, resulting in a haircut for the two-year zero of $\$90.00 \times 2.2\% = \1.98 and a haircut for the 10-year zero of $\$60.00 \times 11\% = \6.60 . Therefore, the two-year zero has moved by $\$1.98 - \$1.80 = \$0.18$ and the 10-year zero has moved by $\$6.60 - \$6.00 = \$0.60$, a ratio of $\$0.60/\$0.18 = 3.33$, which could also be derived as a ratio of the durations multiplied by the present values: $(10 \times \$60)/(2 \times \$90)$.

If you want to hedge against small moves in a credit spread, you would sell short \$30 million 10-year bonds against a long position of \$100 million two-year bonds. But what happens if XYZ defaults? You have losses on \$100 million balanced by gains on only \$30 million. The right ratio for hedging short-term market movements is an extremely poor ratio for hedging default, due to the severe convexity. The higher the $P_D(B)$ component of the in $P_D(B) \times L_D(I)$ product, the greater the probability of default, and the more significant the convexity risk. For large moves that do not go all the way to default, as might be associated with a credit downgrade, a mismatch in correct hedging ratios will still occur, but it will be less severe. This example demonstrates that risk management utilizing short-term exposures to changes in market price is not sufficient by itself; it needs to be supplemented by an analysis of ultimate default risk.

13.1.2.3 CDS-Bond Basis Risk To understand the basis risk between CDS and bonds, we must first start with the theoretical arbitrage relationship between them and then see what factors might alter it. We will use a simple illustrative example of the arbitrage relationship, a fuller discussion of which can be found in Duffie and Singleton (2003, Section 8.3). Under ideal circumstances, the spread above LIBOR on a floating-rate bond issued by a corporation (call this spread S) ought to be equal to the CDS spread for the same maturity (call this spread C). If the purchaser of the bond also purchases a CDS of the same tenor, his return if the issuer does not default is $\text{LIBOR} + S - C$ each year plus the return of his principal. If the issuer does default, he can exchange the bond for par under the terms of the CDS. Since the investor always gets back his principal, he has an investment with no credit risk

on which the return ought to be LIBOR; hence $\text{LIBOR} + S - C$ should equal LIBOR, so S should equal C .

In practice, very few companies issue floating-rate bonds, but an asset swap can be used to turn a fixed-rate bond into a close approximation of a floating-rate bond. So the CDS spread ought to equal the spread over LIBOR that the fixed-rate coupon can be exchanged for in the interest rate swap market, which is the spread between the coupon rate and the swap rate for the bond's tenor.

So why should the actual basis between a CDS spread and a bond's coupon spread to the swap rate be different than zero? Partly it is because the asset swap is not a perfect substitute for a floating-rate bond, and partly it is because of features of the CDS that have not been accounted for in the above idealization, such as the cheapest-to-deliver option discussed in Section 13.1.1.2. The seminal article on the CDS-bond basis is Lehman Brothers' "Explaining the Basis: Cash versus Default Swaps" by O'Kane and McAdie (2001). It analyzes many factors that potentially could make the CDS spread greater than the bond spread ("increase the default swap spread" in the terminology of the paper) or make the bond spread greater than the CDS spread ("decrease the default swap spread"). A summary can be found in O'Kane (2008, Chapter 5).

Other sources worth consulting are:

- DeWit (2006). DeWit's discussion of factors driving the basis in Section 2 of his paper leans heavily on O'Kane and McAdie, but he adds some analysis and a very comprehensive set of footnotes with references to both empirical and theoretical articles. Table 6 gives a concise comparison of empirical research on the size of the basis, which centers around 5 to 10 basis points with CDS spreads higher than bond spreads. DeWit states: "While we define 14 different economic basis drivers, it is our understanding that four of them (i.e. the CDS cheapest to deliver option, difficulties in shorting cash bonds in a context of structural demand for protection, relative liquidity in segmented markets, and synthetic CDS issuance) are the main determinants of the CDS-bond basis."
- Hull, Predescu, and White (2004) also present empirical evidence that supports the same conclusion as DeWit's Table 6.
- Duffie and Singleton (2003, Section 8.3) analyzes the CDS-bond basis. They are less inclusive than O'Kane and McAdie in considering all possible influences, but are worth looking at for the depth of their analysis of the impact of the difficulty in shorting bonds.
- The historical relationship of CDS trading about 5 or 10 basis points higher than bond spreads was severely disrupted by the 2007–2008 crisis, with spreads going negative by 250 basis points for

investment-grade firms and by 650 basis points for high-yield names; see Bai and Collin-Dufresne (2011) for a detailed discussion of both market behavior and possible causes. The two major drivers of this disruption appear to be:

1. **Funding cost.** Many holders of cash bonds were now funding at substantially higher rates than LIBOR. While the high bond spread relative to CDS spreads would then seem to offer an arbitrage opportunity to those who could still fund at LIBOR, there may have been little appetite for such arbitrage in the current environment (a Reuters article “Popular US Credit Trade Turns Sour” of December 13, 2007, stated that “lack of financial balance sheet capacity and a general unwillingness to lend has prolonged the negative basis”).
2. **Heightened concern for counterparty risk.** If a CDS is not fully collateralized, the buyer of CDS protection may be unwilling to pay the full cost of default risk.

13.1.3 Risk Reporting for Market Credit Exposures

A good starting point for risk reporting of market credit risk is to closely parallel the reporting guidelines for forward risk given in Section 10.4. As with forward risk, key questions for market credit risk involve selection of maturity buckets and selection of summary statistics, such as exposure to a parallel shift in the credit spread curve and exposure to linear tilt of the credit spread curve. A measure of credit spread duration is calculated in close parallel to the duration measure for rates and serves as an alternative to the value of a basis point shift in the credit spread curve. Because economic events that have an impact on default probabilities often impact the credit spreads of more vulnerable firms more than those of higher credit quality, many firms utilize a measure of percentage change in credit spread as an alternative to or supplement to a measure of impact of a parallel shift in credit spread. For example, a measure of a 5 percent increase in credit spreads would add together the impact of a 5-basis-point increase in a credit that currently has a 100-basis-point credit spread with the impact of a 25-basis-point increase on a credit that currently has a 500-basis-point credit spread.

There are two key added factors for the measurement of credit spread exposures relative to rate exposures. One is that credit spreads have far more characteristics to be taken into account when grouping exposures—geographic, industry, and credit quality. The second is the importance of price jumps and convexity for credit spreads, which is of little importance for forwards.

Let's look at grouping characteristics first. As with equity spot risk in Section 9.3, grouping of exposures and limits by geography and industry make sense. For corporate credit, equity exposure and credit spread exposure are two aspects of risk exposure to corporations, so the groupings used should be very similar. All levels of management should see total net credit exposure along with exposure to major geographic regions (e.g., United States, Western Europe, developed Asia, emerging markets) and major industry groups, while lower levels of management should see more detailed net credit exposures by country and specific industries. Reporting and limits for exposure to individual borrowers is also needed. Finally, grouping of exposures and limits are needed for credit quality, with rating agency grades, such as Aa, A, Baa, and Ba, often being used.

Given the large impact of convexity on credit spread exposures, as discussed in Section 13.1.2.2, it is very important to have measures and limits that capture this risk. Measures and limits that capture default risk will be discussed in Section 13.2. For large credit spread shifts, the most intuitively appealing are measures of and limits on the amount that can be lost in the event of very large shifts in credit spread that might be associated with a major shift in the economic environment. So you might have a measure of exposure to a 1 percent shift in credit spreads to control for ordinary market moves, and a measure of exposure to a 10 percent shift in credit spreads to control for a large move. There is an obvious parallel to the delta and convexity limits on options positions, discussed in Section 11.4.

13.2 MODELING SINGLE-NAME CREDIT RISK

Models of single-name credit risk are important for several reasons:

- If you have exposure to a single-name credit instrument for which you can't obtain a liquid market price, you will need a model to value it.
- Even when you can obtain a liquid market price, comparison to a modeled price can be useful in informing trading decisions.
- Single-name credit instrument models serve as important inputs to credit portfolio and multiname credit instrument models. Since credit portfolios and multiname credit instruments need to be evaluated over long-term horizons, just having a liquid price for constituent pieces is not adequate—a model of possible price evolution is also required.

The key element in modeling any single-name credit instrument is modeling expected default loss, since, absent default loss, the instrument is just

an interest rate instrument, whose modeling we have already studied in Chapter 10. Referring back to Equation 13.1 in Section 13.1.2, the default loss on a credit instrument can be written as

$$D(I) = P_D(B) \times L_D(I) \times A_D(I)$$

that is, the product of probability of default, loss given default, and the amount that will be owed conditional on default. The best way of organizing the modeling of default loss is to model these three components separately. Partly, this is just an aid to clear thinking. Partly, it is motivated by probability of default being a function of the borrower, independent of the instrument, while the other two components are instrument dependent. And partly, this is a matter of expertise: Those who are most expert in modeling loss given default may be lending officers with experience in loan work-outs of borrowers threatened by bankruptcy, while probability of default may best be modeled by those with direct knowledge of a particular firm or industry.

Our discussion in this section is accordingly separated into sections on estimating probability of default (13.2.1), estimating loss given default (13.2.2), and estimating amount owed conditional on default (13.2.3). Section 13.2.4 looks at information relative to defaults that can be derived from prices for equity and equity options utilizing an option-theoretic approach, a topic that cuts across both probability of default and loss given default.

13.2.1 Estimating Probability of Default

Default probability is the most critical and most intensely studied of the components of single-name credit risk. Almost all firms that deal in credit risk instruments will want to form their own assessments of default probability (only if credit instruments are only a small portion of the investment portfolio and are almost all liquid might a firm be satisfied with just basing this assessment on input from an outside service). Firms with heavy investment in credit instruments, such as traditional banks, will devote considerable resources to their own determination of default probability. But all firms should be aware of and make use of independent assessments of default probability, both as input to their own judgments and as reality checks.

This is particularly true when the credit valuation is for a business with which the lending firm has a long relationship and detailed, intimate knowledge of the business's management and operations. Caution needs to be exercised in such cases, since close, longtime relationships can breed complacency and a reluctance to acknowledge unwelcome changes. It is important to have an internal review mechanism in which internal credit ratings that show lower default probabilities than agency default probabilities or

those derived from the models or markets are challenged. The review mechanism needs to be run by people with good experience in the credit area but who don't have the direct client involvement that may lead to complacency.

We have divided up the possible sources of independent default probabilities into two broad categories. The first, and most widely used, is direct comparison to rating agency evaluations. The second is the use of statistical modeling that may take as input borrower-specific information, gauges of the broad economy, and market prices. This last category will lead us into the area of option-theoretic models, discussed in Section 13.2.4. We discuss these two categories in turn.

13.2.1.1 Rating Agency Evaluations The primary output of rating agency evaluation of individual borrowers is a letter grade. Translating letter grades into default probabilities requires some analysis, but the ratings agencies provide an abundant amount of historical data that can be utilized to make this conversion. While all of the rating agencies provide such historical data, I will, for convenience, make all my references in this section to the Moody's data, which is updated regularly and appears to be easily available to the public on the web, through the National Association of Insurance Commissioners at www.naic.org. All the data quoted and used in tables comes from Moody's (2011a).

It should be noted that any use of historical rating agency data to translate from current agency ratings to default probabilities does rest on the assumption that the ratings assignment process has been reasonably stable and consistent over time. Arguments for this being a reasonable assumption can be found in the section "The Rating Process" in de Servigny and Renault (2004, Chapter 2), for example, "The criteria according to which any assessment is provided are very strictly defined and constitute the intangible assets of ratings agencies, accumulated over years of experience. Any change in criteria is typically discussed at a worldwide level."

The translation of rating agency grades to default probabilities generally starts with transition matrices that show the probability over a fixed time period that a credit rated in one category at the beginning of the period will default during the period or will transition to another credit rating category at the end of the period. Tables 13.1 and 13.2 show a sample one-year transition matrix and a cumulative transition matrix that only looks at default. Rating agencies also publish matrices covering many different transition periods (for example, two-year transitions, three-year transitions, and so on); matrices with finer credit rating graduations; and matrices based on subsets of this historical data.

There are a variety of approaches in using this data to convert agency ratings into default probabilities. Here are some of the major differences. As the one-year transition matrix in Table 13.1 shows, some borrowers who

TABLE 13.1 One-Year Transition Matrix, 1970–2010

Initial Rating	Rating at Year-End										
	Aaa	Aa	A	Baa	Ba	B	Caa	Ca-C	Default	Withdrawn	
Aaa	87.40%	8.63%	0.60%	0.01%	0.03%	0.00%	0.00%	0.00%	0.00%	3.34%	
Aa	0.97%	85.62%	7.97%	0.36%	0.05%	0.02%	0.01%	0.00%	0.02%	5.00%	
A	0.06%	2.69%	86.76%	5.27%	0.49%	0.11%	0.03%	0.00%	0.05%	4.53%	
Baa	0.04%	0.18%	4.52%	84.52%	4.11%	0.78%	0.17%	0.02%	0.18%	5.48%	
Ba	0.01%	0.06%	0.37%	5.64%	75.76%	7.24%	0.53%	0.08%	1.10%	9.21%	
B	0.01%	0.03%	0.13%	0.34%	4.76%	73.52%	5.77%	0.67%	4.23%	10.54%	
Caa	0.00%	0.02%	0.02%	0.14%	0.46%	8.26%	60.09%	4.10%	14.72%	12.18%	
Ca-C	0.00%	0.00%	0.00%	0.00%	0.32%	2.37%	8.88%	36.27%	35.45%	16.70%	

Source: Moody's (2011a, Exhibit 27).

TABLE 13.2 Cumulative Default Rates, 1970–2010

Initial Rating	Number of Years							
	1	2	3	4	5	7	10	15
Aaa	0.00%	0.01%	0.01%	0.04%	0.10%	0.24%	0.49%	0.92%
Aa	0.02%	0.06%	0.10%	0.18%	0.27%	0.44%	0.62%	1.26%
A	0.05%	0.18%	0.36%	0.55%	0.76%	1.24%	2.14%	3.66%
Baa	0.18%	0.51%	0.93%	1.43%	1.95%	3.03%	4.90%	8.85%
Ba	1.16%	3.19%	5.60%	8.15%	10.45%	14.44%	20.10%	29.70%
B	4.47%	10.43%	16.33%	21.51%	26.17%	34.72%	44.57%	56.35%
Caa–C	18.16%	30.20%	39.71%	47.32%	53.77%	61.18%	72.38%	76.16%

Source: Moody's (2011a, Exhibit 34).

receive a rating at the beginning of a period are no longer tracked by the end of the period because they have asked the agency to withdraw its rating. In projecting transition probabilities, a choice must be made between assuming that a request for rating withdrawal indicates anticipation of a downgrade and assuming that a request for rating withdrawal carries no information content or some intermediate assumption (see de Servigny and Renault 2004, Appendix 2A).

- Trade-offs exist between using multiyear default data based on the direct observation of cumulative default rates versus generating multiyear cumulative default rates by the matrix multiplication of one-year transition matrices. The direct use of cumulative default rates suffers from a diminishing data pool for longer tenors and greater potential inaccuracy from withdrawn ratings (firms whose ratings are no longer tracked) (see Gupton, Finger, and Bhatia 1997, Section 6.3.2). Matrix multiplication assumes a Markovian process, where no serial correlation exists between transitions. Alternatively, it could be desirable to derive one-year transition matrices that are consistent with observed longer-term cumulative default and transition behavior (see Gupton, Finger, and Bhatia 1997, Section 6.4). However, there is data suggesting that serial correlation between transitions does exist (see Bahar and Nagpal 2000).
- Default probabilities for tenors that fall in between those for which transition matrices are published can be interpolated (see de Servigny and Renault 2004, Appendix 2A).
- Rating agencies are very frank about the fact that their ratings represent through-the-cycle as opposed to point-in-the-cycle ratings (see de

Servigny and Renault 2004, “Time Horizon for External Ratings” in Chapter 2). Ratings are not adjusted just because a movement from an expansionary phase of the economic cycle to a recession increases the likelihood of defaults. Conversion to default probabilities that accurately reflect the current economic environment can be made using data such as that presented in Table 13.3, which shows how five-year default probabilities differed by starting year. This data could then be correlated with information on the stage of the economic cycle each five-year period represents.

- Default rates and transition matrices could be adjusted for the current stage in the economic cycle, based on historical observation of differences during recession and growth periods.
- Large lending firms may have their own internal data on defaults and transitions that they may want to use to supplement the publicly available data that comes from the ratings agencies. However, even if this data has been well maintained, a trade-off exists between using data that is more relevant to the particular class of borrowers who are customers of a particular firm and the loss of accuracy that comes from the utilization of a smaller sample.
- If default and transition data is available broken out by country and industry, this could be used to refine the data available from the ratings agencies. One criticism of ratings agency data is that they are largely based on experience with U.S. firms; see the sections “Quality of Transition Matrices over Time and Region” and “Industry and Geography Homogeneity” in de Servigny and Renault (2004, Chapter 2). However, the same points about small data samples raised in the last bullet may be relevant here.
- The tables and discussion in this section have referred only to corporate borrowers. The rating agencies publish comparable transition matrices for sovereign government borrowers (see Moody’s 2011b) and other government borrowers, such as municipalities (see Moody’s 2010).
- Default and transition data from different sources can be blended, such as averaging S&P and Moody’s data, or rating agency and private data.

A frequently expressed concern is that agency credit ratings are not updated often enough to fully reflect the probability of default. It reflects the nature of the rating process, which, because of the serious consequences to a firm’s financial health a ratings change can entail, requires that changes be thoroughly deliberated and well documented. This may supply the motivation to supplement this source of independent default probabilities with one of the two other sources we will now discuss.

TABLE 13.3 Five-Year Default Rates
based on data from Moody's (2011) Exhibit 42

	Aaa	Aa	A	Baa	Ba	B
1970	0.000%	0.000%	0.413%	1.412%	7.058%	22.586%
1971	0.000%	0.000%	0.736%	1.079%	3.827%	3.846%
1972	0.000%	0.000%	0.336%	2.061%	3.549%	7.018%
1973	0.000%	0.000%	0.000%	1.941%	3.206%	3.774%
1974	0.000%	0.000%	0.000%	1.784%	4.050%	7.143%
1975	0.000%	0.000%	0.000%	0.819%	3.847%	9.566%
1976	0.000%	0.000%	0.000%	0.910%	3.731%	4.082%
1977	0.000%	0.000%	0.000%	0.594%	2.903%	15.249%
1978	0.000%	0.810%	0.000%	1.406%	4.559%	22.612%
1979	0.000%	0.797%	0.576%	2.046%	5.980%	17.423%
1980	0.000%	0.000%	0.832%	1.706%	8.730%	28.033%
1981	0.000%	0.000%	0.267%	3.365%	11.708%	27.985%
1982	0.000%	0.000%	1.118%	2.477%	18.565%	30.063%
1983	2.395%	0.487%	0.262%	3.776%	14.205%	28.702%
1984	1.449%	1.832%	1.504%	1.729%	18.146%	27.363%
1985	0.000%	0.789%	2.547%	2.876%	18.452%	30.568%
1986	0.000%	1.230%	1.865%	5.813%	20.319%	34.910%
1987	0.000%	0.394%	1.925%	5.006%	22.846%	40.224%
1988	0.000%	1.001%	1.431%	3.991%	22.673%	39.494%
1989	0.000%	0.618%	0.758%	3.289%	23.439%	43.744%
1990	0.000%	0.000%	0.000%	0.629%	18.292%	38.319%
1991	0.000%	0.316%	0.000%	0.275%	9.527%	33.874%
1992	0.000%	0.284%	0.000%	0.000%	2.495%	27.086%
1993	0.000%	0.000%	0.000%	0.547%	4.427%	19.208%
1994	0.000%	0.000%	0.000%	0.673%	4.989%	16.683%
1995	0.000%	0.000%	0.000%	1.583%	7.538%	16.651%
1996	0.000%	0.000%	0.122%	1.327%	8.513%	18.447%
1997	0.000%	0.000%	0.336%	2.472%	12.393%	26.515%
1998	0.000%	0.000%	0.540%	3.283%	13.918%	34.987%
1999	0.000%	0.000%	0.759%	3.100%	10.852%	37.442%
2000	0.000%	0.000%	0.903%	2.848%	6.281%	34.574%
2001	0.000%	0.000%	0.729%	2.791%	6.285%	30.406%
2002	0.000%	0.000%	0.355%	1.949%	6.490%	18.057%
2003	0.000%	0.000%	0.000%	0.329%	3.155%	9.995%
2004	0.000%	0.231%	1.355%	0.316%	3.903%	8.633%

(continued)

TABLE 13.3 (Continued)

	Aaa	Aa	A	Baa	Ba	B
2005	0.000%	0.210%	1.388%	1.986%	9.231%	20.057%
2006	0.000%	0.898%	1.496%	1.238%	9.873%	20.838%
Mean	0.104%	0.267%	0.610%	1.984%	9.729%	23.139%
StdDev	0.454%	0.441%	0.677%	1.349%	6.504%	11.394%
Max	2.395%	1.832%	2.547%	5.813%	23.439%	43.744%
Max – Mean/SD	5.041	3.549	2.860	2.838	2.108	1.808

13.2.1.2 Statistical Modeling The seminal concept in statistical modeling of default probabilities was Edward Altman's 1968 Z-score model that related probability of corporate default to firm-specific accounting ratios—the ratio to total assets of working capital, retained earnings, earnings before interest and taxes, and sales—and one market price, the market value of equity. Bohn and Stein (2009, Chapter 4) and Saunders and Allen (2010, Chapter 6) give a good exposition of the current state of these models.

Market prices can be used in statistical models of default probability in one of three ways. The first is the way Altman used the market value of a firm's equity in his Z-score model, as just an independent variable in a regression model or discriminant analysis. The second is to try to bring more theoretical structure to the relationship between equity market prices and default probability, the option-theoretic models we will examine in Section 13.2.4. The third is to try to find a structural relationship between bond and CDS market prices and default probability.

Linking bond and CDS market prices to default probability could be useful in several ways. A bank that is holding too much debt of a particular borrower to be able to consider using the CDS market to liquidate the risk and which therefore must manage the risk using a longer-term portfolio management approach would still be interested in finding out the default probability that is built into the market price—CDS spreads may reflect new information faster than the bank's internal review process and would be valuable as input to the internal process. Even when no liquid market exists for the bonds or CDSs of a particular borrower, it might be possible to construct an index of liquid bonds and CDSs for other borrowers related by similar characteristics (e.g., credit rating, nationality, industry), and deriving a default probability for this index could be similarly valuable input to the bank's internal review process.

There are two barriers that must be overcome in deriving default probabilities from market credit spreads. The first is the one discussed in Section 13.1.2, the inability to separate default probability from loss given

default (LGD). This would need to be addressed by making a reasonable assumption for the LGD and then deriving the default probability implied by the credit spread. The second barrier is the large difference between actual default probabilities and those implied by market rates, due largely to the systematic risk embedded in credit exposure (this will be discussed further in Section 13.4.4). This difference has been studied extensively over the past few years; good summaries are in Hull (2012, Section 23.5) and Amato and Remolona (2003). Actual default probabilities can be inferred from market-implied default probabilities based on observed historical relationships.

The downside to this latter approach is that changes in debt prices may reflect many factors other than changes in market sentiment about default probability; technical liquidity factors or changes in the willingness to take on systematic risk can dominate. And even when a borrower does have liquid bonds and CDSs, they may not be very liquid and may not provide an up-to-date assessment of market sentiment on the firm's credit risk. Stock prices are generally far more liquid and less subject to, though not immune to, being impacted by technical liquidity factors (and equity is certainly subject to the same buffering as debt by changes in willingness to take on systematic risk). The greater liquidity of stock prices is a major driving factor behind the use of the option-theoretic models for credit.

13.2.2 Estimating Loss Given Default

De Servigny and Renault (2004, Chapter 4) and Bohn and Stein (2009, Chapter 5) are good introductions to the general topic of estimating loss given default (LGD).

Statistical estimates of LGD have been published by the credit rating agencies. A few other published studies are available as well. De Servigny and Renault (2004, Chapter 4); Altman, Resti, and Sironi (2001, Appendix III.1); and Gupton, Finger, and Bhatia (1997, Chapter 7) offer good discussions of the public data available. Table 13.4 provides results from the Moody's study for bond defaults occurring from 1982 to 2010, as reported in Moody's (2011a). Distinctions are drawn based on the relative seniority of debt, with bank loans regarded as a separate seniority class from bonds. Published studies usually show recovery rates, which are 100 percent minus the LGD rate, but I have translated into LGD.

The measurement of historical LGD can be performed in two different ways. One is to observe the drop in market prices for an instrument about one month after the announcement of default, and is shown in the column labeled "Measured by Postdefault Trading Prices" in the table. The second is to track all cash eventually received in the settlement of claims and to present value these future receipts back to the date of default, utilizing a

TABLE 13.4 Comparison of Rates of Loss Given Default

Seniority Class	Measured by Ultimate Recoveries (1987–2010)	Measured by Postdefault Trading Prices (1982–2010)
First lien bank loans	19.7%	34.2%
Second lien bank loans		70.9%
Senior unsecured loans		52.2%
Senior secured bonds	36.5%	49.2%
Senior unsecured bonds	50.8%	63.3%
Senior subordinated bonds	70.6%	69.3%
Subordinated bonds	70.7%	68.7%
Junior subordinated bonds	81.6%	75.3%

Source: Based on Moody's (2011a, Exhibits 7 and 9).

discount rate that suitably reflects the uncertainty of recovery. This measure is shown in Table 13.4 in the column labeled “Measured by Ultimate Recoveries.” Gupton, Finger, and Bhatia (1997, Section 7.1) cite academic studies that conclude that the “bond market efficiently prices future realized liquidation values,” supporting a rough equivalence of these two methods. This conclusion is consistent with the data in Table 13.4. Bohn and Stein (2009, Chapter 6) cite a Moody's study by Varma and Cantor that “determined that the single B bond spread provided a reasonable proxy for the discount rate that, on average, equated” these two measures. Which measure is more relevant depends on usage. In the context of the management of liquid credit instruments in Section 13.1, postdefault trading prices would be more in line with the exit price approach for liquid instruments. Managers of less liquid credit portfolios would have more flexibility in deciding which method of recovery was more promising for each default event.

All losses should be expressed as a percentage of par, given that bankruptcy law uses par amount of the instrument as the basis for a claim (as discussed in Section 13.1.2.1). Volatility of LGD rates is an important issue for the credit portfolio simulations discussed in Section 13.3.2. Tables 4.4 and 4.5 of de Servigny and Renault (2004) display statistics on volatility of LGD rates by seniority class, showing standard deviations in the 25 to 35 percent range.

Parallel to our discussion on the estimation of the risk of default, firms may want to supplement published data on LGD with their own internal data. This is particularly an issue with non-U.S. debt and bank loans. Published data on loss given default is heavily weighted toward the U.S. market, but bankruptcy laws and procedures differ substantially by country and may thus be expected to impact recovery rates. Recovery rate has also been shown

to differ significantly by industry; see de Servigny and Renault (2004, Chapter 4) for data in Table 4.5 and discussion; in particular, de Servigny and Renault suggest that “what may appear as an industry effect may actually reflect differences in collateral quality offered by firms in various industries.” The lower loss given default rate on bank loans can be presumed to be due to the attention banks pay to the negotiation of security against default. However, this attention may vary between banks and, even within a bank, by loan type.

Firms putting together their own internal data on LGD must be careful in compiling the data on ultimate recoveries. Gupton and Stein (2005, Section 4.3.1) point to a 1999 Moody’s study “Debt Recoveries for Corporate Bankruptcies” by David Hamilton and Lea Carty showing that “15% of the value of recoveries for Senior Secured Loans came in the form of equity of the defaulted form.” Gupton and Stein then comment:

Since these payments with equity interests (e.g., common stock, preferred, and warrants) commonly do not trade, their value will be unclear and unrealized for years. When these equity values are eventually realized/known (often well past the write-off date), it would be atypical for a bank’s accounting system to track flows back to the original charge-off. When we assist clients in databasing their own institution’s LGD histories, we have always found it necessary to examine archived paper records. The full tracking of default resolution realized values (cash flows) has been far more informative than sourcing simply the accounting write-offs.

Economic modeling of LGD has not received as much attention as economic modeling of probability of default. Jacobs and Karagozoglou (2011), Altman and Kalotay (2010), and Bohn and Stein (2009, Chapter 5) each present economic models for LGD along with discussion of the relevant literature. Moody’s KMV has developed a commercial economic forecasting model for LGD; see Gupton and Stein (2005). Even when forecasts are based on the judgments of experienced credit managers, it is still advisable to be aware of the economic models, at least for sensitivity to the factors that have proved most important. Along with loan structure and ranking of collateral, Bohn and Stein find macroeconomic environment (state of the economy, industry) and firm leverage among the significant factors. Jacobs and Karagozoglou also find firm size to be significant. Gupton and Stein also utilize KMV’s distance-to-default measures (discussed in Section 13.1.4) for the firm, the industry average, and the geographic region average.

An issue that has drawn significant recent attention is the correlation between the occurrence of default and the rate of loss given default. This is the focus of a report submitted by Altman, Resti, and Sironi (2001) to the

International Swaps and Derivatives Association. This study finds significant negative correlation between the occurrence of default and recovery rate, which translates to a strong positive correlation between the occurrence of default and loss given default. This is not surprising on economic grounds, since an economic recession is likely to trigger more defaults while also negatively impacting the ability of a bankrupt firm to realize value on its remaining assets. This correlation has much the same effect as an increase in the level of correlation between defaults, since both result in more clustering of default losses. For example, if we're projecting the possible default losses for the next year, we might experience a good period for the overall economy that leads to few defaults and small losses on the defaults that do occur, or we might experience a recession that leads to many defaults and a high level of losses on these defaults. To the extent default losses cluster, it implies the need for added capital to guard against large losses, as discussed in Section 13.3.2, and a lower valuation of the senior tranches in CDOs, as discussed in Section 13.4.1.

13.2.3 Estimating the Amount Owed at Default

For loans and bonds, amount owed at default is simply the par amount. But for lines of credit and counterparty credit on derivatives, the amount owed at default needs to be modeled. We will consider the modeling of the amount owed at default for counterparty credit on derivatives in Chapter 14. Here we will confine our discussion to lines of credit.

Lines of credit enable a borrower to draw funds as needed up to some maximum amount, subject to various terms and conditions. From a completely pessimistic view, $A_D(I)$ would be set for a credit line equal to the maximum amount that can be drawn, since just prior to default a borrower will likely try to maximize the use of all available sources of credit. However, this fails to take into account some of the contractual terms that the lender can employ to limit credit line usage when the credit rating of the borrower is declining. It is thus possible that $A_D(I)$ will be less than the maximum amount that can be drawn.

Two principal forms of credit lines are available—those used for working capital and those used as backstops for commercial paper issuance.

Working capital credit lines give a borrower the flexibility of only paying full interest on the amount of funds it needs at a particular point of time without losing the security of knowing that it can draw down a precommitted amount as needed.

Commercial paper backup lines act as a safety net for commercial paper issuers. Commercial paper issuance typically occurs for very short time periods, often only a few days, to accommodate the liquidity needs of commercial paper investors. The tenor of the commercial paper is usually shorter

than the borrowing need of the commercial paper issuer, leaving the issuer vulnerable to an inability to roll the paper over at maturity, but also leaving the investor vulnerable to not being paid back in the event of rollover failure. The backup line gives assurance to both the borrower and investor in the event of a liquidity squeeze. A backup line is consequently insisted on by rating agencies as a prerequisite for an investment-grade credit rating on a firm's commercial paper. Usage on commercial paper backup lines is virtually zero, except in the rare case of rollover difficulty.

In measuring the loss given default of credit lines, average usage is obviously of little value, since it fails to deal with the high correlation between line usage and credit deterioration. The key is how much usage will there be if default occurs. As noted previously, backup line usage averages close to zero, but when the lines are used, it is because credit difficulties make rolling commercial paper problematic. If only 1 percent of all commercial paper issuers default, but all of these have their lines drawn by 100 percent just prior to default, and if 0 percent usage appears on the remaining 99 percent of issuers, then the overall line usage will be only 1 percent, but default losses will be just as great as if overall line usage is 100 percent.

If credit lines are viewed simply as an option to draw funds exercisable by the borrower, then line usage should be assumed to be 100 percent in the event of default. However, this option is not unconstrained, given that covenants that form part of the contract for the line give lenders the opportunity to reduce line availability in the event of credit deterioration. There will, on one hand, be competitive pressures on the bank not to exercise its full rights under these covenants to avoid damaging the particular relationship and to maintain a reputation with customers as being reliable in a crisis. On the other hand, a bank can pressure a customer to renegotiate loan terms. Araten and Jacobs (2001) aptly describe credit line usage in the event of default as "the outcome of the race between the bank and the borrower with regard to the draw-down of unused commitments in adverse circumstances."

When a result is the product of complex behavioral assumptions, it is not surprising to see that the dominant method of analysis is historical statistical study. Araten and Jacobs (2001) published the most complete analysis based on a study of 399 defaulted borrowers at Chase Manhattan Bank over a 5³/₄-year period, ending in December 2000. Their main results are shown in Table 13.5.

As would be expected, average usage upon default rises with the time elapsed between when a line is committed and when default occurs. This is because the longer the time period elapsed, the more likely that a borrower who started as higher grade and subject to fewer covenants has slipped downward in credit grade. Similar reasoning explains the finding

TABLE 13.5 Average Usage Conditional on Default by Facility Risk Grade and Time to Default for Revolving Credits

Facility Risk Grade	Number of observations in parentheses					
	Time to Default (in Years)					
	1	2	3	4	5-6	Total
AAA/AA		12.1% (1)				12.1% (1)
A	78.7% (3)	75.5% (6)	84.0% (1)			77.2% (10)
BBB+/BBB	93.9% (1)	47.2% (7)	41.7% (5)	100% (2)		55.5% (15)
BBB-/BBB-	54.8% (18)	52.1% (20)	41.5% (9)	37.5% (3)	100% (2)	52.2% (52)
BB	32.0% (81)	44.9% (84)	62.1% (45)	76.0% (17)	68.3% (4)	46.4% (231)
BB-/B+	39.6% (129)	49.8% (100)	62.1% (37)	62.6% (25)	100% (4)	50.1% (295)
B/B-	26.5% (86)	39.7% (22)	37.3% (5)	97.8% (2)		30.7% (115)
CCC	24.5% (100)	26.7% (14)	9.4% (1)			24.6% (115)
Total	32.9% (418)	46.6% (2,54)	62.1% (103)	68.7% (59)	71.8% (59)	43.4% (834)

that average usage upon default tends to rise with a higher initial credit rating. Of course, it is less likely that a higher-rated credit will default compared to a lower-rated credit, but for those who do default, the lower level of covenants results in higher usage.

13.2.4 The Option-Theoretic Approach

Before expounding on the option-theoretic approach, let us review why it would be very useful to have a model that relates a firm's equity price to credit spreads, default probability, and loss given default. First, as noted toward the end of Section 13.2.1.2, the generally greater liquidity and more frequently available quotes of equity prices relative to debt prices makes this an attractive potential driver of inputs to portfolio credit models. Second, the greater availability of historical stock price data makes it attractive as a driver of default correlation models, as we will see in Section 13.3.1. Third, credit spreads derived from equity prices can be a useful input to trading decisions about which credit instruments represent good investment values. Fourth, models of correlations between equity prices and credit spreads can be valuable tools in building models of products, such as convertible bonds, that are hybrids of equity and debt. Fifth, models of correlations between equity prices and credit spreads can be useful input to the creation of stress scenarios. And sixth, certain trading strategies, termed capital structure arbitrage, use option-theoretic analysis to identify mispriced relationships between debt instruments and equity options; Morini (2011, Section 11.2) offers an extensive discussion of these strategies and possible difficulties they may encounter.

In the option-theoretic approach, a firm's equity is viewed as a call option on the value of the firm's assets with a strike price equal to the face value of the firm's debt. This is equivalent to viewing the equity owners of a firm as having a put option to pay off the debt holders with either the face value of the debt or the total value of the firm's assets, whichever is smaller. So the total economic value of the firm's debt to the debt holders must be the face value of the debt less the value of this put option.

Let us first look at a very simple version of the options model, basically corresponding to the original Merton model, which can be found in Hull (2012, Section 23.6). It is extremely useful as a first approximation, since we will see that it provides a precise relationship between all of the elements we are trying to link with very little computational burden. This model has four key simplifying assumptions:

1. The firm has only a single class of debt outstanding, a zero coupon debt, and the firm will not issue any new debt before this debt matures.

2. If the firm defaults, this will only occur at the time of the maturity of this debt.
3. The firm's behavior, such as the riskiness of its investments, will not be impacted by how close it is to default.
4. No intermediate payments, such as dividends, will be made to equity holders.

At the price of these simplifying assumptions, the model requires only four inputs—the time to maturity of the debt, the market value of the firm's assets, the present value of the firm's debts, and the volatility of the firm's assets. The model can give explicit formulas, in terms of these four inputs, for the probability the firm will default, the loss given default, the required interest rate spread over the risk-free rate for the firm's debt, and the market value of the firm's equity and debt.

Using notation close to that in Hull, we'll denote:

V_0 : The current market value of the firm's assets

D_0 : The present value of the firm's debt, which matures at time T , discounted at the risk-free interest rate

σ_V : The volatility of the firm's assets

P_D : The probability of default

L_D : The loss in the event of default

Viewing the equity as a call option on the firm's value with a strike price of the face amount of the debt, we can write a formula for the current market value of the firm's equity as:

$$E_0 = V_0 N(d_1) - D_0 N(d_2) \quad (13.2)$$

where

$$d_1 = [\ln(V_0/D_0) + \sigma_V^2 T/2] / \sigma_V \sqrt{T}$$

$$d_2 = d_1 - \sigma_V \sqrt{T}$$

The current market value of the firm's debt is just $V_0 - E_0$.

Following the standard Black-Scholes analysis, $N(d_1)$ is the delta, the partial derivative of E_0 with respect to V_0 , and $N(d_2)$ is the probability that the strike price will be exceeded at time T . But this is the probability that the firm will not default so:

$$P_D = 1 - N(d_2) \quad (13.3)$$

If no default occurs, the debt holders receive the face value of the debt and, if default does occur, the debt holders receive the recovery rate times the face value of the debt, so we can write the market value of the debt as:

$$V_0 - E_0 = [(1 - P_D) + P_D(1 - L_D)]D_0 \quad (13.4)$$

Substituting from Equations 13.2 and 13.3,

$$V_0[1 - N(d_1)] + D_0N(d_2) = \{(1 - P_D) + [1 - N(d_2)](1 - L_D)\} D_0 \quad (13.5)$$

Solving this equation for L_D , we get:

$$L_D = 1 - (V_0/D_0)[1 - N(d_1)]/[1 - N(d_2)] \quad (13.6)$$

If the debt were truly risk free, its market value would be D_0 . The credit spread on a zero coupon instrument can be written as s , where the market value (MV) of the instrument is the face amount (F), discounted by $r + s$, where r is the risk-free rate.

Thus,

$$\begin{aligned} MV &= Fe^{-T(r+s)} \\ e^{-Ts} &= MV/Fe^{-Tr} \\ s &= -\ln(MV/Fe^{-Tr})/T \end{aligned} \quad (13.7)$$

We know the market value of the debt is $V_0 - E_0$ and the present value of the debt discounted by the risk-free rate, Fe^{-Tr} , is D_0 .

Thus,

$$s = -\ln [(V_0 - E_0)/D_0]/T = \ln[D_0/(V_0 - E_0)]/T \quad (13.8)$$

Two of the four required inputs, T and D_0 , are easy to determine, provided all the firm's debts are reported in some publicly filed statement. To use the model as an approximation when several maturity dates are available for debt and the debt has scheduled coupon payments, T can be calculated as the weighted average duration of the debt.

In theory, you could obtain V_0 by summing the market prices of all the firm's equity and debt and estimate σ_V by looking at the historical volatility of this sum. In practice, most firms have some amount of debt that is not publicly traded and for which a market price would therefore not be available.

Inputs that can be obtained easily are the market price of equity, E_0 , and the volatility of equity price, σ_E , which can be based on both historical

observation and implied volatility from equity options. To obtain V_0 and σ_V from E_0 and σ_E , solve the simultaneous equations:

$$E_0 = V_0 N(d_1) - D_0 N(d_2) \quad (13.9)$$

and

$$\sigma_E E_0 = N(d_1) \sigma_V V_0 \quad (13.10)$$

The latter equation can be derived from Ito's lemma and the fact that $N(d_1)$ is the partial derivative of E_0 with respect to V_0 . The **MertonModel** spreadsheet takes E_0 , σ_E , D_0 , and T as input and solves for V_0 , σ_V , P_D , L_D , MV , and s .

Whenever I have tested this model out on real data, the result has always been the same—reasonable values for P_D but unreasonably low values for L_D and for s —values produced for L_D would be around 10 percent when real experience with loss given default is usually 50 percent or greater, as can be seen in Table 13.5.

To explore which of the simplified assumptions of the model considered thus far is leading to this divergence from reality, we could move to a Monte Carlo model that reproduces many possible future paths of the firm's asset value. The growth rate of the asset value assumed would be the risk-free rate by the usual risk-neutral valuation argument. It is easy in the context of a Monte Carlo model to build in payments due to different maturities of debt with coupons, build in rules for when default will occur (such as when the net worth of the firm is below a certain threshold), and build in rules for the distribution of asset value in the event of default to different seniority levels of debt. It is also easy to build in behavioral rules for the firm's response to different levels of net worth (such as increasing asset volatility as the net worth gets close to the default threshold or issuing new debt as it gets further from the default threshold) and build in rules for dividend policy. By summing over all paths in the Monte Carlo model, it is easy to compute the expected default rates by time period, recovery rates in the event of default by time period and seniority level, and the market value of equity and of each combination of maturity and seniority level of debt. Required spreads over the risk-free rate for each combination of maturity and seniority level of debt can be computed from the market value. When the assumptions of the simple options model are input to the Monte Carlo model, the same result is obtained as from the simple model.

When this model is implemented, we can see what is driving the unrealistic L_D and s outputs. If the default threshold is set greater than zero and if asset values are assumed to follow paths without jump processes, then the

required spread over the risk-free rate can be driven as close to zero as desired by increasing the frequency with which observations of the asset value are taken. Increasing the frequency of observation increases the probability of default, but it also causes the loss in the event of default to approach zero by dividing up the assets of the firm among the creditors while they are still sufficient to pay off the creditors in full. This shows that the key issues in determining default loss are behavioral rather than financial; that is, they depend critically on how transparent the operations of the firm are to creditors and how much control the creditors can exercise in forcing bankruptcy in a timely fashion. This may differ significantly by government jurisdiction. The role governments may play in providing help for firms close to default may also differ.

There are two ways forward from this impasse. One is to focus on models that do incorporate jump processes. The other is to stick with a simple model but treat it just as a heuristic that can be input to a statistical analysis. We will explore both in turn.

13.2.4.1 Jump Process Models Many such models have been proposed. A good summary with references and discussion for a variety of such models and is Bohn and Stein (2009, Chapter 3).

It is important to distinguish between two reasons why a jump process may exist. One is that the asset value of the firm may follow a jump process. The other is that there can be discontinuities in the asset value threshold that will lead to default. CreditGrades (2002) in documenting a model of the second type states, in the introduction to Chapter 2, “In our approach, we model the uncertainty in the default barrier, motivated by the fact that we cannot expect to know the exact leverage of the firm except at the time the firm actually defaults. The uncertainty in the barrier admits the possibility that the firm’s asset value may be closer to the default point than we might otherwise believe.” The advantage of this approach is that it is consistent with the term structure of credit spreads observed in the market; without uncertainty around how close a firm currently is to the default point, one would expect to see much lower short-term credit spreads than are actually observed.

The **JumpProcessCredit** spreadsheet implements a jump process model closely related to the one documented in CreditGrades (2002) and also documented in Schonbucher (2003, Section 9.5). This model has advantages similar to the Merton model, in requiring very few inputs and being relatively easy to understand. The model assumes the same sort of stochastic evolution of total firm asset value as the Merton model, but assumes that default could occur at any time the asset value falls below a default barrier. Inputs are E_0 and σ_E , as in the Merton model, along with the risk-free rate

and both the mean and standard deviation of the default barrier. This input for the default barrier takes the place of the present value of debt that is input to the Merton model. Unlike the Merton model, this model does not attempt to compute a loss given default rate; this is assumed to be estimated by statistical means for each class of debt, as per Section 13.2.2. The model outputs probability of default and credit spread for any desired time period; unlike the Merton model, it is not restricted to a single time period corresponding to the tenor of existing debt.

As our quote from CreditGrades (2002) in the paragraph before last indicates, standard deviation of the default barrier is assumed to represent uncertainty about the current level of the default barrier, and hence is independent of time period. Section 2.2 and Figure 2.2 of CreditGrades (2002) show that a standard deviation of 30 percent for the default barrier, derived from historical statistics on actual recovery data, produces a term structure of credit spreads that is consistent with market observations.

The CreditGrades model does not have an input for the mean of the default barrier. Instead it is assumed to be the face value of outstanding debt multiplied by the historical average loss given default, averaged over all of the outstanding debt of the firm. While CreditGrades (2002, Section 2.2) presents an argument for the plausibility of this assumption, there is a range of values for the default barrier that would also be plausible given the criteria of the CreditGrades argument. I have chosen in the **JumpProcessCredit** spreadsheet to leave the mean default barrier as a user input. Users can choose the CreditGrades assumption or experiment with default barrier levels that seem to fit the historical credit spreads of a particular issuer or category of issuers (say a grouping by industry and country). Exercise 13.2 is designed to give you an understanding of the differences between the Merton model and the jump process model in results and in sensitivities to inputs.

13.2.4.2 Statistical Analysis Even simple options models can still play a useful heuristic role in helping to understand the default process. This is the role they play in the models of Moody's KMV, whose analysis is widely utilized among investors in credit instruments. Crosbie and Bohn (2003) summarize the KMV methodology. De Servigny and Renault (2004, Chapter 3) in the section "KMV Credit Monitor Model and Related Approaches" provide a brief review of the model, along with some reservations.

The KMV approach is to utilize a model somewhat like the simple Merton model we first discussed, but the objective is to use it not to try to directly measure default probability, but rather to produce a measure called *distance to default*, which is then used to project default probabilities based on an empirically fitted statistical model. Technically, the model utilized by

KMV treats equity as “a perpetual option with the default point acting as an absorbing barrier for the firm’s asset value” (see Crosbie and Bohn 2003, Section 3). The insight behind this is that, whereas the behavioral nature of default requires the use of statistical observation of past experience, the options model output can be a valuable input to this process when used comparatively to judge which firms are relatively more likely to default than others. In this approach, statistical models, not option-theoretic ones, are employed in estimating loss in the event of default.

KMV presents the following points in favor of this use of the option model:

- Because the model is based on equity market prices, which are continuously observable, it is more likely to represent the latest available information than the ratings of just a single firm’s credit officers or a rating agency or on statistical models based on accounting information that is only available periodically. It can also be applied to any public company, even one that does not have publicly rated debt, since it is based on equity prices.
- The model takes into account both the capital structure of a firm and its business and industry risk. Capital structure is represented by the *leverage*, the ratio of total firm value to equity. Business and industry risk is represented by the volatility of asset values. (For example, you can expect much more volatility from a firm in a high-tech industry than a utility, or much more volatility from a firm in an emerging market country than one in an established industrial country.)

The distance to default is measured by the number of standard deviation movements it would take to put a firm at the point where default is a serious possibility. In terms of the simple model we presented, it would be $(V_0 - D_0)/(V_0 \sigma_V)$, which is calculated in the **MertonModel** spreadsheet. The actual model used by KMV to calculate the distance to default is more complex than our simple model in several ways. To highlight a few:

- Our simple model assumes that default can occur only when firm asset value is insufficient to make a required payment. The KMV model recognizes that firms can be forced to default when their asset values decline sufficiently below the present value of required future payments. Based on empirical studies, KMV has set the default point, which in our model is D_0 , as the sum of short-term debt, representing required current payments, and one-half of long-term debt, representing payments that will be required in the future. In this way, assets can decline below the required future payments by some amount, but not too far,

before default is threatened. De Servigny and Renault (2004, Chapter 3) note that this is a purely empirical rule of thumb that “does not rest on any solid theoretical foundation. Therefore there is no guarantee that the same rule should apply to all countries and jurisdictions and all industries. In addition, little empirical evidence has been shown to provide information about the confidence level associated with this default point.” This critique should be compared to the response to the question “Are default probabilities applicable across countries and industries?” in Crosbie and Bohn (2003, Section 6).

- The KMV model can handle more liability classes than just straight debt and equity; it can also accommodate hybrid classes—convertible debt and preferred stock.
- KMV regards Equation 13.10 as too simplistic, since it does not take into account the impact of varying leverage levels through time on the relationship between equity volatility and asset volatility. KMV uses a more complex model to reflect this factor. In particular, the concern is that for a firm whose performance is trending downward, the decline in equity value will result in current leverage being higher than its leverage has been in the past. If asset volatility is estimated from its historical equity volatility and its current leverage, this will tend to understate historical asset volatility, resulting in understating the default probability. The converse of this effect will result in overstating the default probability for a firm whose performance is trending upward. As Crosbie and Bohn (2003, Section 4) state, this “biases the probabilities in precisely the wrong direction.”

KMV’s solution is a more granular approach in which a time series of historical daily asset returns is constructed from historical daily equity returns and Equation 13.2, based on an initial guess at σ_v . These daily asset returns can then be used to compute a new guess at σ_v , leading to a new series of daily asset returns. The process is repeated until it converges (see Crosbie and Bohn 2003, Section 4).

Many aspects of KMV’s methodology are proprietary and undisclosed, but the results they have published have had a major impact on firms that manage credit risk, both as a source of information and as an inspiration for their own research. De Servigny and Renault (2004, Chapter 3) note that “Many banks have developed their own systems to extract early warning information from market variables. Many variants can be found that extract the volatility of the firm from either equity time series, implied volatilities in options markets, or even spreads. . . . Equity-based models reflect the market’s view about the probability of default of specific issuers and therefore can provide valuable early warning signals. Unfortunately they are no panacea,

as they also reflect all the noise and bubbles that affect equity markets. Overall, they can be seen as a useful complement to an analysis of a firm's fundamentals." Bohn and Stein (2009, Chapter 3) in the section "Modifying BSM" provide references to empirical research that "cast[s] doubt on the practical viability of structural models" but observe that "numerous financial institutions around the world have successfully implemented and tested credit risk management systems based on the structural framework." (In this context, "structural" is equivalent to what we have been calling "option-theoretic"—models that are not just based on statistical linkages but utilize options theory to link default probability to equity prices.)

Altman, Fargher, and Kalotay (2010) present results supporting the use of statistical models of default probabilities that combine equity market information with traditional accounting variables (of the type discussed in Section 13.2.1.2). They provide references to other published models utilizing equity market based inputs with a discussion of comparative results. Bohn and Stein (2009, Chapter 3) in the section "Modifying BSM" also observe that a "promising point of departure is that of the hybrid approach, where characteristics of both structural and reduced-form models . . . or structural and econometric approaches are combined," and provide many references to published hybrid models.

13.3 PORTFOLIO CREDIT RISK

In Section 13.2, we have established the main building blocks that are needed for analyzing portfolio credit risk. The remaining building block is estimation of correlations between defaults, which we will investigate in Section 13.3.1. We will then turn, in Section 13.3.2, to Monte Carlo simulation models that bring all of these building blocks together, and look at computational alternatives to full simulation in Section 13.3.3. Finally, in Section 13.3.4, we will examine how simulation models and the tools of Sections 13.1 and 13.2 can be used in the management and reporting of portfolio credit risk.

13.3.1 Estimating Default Correlations

Let's begin with some points on which almost everyone who has worked on this topic can agree:

- Strong evidence supports a positive correlation between defaults—that is, that defaults tend to occur in clusters. For example, Table 13.6, from Moody's (2011a), of default percentages by year for ratings categories

TABLE 13.6 Default Percentages by Year

Year	Rating		
	Baa	Ba	B
1982	0.32%	2.78%	2.35%
1983	0.00%	0.91%	6.36%
1984	0.36%	0.83%	6.75%
1985	0.00%	1.41%	7.48%
1986	1.01%	2.05%	11.60%
1987	0.00%	2.73%	6.49%
1988	0.00%	1.26%	6.20%
1989	0.60%	3.04%	8.72%
1990	0.00%	3.41%	15.47%
1991	0.28%	4.89%	12.36%
1992	0.00%	0.31%	9.22%
1993	0.00%	0.57%	4.56%
1994	0.00%	0.25%	4.06%
1995	0.00%	0.73%	4.26%
1996	0.00%	0.00%	1.37%
1997	0.00%	0.19%	1.94%
1998	0.12%	1.00%	4.15%
1999	0.11%	1.32%	3.81%
2000	0.39%	0.72%	4.90%
2001	0.20%	1.39%	6.03%
2002	1.10%	1.38%	9.57%
2003	0.00%	1.00%	4.53%
2004	0.00%	0.41%	2.11%
2005	0.18%	0.00%	0.84%
2006	0.00%	0.20%	1.03%
2007	0.00%	0.00%	1.18%
2008	0.47%	1.16%	2.07%
2009	0.86%	2.40%	7.41%
2010	0.00%	0.00%	0.48%

Source: Based on Moody's (2011a, Exhibit 31).

Baa, Ba, and B, shows much higher default rates in recession periods, such as 1989 to 1991 and in 2001 to 2003, than in periods of economic growth, such as 1993 to 1997.

- Estimating this correlation based on the joint default history of firms with the same credit rating is unsatisfactory. Grouping together all firms with the same credit rating ignores factors such as whether firms are in the same industry or whether firms are located in the same geographical region, but these factors are widely believed to influence joint default correlation. See Gupton, Finger, and Bhatia (1997, Section 8.2).
- The direct estimation of joint default correlation by examining historical defaults categorized by rating, country, and industry is not a feasible approach. Default is a relatively rare event and with this fine a segmentation, there would not be enough observation to allow robust statistical inference. A way around this impasse is to estimate correlation for a variable that can be more frequently observed and can then be utilized to produce default correlations.

For KMV, asset returns are a very natural choice for such a variable, since they are directly tied to defaults through the distance-to-default measure and its statistical relationship to default probability. KMV utilizes the methodology we discussed in Section 13.2.3 to delever equity returns directly observed in the market and compute asset returns. It is an easy step from creating a time series of asset returns for a large universe of borrowers to computing correlations between asset returns for those borrowers. Monte Carlo simulations of correlated movements in asset returns, which we discuss in the next section, can then be used to calculate the percentage of cases that result in joint default, enabling a default correlation to be computed. The actual methodology employed by KMV does not directly calculate asset return correlations between pairs of borrowers. Instead, a factor analysis is used in which composite asset returns are calculated for sectors—countries and industries as well as groupings of countries and industries. Historical asset return correlations can then be computed between sectors. Asset return correlations between borrowers can then be easily computed based on the statistical relationship between each borrower's asset return and those of the country and industry sectors. The KMV approach to correlations is described in more detail in the section on “Model of Value Correlation” in Kealhofer and Bohn (2001) and in the section “Calculating Correlations Using Moody's KMV Portfolio Manager” in Chapter 8 of Saunders and Allen (2010).

The CreditMetrics approach to estimating default correlations is very similar to KMV's, except that correlation between equity returns is used as a proxy for correlation between asset returns. Gupton, Finger, and Bhatia (1997, Section 8.5) provide great detail on this process.

The default probability implied by credit spreads is another natural candidate to be used. The drawbacks are that this involves a much smaller universe of borrowers for whom liquid public debt prices are available relative to the number of borrowers for whom liquid equity prices are available and that implied default probabilities from public debt prices significantly overstate actual default probabilities, as we discussed in Section 13.2.1.2. As a result, firms that do decide to use market implied default probabilities as indicators of relative credit quality, but may choose to adjust the overall default probability by a factor that lowers these probabilities to anticipated rates of actual default, following our discussion in Section 13.2.1.2.

Whichever variable is used to provide the linkage, the key to transforming shorter-term correlations into longer-term default correlations is a simulation of movements through time, which we will come to in the next section. Because of the relative infrequency of default, even high short-term correlations transform into much smaller default correlations.

While much of the early work on building default correlation relationships focused on estimating correlation coefficients, the recent trend has been to also place a lot of attention on the shape of the correlation relationship, the copula. For example, it is frequently the case that large moves in changes in default probability are more closely correlated than smaller moves. The section on “Copulas” in Bohn and Stein (2009, Chapter 8) and Duffie and Singleton (2003, Section 10.4) give an introduction to this topic in the context of estimates from historical data. The assumption that correlation is the same for all sizes of changes in probability is known as the *Gaussian copula* assumption. Estimates based on market prices of credit correlation products, such as CDOs, will be discussed in Section 13.4.2.

One recent trend has been to utilize *frailty analysis*, a technique borrowed from medical research, to correct for underestimation of correlation due to undetected factors that can have a common impact on many borrowers. A good explanation can be found in Duffie et al. (2009), which provides a frailty model applied to corporate defaults, a statistical analysis of the degree of correlation underestimation that may occur if this correction is not accounted for, and references to related literature. Another recent trend has been to introduce modeling of *contagion*, the impact that defaults by one or more firms may have in increasing the default probabilities of remaining firms. Rullière, Dorobantu, and Cousin (2010) provide a recent model of this effect with references to related literature.

13.3.2 Monte Carlo Simulation of Portfolio Credit Risk

Where we stand, based on the Sections 13.2.1, 13.2.2, 13.2.4, and 13.3.1, is that a variety of methods have been presented for estimating default

probability and loss given default for individual borrowers and for estimating the short-term correlation between borrowers for some variable that can be linked to longer-term defaults. We now focus on analyzing methods that can provide this linkage and also produce calculations of portfolio risk very similar to the portfolio risk measures that were provided for market risk in Chapter 11.

Let's begin by assuming that all of this analysis will be provided by Monte Carlo simulation. For many of the same reasons stated in our analysis of VaR in Section 11.1, simulation is the most accurate method of generating portfolio risk measures. It has the flexibility to incorporate almost any assumption about statistical distributions we want to make. Later, in Section 13.3.3, we will discuss possible shortcuts to the calculation of portfolio risk under more restrictive distributional assumptions. However, simulation is always the benchmark against which the accuracy of other approximations can be tested. The reason why we only consider Monte Carlo simulation for credit risk, while we considered the alternatives of Monte Carlo simulation and historical simulation for market risk VaR, is that the longer time periods involved in credit risk simulations mean that not enough nonoverlapping historical data points will be available to derive a historical simulation.

A Monte Carlo simulation will follow a key variable, whether it is asset value, macroeconomic factors, or default probability, for each borrower to whom credit has been extended. The simulation will be based on assumptions about the volatility of asset returns or transaction matrices for default probabilities and assumptions about short-term correlations between the borrowers (both correlation coefficients and copulas, as discussed in Section 13.3.1). If asset value is being used as the key variable, it must be converted into default probabilities, using a statistical relationship such as the one developed by KMV between an asset's distance to default and probability of default. Defaults can then occur at random, based on the probability of default. In the event of default, a random sample is drawn based on the mean and standard deviation of the loss given default for a given seniority class of instrument (if instruments of different seniority are outstanding to the same borrower, a correlation should be enforced between the degree to which the loss given default on each instrument exceeds or is below the average). Correlations between default probability and LGD, as discussed in Section 13.2.2, can be specified as part of the simulation. Detailed descriptions of such simulation models can be found in Duffie and Singleton (2003, Chapter 10), Schonbucher (2003, Chapter 10), and Bohn and Stein (2009, Chapter 8).

A Monte Carlo model meeting this description has many possible applications. It can be used with just two borrowers to translate a short-term correlation of assets values or credit spreads into long-term default correlations.

It can also be used with an entire portfolio of assets to generate statistics on expected credit losses and the full distribution of credit losses, such as losses at the 99th percentile. It can be used for valuing a tranching CDO by tracking losses to each tranche along each of the paths and then calculating the expected losses on each tranche, as we will discuss in Section 13.4.2.

Thus far we have been discussing a Monte Carlo simulation that only deals with a single time period in which the relevant outcomes are for each credit to either default or not default. This can be extended in one of two directions. One direction would be the simulation of an end-of-period change in credit grade and credit spread in addition to default. This extension requires a tie between the key variable being simulated and credit grade and credit spread. This relationship is straightforward for implied default probabilities and is provided for asset values by KMV's statistical linkages of the distance to default to credit rating. KMV has also developed a linkage between asset values and credit spreads (see Agrawal, Arora, and Bohn 2004), partially based on the capital asset pricing model and partly on statistical relationships. The other direction would be multiperiod simulation. Multiperiod simulation could be achieved by just computing default loss distributions at different points along the simulation path. However, full accuracy requires some simulation of possible changes in the overall economic climate, factoring in features such as the increased probability of an economic downturn following a period of sustained economic growth (see Wilson 1997).

De Servigny and Renault (2004, Chapter 6) give a thorough discussion of four commercially available Monte Carlo simulation models for credit portfolios: the CreditMetrics model documented in Gupton, Finger, and Bhatia (1997); the Moody's KMV Portfolio Manager model documented in Kealhofer and Bohn (2001); Standard & Poor's Portfolio Risk Tracker mode documented in de Servigny et al. (2003); and Financial Analytics's CreditPortfolio View documented in Wilson (1997). De Servigny and Renault's summary in their Tables 6.1 and 6.3 is particularly useful in showing at a glance the similarities and differences in these four models.

De Servigny and Renault's analysis shows greater similarities than differences. All four models simulate stochastic evolution of default probabilities and use stochastic LGD rates, but only Portfolio Risk Tracker includes correlation between default probabilities and LGD rates. All four models handle correlations between defaults by simulation of common factors that drive default probabilities. They differ on the common factors that drive default probabilities—country and industry factors for CreditMetrics, Portfolio Manager, and Portfolio Risk Tracker, and macroeconomic factors for CreditPortfolio View. Both CreditMetrics and Portfolio Manager derive the relationships that drive these common factors from equity markets (as

discussed in more detail in Section 13.3.1), while Portfolio Risk Tracker offers user flexibility to base correlations on equity, credit spread, or empirical data. Their biggest differences involve outputs—Portfolio Manager focuses on the distribution of default losses, CreditMetrics and Portfolio Risk Tracker on the distribution of changes in market value that can be driven by both defaults and ratings changes, while Credit Portfolio View gives a choice between these two distributions.

While information on changes in market value can be very useful information for credit portfolio managers, as we will discuss further in Section 13.3.4, I am extremely suspicious of any approach to portfolio credit risk that does not include a focus on distribution of ultimate default losses. Since portfolio credit risk is a long-term risk not amenable to management with liquid market instruments, as discussed in the introduction to this chapter, the approach of Sections 6.1.2 and 8.4 should govern. Any modeling that relies on future market value is assuming a future liquidity that cannot be relied on. If no statistics on ultimate default losses are available from the model being used, then there needs to be a ready way of translating model output on market value changes into distributions of ultimate default losses.

A Monte Carlo simulation of individual loans becomes computationally infeasible for loan portfolios with very large numbers of very small loans. This is certainly true for retail loans such as home mortgages or credit cards. In such cases, the portfolio needs to be analyzed into segments that can be treated as roughly homogeneous. For example, a portfolio of home mortgages could be divided into segments grouped by geography and home value. Each segment now must be treated as a single loan in the Monte Carlo simulation of the entire firm's loan portfolio. But unlike a true individual loan, which is in either one of two states, default or nondefault, a grouping of small loans must be represented by a percentage of loans that default in a particular time period. An analysis of the history of default patterns can establish statistics to drive the simulation, including correlations with default levels between two segments and between the segment and individual loans being simulated. The best way to derive historical correlations may be through a mutual dependence on macroeconomic factors, such as growth rates in the economy; see Wilson (1997) and Bucay and Rosen (2001). For a general overview of a simulation of defaults on retail credits, see Risk Management Association (2000).

One way in which the use of simulation for credit risk differs from its use for market risk VaR is that the expected value of the distribution plays a significant role. Since market risk VaR is computed over very short time periods, expected value can either be ignored or else has only a minor impact. The far longer time horizon of credit risk simulation requires that expected credit loss be accounted for. Expected credit loss should be taken

as a charge against earnings, either in the form of a reduction in valuation for a portfolio that is marked to market or in the form of a loan loss reserve for a portfolio that uses accrual accounting. Risk should be measured and capital should be allocated based on the unexpected losses—the distribution of returns around the expected losses.

Parallel to the discussion in Section 7.2 of stress testing as a complement to VaR for market risk, it is often desirable to complement the statistical analysis of credit risk for a portfolio with a stress test based on economic insight, for example looking at the impact of an unusually prolonged global recession. This is especially true in evaluating the risk of credit concentration to firms doing business within a particular country. Credit concentration within a country leads to the risk of correlated outcomes since all firms may be impacted by how well the country's economy performs. This type of correlation risk is very much the same type of risk as the risk of credit concentration within an industry or within a geographical region of a country. All of these correlation risks can be reasonably measured by statistical means. But country risk has an additional dimension. The possibility exists that all firms, individuals, and government bodies within a given country will be prohibited from meeting their contractual obligations. This can arise from the imposition of exchange controls by the government as a defensive measure against adverse currency flows, or from government renunciation of foreign debts, or from disruption of normal contractual relationships due to war or revolution. This form of risk represents a major political discontinuity that statistical analysis of historical economic data will shed little light on. It can best be quantified by looking at the extent of damage in past incidents of political disruption in other countries combined with subjective assessment of the likelihood of occurrence based on economic and political insights into the current conditions within a particular country. See Bouchet, Clark, and Gros Lambert (2003).

Using Monte Carlo methods to design credit portfolio stress tests, parallel to the discussion in Section 7.2.3, is specifically addressed in Breuer and Csiszar (2010, Sections 3.3 and 3.4).

13.3.3 Computational Alternatives to Full Simulation

Portfolio credit risk analysis does not have the same need for rapid turnaround that models used for trading liquid instruments do. Changes in the portfolio occur more slowly, you don't need to respond to the needs of a trading desk requiring an up-to-date picture at the start of each trading day, or to contribute to the daily VaR market risk calculations. So there is much greater tolerance for full simulation runs that may require many hours or even a few days to produce statistics. Even so, there will be a desire to see

the impact of alternative strategies in building the credit portfolio or in buying protection that may need a speedier approximation to accommodate multiple runs. And definitely the need to produce marginal risk analysis for the impact of a proposed new loan, discussed in Section 13.2.4, will require a fast approximation technique.

Fortunately, approximations can easily be tested for accuracy against the full simulation, following the prescriptions of Section 8.2.5. And fortunately a great deal of clever mathematics has been developed to produce good approximations in reasonable time. Much of this work has been in support of credit derivatives, such as CDOs, which are traded in a market environment and have an even greater demand for quick estimation, as we will see in Section 13.4.2. But, since full simulation models for portfolio credit risk and for CDOs are virtually identical, portfolio credit risk can benefit from these quantitative advances.

The two most important ideas that have been introduced for approximating full simulations are the *large homogeneous portfolio* (LHP) approximation and the Vasicek model that utilizes only a single factor to drive correlation relationships.

The LHP approximation looks very similar to the approach to simulation modeling of large numbers of very small loans discussed in Section 13.3.2. Loans are grouped by common characteristics (this could include industry, geography, credit rating, estimated probability of default) and each group is simulated as if it were a single loan, but instead of being represented by just two states (default or nondefault), the representing state is the percentage of loans within the group that have defaulted as of a given time step. Within each group, the loans are treated as homogeneous (i.e., all having the same default probabilities, LGD probabilities, and default correlations with loans in other classes). The number of loans within each class is treated as large enough that the class can just be represented by an overall percentage of default without worrying about the actual sizes of individual loans within a category. The less equal the loan sizes are, the less accurate this assumption will be; for example, if there were one very large loan, its default would cause a jump in the percentage of defaults for the category. The fewer categories you use, the faster the simulation will run, but the less accurate the approximation to the full simulation model will be.

The Vasicek model utilizes only a single factor, roughly corresponding to the state of the world economy, to drive the simulation. This is obviously a major approximation, since much of the detail of correlation relationships based on industry-specific and geography-specific factors will now be lost. But the resulting simplification allows calculations to be performed by much quicker numerical integration methods, as opposed to simulation; see Schonbucher (2003, Section 10.4.3) for details. Even in cases where the

decision is to rely on a fuller simulation for risk reporting, this much faster calculation allows quick estimation of sensitivities to input variables that can be valuable for building intuition for portfolio managers. The Vasicek model is a particularly useful approximation for building intuition because of its strong emphasis on the separation of systematic risk and idiosyncratic risk, as we will see in the following pages.

Even quicker numerical integration can be achieved by combining the Vasicek model with the LHP approximation. The most frequently encountered version of this combination also assumes a Gaussian copula (see Section 13.3.1) for the correlations. This version, the now infamous Li model (see Section 5.2.5.3), has been well documented, for example Hull (2012, Section 23.9) or Schonbucher (2003, Section 10.4.4), as well as the original Li (2000).

The Vasicek model operates by keeping track of all default correlations using a single common factor and calculating losses corresponding to each possible value of this common factor. Tying computations to this common factor, which could be thought of as the state of the economy or, for mortgage portfolios, the level of housing prices, is what makes the Vasicek model so appealing for gaining an intuitive grasp of the impact of systematic risk (we will discuss this further in Section 13.4.4). What makes computation so easy is that all idiosyncratic risk is incorporated through a simple formula applied to each level of the common factor. We will quickly look at how this is done.

Since the portfolio is assumed homogeneous, there are only three input variables to describe the underlying credit portfolio of the CDO: the expected default percentage of this portfolio, D ; the recovery rate in the event of default, R ; and the correlation between changes in asset values, ρ . All the other inputs reflect the structure of the CDO—the attachment and detachment point of each tranche.

The model assumes that each individual credit has associated with it a standard normally distributed variable x_i that reflects the *distance to default* of the credit. There is a threshold value such that, if x_i goes below the threshold, default will take place. Since we know that the probability of default of each credit is D , this threshold must be $N^{-1}(D)$, where N is the cumulative standard normal distribution and N^{-1} is its inverse. The variable x_i may be written as:

$$x_i = \sqrt{\rho}M + \sqrt{1-\rho}Z_i$$

where M is a common factor affecting all defaults and Z_i is a factor affecting only credit i . The variable M and all the Z_i variables are assumed to have independent standard normal distributions, so that the relationship assures

that all pairs of credits have a correlation of ρ and that all x_i variables have a standard normal distribution.

The probability of default of any individual credit is $x_i < N^{-1}(D)$, which becomes

$$\sqrt{\rho}M + \sqrt{1-\rho}Z_i < N^{-1}(D)$$

or

$$Z_i < (N^{-1}(D) - \sqrt{\rho}M) / \sqrt{1-\rho}$$

so that the probability of default is:

$$N\left[(N^{-1}(D) - \sqrt{\rho}M) / \sqrt{1-\rho}\right]$$

The next step is to numerically integrate over many different possible values of M . For each one we can calculate the percentage of total defaults, which multiplied by $(1 - R)$ gives the percentage of total losses. We can easily calculate the losses due to each tranche, utilizing the attachment and detachment points, corresponding to each value of M . We make use of the LHP assumption to treat these loss estimates as exact (given the value of M), rather than just the central point for a probability distribution. We then use the probability distribution of the values of M to infer probability distribution of the tranche losses.

The assumption of a Gaussian copula is not at all necessary for the quick numerical integration technique to work; see O’Kane (2008, Sections 21.5 and 21.6) and Schonbucher 2003, Section 10.8.2) for examples. Similar, but more computationally intense, integrations can be used for multifactor approximations (see Schonbucher (2003, Sections 10.4.5 and 10.4.6). The CDO spreadsheet on the website for this book will allow you to experiment with a Vasicek model that uses the LHP approximation but with several choices for the copula. O’Kane (2008, Sections 16.4 and 18.6) analyzes the accuracy of these approximations.

The LH+ model is an interesting compromise (see O’Kane 2008, Section 17.3). It’s a Vasicek model that uses the LHP approximation for the entire portfolio except for a single loan that is individually modeled. It can still make its computations using a fast numerical integration, but as O’Kane says, it “allow[s] us to understand the interplay between the characteristics of the single credit and those of the overall portfolio.” This can obviously be very valuable when making decisions about how to place an internal price on a new loan or whether to buy credit protection against an existing one.

Alternatives to the LHP assumption can achieve speeded approximation while retaining more detail about the structure of individual loans within the portfolio. O’Kane (2008, Chapter 18) provides a good overview of these approximation techniques. A good starting point for learning about these models would be Hull and White (2004), which is particularly clear in its exposition, presents two models that are relatively easy to implement, and provides references and comparisons to other similar approaches in the literature. Both models work with individual loan data and utilize recurrence relationships in place of Monte Carlo simulation for calculation. The first calculates probabilities of exact loss percentages from the recurrence relationships using Gaussian quadrature. The second approximates the chances of losses falling into user-specified probability buckets. Two other well-known models along comparable lines are Andersen, Sidenius, and Basu (2003), which also provides many references to similar approaches, and Laurent and Gregory (2003), which utilizes a fast Fourier transform. Bluhm, Overbeck, and Wagner (2002, Chapter 4) is a good exposition of the commercially available CreditRisk+ model that utilizes recurrence relationships and probability-generating functions in place of Monte Carlo simulation. Since credit risk calculations are focused on occurrence of default, which is a low-probability event, improvements in accuracy relative to computation time can be gained utilizing importance sampling, a technique that focuses more of the simulation paths on those probability regions where default is more likely to occur. Glasserman and Li (2005) is a key paper in this area. Glasserman (2004, Section 9.4.3) covers similar material. Giesecke and Shkolnik (2011) is a recent contribution and provides many references to similar approaches.

13.3.4 Risk Management and Reporting for Portfolio Credit Exposures

A traditional bank managing a large portfolio of credit risk will need to find the proper balance between the illiquidity of much of its portfolio and the liquid instruments that can allow it to manage some of its risk. On the one hand, the illiquidity of some borrowers and the size of exposures to even liquid borrowers will preclude any chance of using liquid instruments to eliminate all of the exposure. On the other hand, a combination of loan sales, purchases of credit insurance through CDSs, and packaging of some credit in CDOs does permit some choices on the composition of the portfolio.

Choices about sales of existing loans or purchases of credit insurance against current positions are not the only tools available to a credit portfolio manager. A key tool is the internal pricing for taking on new credit exposure. When credit managers have a more favorable view of the credit prospects of a particular borrower than the market does, they will convey

this to the firm's relationship managers by quoting internal prices that reflect narrower credit spreads than those quoted in the CDS market. (The same will apply for particular classes of loans where the credit managers' view of the combination of default probability and LGD is more favorable than is reflected in the market.) As we noted in Section 13.2.1, it is important that the views of the firm's credit managers be challenged when they are more favorable than ratings agencies or market prices imply, but when the credit managers' judgment is sustained after such a challenge, it is an appropriate strategy for the firm to encourage such lending—it may very well be that long experience with particular borrowers or industries or particular expertise gives the firm an edge that it should be taking advantage of.

Conversely, when the firm's credit managers have a less favorable view of the credit prospects of a particular borrower than the market does, they will want to discourage relationship managers from extending new credit. But this does not necessarily involve quoting internal prices that reflect wider credit spreads than those quoted in the CDS market. To the extent that the CDS market for the borrower is liquid, the internal quote may just reflect the CDS spread, with the credit managers intending to purchase CDS protection against any new credit extensions to the borrower. But this strategy must be accompanied by some type of limit on the amount of lending that can be offered at this market price, gauged to the liquidity of the market.

Internal price quotes more favorable than market quotes are not confined to the situation in which credit managers have a favorable view on a borrower. It may also reflect the composition of the firm's credit portfolio. If lending to a particular borrower offers diversification of the portfolio, perhaps in terms of industry or geography, this may get reflected in a lower internal price for credit than the market. And conversely, if lending to a particular borrower will add to existing portfolio concentrations, the credit managers may quote an internal price equal to the market price and plan to offset the loan with CDS protection, even when they have a more favorable view of the borrower than the market does. This approach clearly requires reporting from the credit portfolio model that can assess the marginal impact new lending to a particular borrower will have on the firm's overall credit risk. We will address this information need shortly, when we discuss credit portfolio reporting requirements.

Given our discussion of the large differences between actual and market-implied default probabilities toward the end of Section 13.2.1.2, one might suspect that this would motivate relationship managers to prefer internal loan pricing to market loan pricing. In my experience, this does not turn out to be much of a factor. The reason is that the internal pricing's charge for capital usage is roughly equal to the difference between default experience and market-implied probabilities. Looking at Table 13.5, you

can see that the difference between average and worst-case default experience is greater than the average default experience for every credit rating above B. Since allocated capital is roughly comparable to this difference, and since charges against earnings for capital allocations are on the order of 15 percent per year, you can see that internal pricing is unlikely to look more favorable than market pricing just based on this factor. Rather, it is differences between market and internal assessments of credit quality for a given name and portfolio composition considerations just discussed that usually drive preferences between internal and market pricing.

Risk reporting for portfolio credit risk is similar to risk reporting for market risk VaR, and many of the recommendations of Sections 7.1.2 and 7.3 can be followed with little modification. The reporting of VaR at different percentiles, the use of shortfall VaR as an alternative to VaR to better capture tail risk and to avoid issues of instability and negative diversification effects, the reporting of exposures by product and business line, and the use of both marginal and stand-alone measures of risk all carry over quite well to portfolio credit risk. Reporting of exposures by product and business line will help to identify lines of business that should be expanded or whose growth may need to be slowed and to identify priorities for parts of the portfolio that require hedging through loan sales, CDSs, and CDOs. More thoughts along these lines will be found in the section on “Improving Portfolio Performance” in Bohn and Stein (2009, Chapter 8).

In making decisions among competing priorities for portfolio hedging, credit managers will need to consider the risks of delay not just in terms of possible defaults, but also in terms of ratings downgrades and widening of market credit spreads. So, while I would still insist on the illiquidity of credit portfolios requiring output based on eventual defaults (see Section 13.3.2), output based on the impact of changes in market value is also needed. This output should also include reports on market value sensitivity to shifts in the economic environment, along the lines of the sensitivity measures for liquid credit discussed in Section 13.1.3. O’Kane (2008, Section 17.2) has an extensive analysis of sensitivity measures for credit portfolio risk in the context of CDO tranches, which will be discussed in Section 13.4.3. The caveats addressed there about the limitations on the usefulness of these sensitivity measures due to illiquidity of credit markets also apply here.

One major difference between market risk VaR and credit portfolio risk models is the need to measure the marginal risk contribution of individual loans, in line with our previous discussion. No marginal measure that is this fine-grained is required for market risk. Since it would be computationally infeasible to generate this information by running the full simulation for each loan, finding computational shortcuts, addressed in Section 13.3.3, is critical for credit portfolio modeling.

13.4 RISK MANAGEMENT OF MULTINAME CREDIT DERIVATIVES

13.4.1 Multiname Credit Derivatives

Multiname credit derivatives are baskets that bundle together credit exposure to several debt issuers into a single instrument. From the side of credit protection sellers, these instruments offer an opportunity to obtain exposure to a diversified basket of corporate debt—it can be quite difficult for an investor to put together such a basket on his own, owing to the relative illiquidity of corporate bond markets. So a market maker who has the ability to source such a basket of debt can get paid a good spread for selling it in a convenient form. From the side of credit protection buyers, these instruments offer a chance to offset portfolio credit risk and can play a significant role in the management of portfolio credit exposure discussed in Section 13.3.4.

There are several forms such credit baskets can take. One that has become particularly popular is to create a derivative tied to a credit market index such as the CDX and iTraxx indexes (see the discussion in Hull (2012, Section 24.3). Since these indexes are calculated and disseminated in a very public manner, they provide good transparency. A spread is set that the credit protection buyer will pay to the credit protection seller. Every time one of the components of the basket defaults, settlement is made on that portion of the basket, following the same rules as settlement of individual CDSs (with a strong bias toward cash settlement). The choice of index components has been designed to balance liquidity of individual names and diversity of credit exposure. Many different strategies for expressing views on relative credit spreads and achieving protection against existing credit risks can be obtained through combinations of long and short exposures to different credit indexes and individual CDSs. Market makers will also construct more tailored indexes for clients with particular needs.

As long as all credit protection sellers in a credit basket share equally in all cash flows (both receipt of credit spread and payment of default losses), the credit basket is just a simple summation of the pricing of individual CDSs and so should be risk managed following the principles of Section 12.4.1.1; there are some technical basis risk issues that create differences between the index basket and a portfolio of the individual CDSs comprising the basket—these are well covered in O’Kane (2008, Chapter 10). But a frequent variant is to structure a credit basket into tranches that receive different portions of the cash flows. The motivation for this structuring is that different classes of investors have different tolerance for credit risk and different institutional restrictions on the type of credit risk they can invest in, and the object is to design a structure that will better fit demand.

Some CDOs are structured this way, but others are just single tranches (called *synthetic tranches*) that have been agreed to between a credit protection buyer and a credit protection seller with an agreed reference portfolio. Credit protection buyers may enter into these single tranches either as a way of buying protection against their credit exposure in a piece-by-piece fashion or in order to express a particular view on credit risk. Credit protection sellers may want to express a particular view on credit risk or offset previously purchased credit protection they no longer need. The modeling and risk management of tranches is very similar, whether they arose as part of a credit basket or as a synthetic tranche, though there are some differences in payment details that will impact modeling; see O’Kane (2008, Section 12.5).

Tranches of CDOs are structured by designating an attachment point and a detachment point. For example, a tranche with an attachment point of 7 percent and a detachment point of 10 percent would not have to make any default payments until default losses in the basket exceed 7 percent of notional principal and then would pay all default losses until default losses in the basket reach 10 percent of notional principal, after which time the tranche no longer exists (it is no longer receiving any credit spread payments and does not owe any obligations on any possible future defaults). In between 7 percent and 10 percent, every time a default takes place, settlement is made on that portion of the tranche, with credit spread payments on that portion ceasing. By market convention, credit spreads paid to the protection seller of a tranche are quoted as percentages of the portion of notional principal that the seller could potentially lose. For example, if an investor sells 7 percent to 10 percent credit protection on a \$100 million basket, his largest potential loss would be $\$100 \text{ million} * (10\% - 7\%) = \3 million . If his credit spread is quoted at 3.47 percent, he will receive $3.47\% * \$3 \text{ million} = \$104,100$ per year until such time as default losses exceed 7 percent. Standardized tranches have been created for both the CDX and iTraxx indexes (see Hull, Table 24.6). The tranche that will receive the first losses, the tranche with 0 percent attachment point, is called the *equity tranche* (since its absorption of losses prior to any losses impacting other tranches is similar to the relation between the equity investors in a corporation relative to the debt holders). The tranche with the highest attachment point is called the *super-senior tranche* (since its expected losses are usually even smaller than the highest-rated AAA corporate debt). Intermediate tranches are called *mezzanine tranches* for those with lower attachment points and *senior tranches* for those with higher attachment points.

Tranching cash flows from a credit basket introduces a new type of risk that did not previously exist in the basket—exposure to default correlation. This can be illustrated by a simple example. Suppose you have a credit basket on which your expected default losses, net of recovery, over its five-year

life are 3 percent of principal. If you assume a very low level of correlation between defaults, then almost all scenarios will involve some group of companies defaulting and very few will involve a large number defaulting. So a 0 percent to 3 percent equity tranche will almost always lose close to its maximum and a 15 percent to 30 percent super-senior tranche will experience zero losses. By contrast, at a very high level of default correlation, some scenarios will involve almost no losses while some will incur very heavy losses. So a 0 percent to 3 percent equity tranche will sometimes lose less than the maximum and so have lower losses than under the low correlation assumption while the 15 percent to 30 percent super-senior tranche will sometimes experience losses and so have higher losses than under the high correlation assumption. This pattern always holds—higher correlation means lower losses for any tranche with a 0 percent attachment point and higher losses for any tranche with a very high attachment point, but you can't tell in advance of detailed calculations how an intermediate tranche will be impacted.

One variant of tranching CDOs is to allocate losses based on number of defaults rather than on the percentage of losses in the portfolio. It is called a *default basket*. For example, a first-to-default tranche absorbs all the losses of the first credit in a basket to default (if any) but loses nothing on any subsequent defaults. Default baskets make sense only when based on a relatively small number of individual credits, anywhere from 2 to 10. O'Kane (2008, Section 12.2) explains the mechanics and basic economics of this product. While its modeling and risk management are closely related to those of standard CDOs, LHP approximations do not make sense, given the small number of credits in the underlying portfolio and the digital nature of loss allocation. Approximation methodology such as the first of the two models in Hull and White (2004), referenced in Section 13.3.3, are more appropriate for default baskets.

More complex variants of CDOs, such as CDO-squareds, constant proportional debt obligations, and options on tranches, will be covered only briefly as part of the next section on modeling of multiname credit derivatives.

13.4.2 Modeling of Multiname Credit Derivatives

The modeling of multiname credit derivatives is extremely similar to the modeling of portfolio credit risk covered in Section 13.3. Indeed, many of the techniques discussed there were originally developed in support of analysis of CDOs and CDO tranches. This is not surprising, since multiname credit derivatives represent an attempt to provide a market for the transfer of portfolio credit risk.

A first-cut sketch of the model for a multiname credit derivative instrument would therefore be to start with the portfolio of credits that comprise the underlying basket referenced by the instrument, model the losses in this portfolio using the tools of Section 13.3, and in this modeling keep track of which losses accrue to which tranches. You can see a simple illustration for a single-factor Vasicek model in the CDO spreadsheet on the website for this book. At each level of the single factor that drives the correlation between defaults of underlying credits, the spreadsheet keeps track of how much loss accrues to each tranche of the CDO. It is then easy to compute a full probability distribution of the losses for each tranche. An exposition of this simple model can be found in O’Kane (2008, Section 16.2) along with an analysis of the sensitivity of model results to input parameters in Section 16.3.

In practice, the allocation of losses to tranches may follow complex rules. This is particularly true for tranches of CDOs based on mortgage securities. So modeling of tranche losses requires more sophistication in the simulation of each individual path. Smithson and Pearson (2008) addresses this issue. More complex multiname credit derivatives may require more detailed modeling of the evolution of losses over time. O’Kane (2008) gives a brief introduction to these products (constant proportional debt obligations in Section 22.3, forward-starting tranches in 22.6, and options on tranches in 22.7) along with an introduction to the more detailed models of evolution of losses in Chapters 23 and 24.

A product that became popular in the explosion of subprime mortgage-based CDOs was the multilevel CDO in which tranches of different CDOs are bundled together to form a portfolio that can itself be tranced, called a CDO-squared (and this process can be repeated to form a CDO-cubed, and so on). The same fundamental modeling approach can be utilized as for single-level CDOs, modeling losses and keeping track of which losses accrue to which tranches and tranches of tranches. However, the computational intensity of keeping proper track of this waterfall may require new techniques for efficient approximation. O’Kane (2008, Section 22.4) is a good introduction to these products and their modeling. As O’Kane illustrates in Figure 22.4, CDO-squareds have tremendous sensitivity—very large changes in the losses to tranches due to very small variations in losses to the underlying credits. In the wake of the 2007–2008 crisis, when a great many AAA-rated tranches of CDO-squareds of subprime mortgages suffered close to 100 percent losses, these products came in for harsh criticism as vehicles for inappropriate levels of leverage. It is doubtful that we will ever see a revival of interest in them (a parallel story to the power options whose unsuitable use in the early 1990s led to the virtual death of the product ever since—see the discussion toward the end of the introductory section in Section 12.1).

Credit portfolio managers are typically dealing with just a single firmwide portfolio. By contrast, traders in CDOs and other multiname credit derivatives are typically dealing with a large number of different reference portfolios. This makes computational alternatives to full simulation, which we covered in Section 13.3.3, even more critical to CDO traders than they are to credit portfolio managers. As we already noted in Section 13.3.3, it was in the context of CDO modeling that many of these computational alternatives were first developed. It also means that CDO traders will typically have a far less intimate knowledge of the characteristics of any particular reference portfolio they are dealing with than a credit portfolio manager will have of her portfolio. This should raise a note of caution concerning the accuracy of CDO modeling, to which we will return in the next section on CDO risk management.

A critical difference between credit portfolio modeling and CDO modeling is that CDO modelers will frequently be trying to fit to market data on where different CDO instruments are trading. The only market data that was considered in our discussion of credit portfolio modeling in Section 13.3 was market data on individual credits; all input on relationship between defaults of different borrowers came from statistics and subjective probabilities. While CDO traders and risk managers must also be aware of the implications of statistical and subjective estimates, input from markets is vital.

Let's consider a typical situation. A trading desk is asked to price a nonstandard tranche on a particular portfolio (by nonstandard, we mean having different attachment and detachment points from more commonly traded tranches). The desk can obtain market prices for standard tranches on the same portfolio. To price the nonstandard tranche, the traders would like to fit the parameters of a pricing model to correctly price all of the standard tranches and then apply the model to a nonstandard tranche. This very closely fits the interpolation model approach of Sections 8.2.6.1 and 8.3. It would be convenient if a simple model, such as the Li model, could fit the standard tranche prices, but this is virtually never possible. The reasons why sound very much like the reasons we discussed in Section 11.6.2 for why a single implied volatility is unlikely to fit market options prices for several different strikes.

For vanilla options, as we saw in Section 11.6.2, it is partly because of different market supply and demand pressures for different strikes, and it is partly because some of the assumptions of the Black-Scholes model are incorrect. The story is similar for CDOs. Tranches with very low attachment points (equity tranches) and tranches with very high attachment points (super-senior tranches) are far less popular with buyers of credit risk than are tranches with intermediate attachment points (mezzanine tranches). We discussed some of the reasons for this in Sections 5.2.2 and 5.2.5, in the context of CDOs of mortgages; similar reasons apply to CDOs of corporate

loans. As for model assumptions, the assumption of a Gaussian copula is often contradicted by historical data (see Section 13.3.1).

There are two basic approaches to fitting market prices for tranches. One focuses on finding a model that more accurately reflects statistical relationships between defaults; the second focuses on pragmatically changing model parameters to achieve a fit. The first has the advantage of trying to build in more economic reality and so is likely to be a more robust model than one that just fits to prices. But the first has the disadvantage that even the most realistic model may not be able to account for supply and demand forces in the market.

The more pragmatic approach of just fitting market prices is a very close analogue to utilizing volatility smiles and skews in fitting option prices; whatever supply and demand dictates determines the implied volatility input for each strike and tenor, and options that can't be directly priced in the market have their volatilities interpolated from those that are directly priced, as in Sections 11.6.1 and 11.6.2. The analogous method for CDOs is called implied correlation skew to parallel the implied volatility skew. O'Kane (2008, Chapter 19) explains this approach in detail. To make the fitting more manageable and arbitrage-free, it is extremely helpful to do all fitting to base tranches, tranches with an attachment point of 0 (i.e., tranches that absorb all losses up to the detachment point). This approach, known as *base correlation*, is explained in detail in O'Kane (2008, Chapter 20). Standard base tranche prices can always be constructed directly from standard tranche prices (e.g., a 0 percent to 10 percent base tranche is just the direct summation of a 0 percent to 3 percent tranche, a 3 percent to 7 percent tranche, and a 7 percent to 10 percent tranche). This approach closely parallels one that has been in use for years in the vanilla options market, where all fitting of volatilities by time period is done for time periods starting at the current date; this avoids interpolations that produce negative implied volatilities and makes for smoother fits.

The more fundamental approach of finding a model that more accurately reflects statistical relationships has a vast multitude of candidate models—at least as many as the different ideas on alternative copulas discussed in Section 13.3.1. O'Kane (2008, Chapter 21) addresses some of the more popular choices for copulas, and discusses issues of calibration and comparison between models.

13.4.3 Risk Management and Reporting for Multiname Credit Derivatives

We will begin with two polar views of risk management and related reporting requirements for multiname credit derivatives and then see how the two

views can be blended. At one extreme, we will focus on the fact that holding a CDO position is very similar to holding a credit portfolio position, so the risk management should look very similar to the approach to risk management of portfolio credit given in Section 13.3.4. At the other extreme, we will focus on the greater liquidity of credit derivatives and look to a risk management approach closer to that for other liquid derivatives in Sections 11.4 and 13.1.3. Which approach will have the greater weight in a blended view will depend a lot on just how liquid the market is for multi-name credit derivatives.

Even if we believed that multi-name derivatives were completely illiquid, we would still need to modify the portfolio credit risk approach of Section 13.3.4 to account for the fact that in a derivatives book your positions encompass sales of credit portfolios as well as purchases. But the basic principle would remain the same: simulate default losses all the way to the maturity of positions and look at the full distribution, both expected losses and tail losses. But there are a few additional points to be considered:

- Some credits may be referenced in multiple tranches. And for multilevel CDO products, such as CDO-squareds, the same tranche may be referenced in multiple higher-level tranches. The simulation engine must have the capability of identifying these multiple references and treating them properly; they must show 100 percent correlation on defaults.
- Traders in CDO tranches will often lack the detailed knowledge about individual underlying credits that a credit portfolio manager would have for credits originated in the firm. The credit portfolio manager might still use a faster simulation method such as an LHP approximation, as addressed in Section 13.3.3, but has the capability of checking the accuracy of the approximation by occasional comparison of results to a full simulation, and then adjusting risk reports to reflect the accuracy of the approximation. A CDO trading desk manager using the same LHP methodology may lack the detailed data on individual underlying credits that would allow the accuracy of the approximation to be checked. Some alternative means of adjusting risk reports for the inaccuracy of approximation must be designed, such as simulating several different possible specifications for the data on underlying credits, calculating the approximation error that each would lead to, and basing a conservative estimate of approximation error on these results.
- Some adjustment in probability distributions would also be appropriate to reflect the lower certainty regarding estimates of default probabilities and loss given default that is associated with less detailed knowledge of individual credits; compare with Rajan (2010, 128–129).

A trader dealing in liquid CDO tranches would want to start with a set of risk measures and limits that looked a lot closer to those of Section 13.1.3, with a focus on exposure to changes in market credit spreads, but supplemented by measures of convexity exposure to large jumps in credit spread and default. But this would need to be modified to take exposure to correlation into account. If the tranches are truly liquid, then it should be possible to manage correlation risk in a manner very close to the management of option risk in Section 11.4, with measures of exposure to changes in correlation levels as well as changes in the shape of the correlation surface (by time bucket and attachment point) and to joint changes in credit spread and correlations. O’Kane (2008, Chapter 17) has a detailed discussion of risk reporting following this approach.

The reporting and risk management of a multiname credit derivative portfolio needs a blend of these two approaches, based on actual degree of liquidity. But no matter how liquid the derivatives in the portfolio, some weight should always be given to the approach of Section 13.3.4, since this is the approach best designed to deal with the impact of defaults. This is a parallel point to the one made in Section 13.1.3 on risk management and reporting for single-name credit instruments; the extreme difference between exposure to credit spread movements and exposure to defaults, illustrated in Section 13.1.2.2, necessitates two different reporting frameworks.

Whatever approach is being taken to risk management of CDOs, there needs to be a strong awareness by both traders and risk managers of the extreme sensitivity of some CDO tranches to systematic risk and to changes in assumptions. This will be highlighted in the next section.

13.4.4 CDO Tranches and Systematic Risk

Among the general principles for risk management in Section 6.1.1 was the need for risk managers to carefully distinguish between systematic (undiversifiable) and idiosyncratic (diversifiable) risks. Earlier in this chapter, we noted the strong impact of systematic risk on the pricing of credit exposure (Section 13.2.1.2). This becomes a particularly important issue for the most senior tranches of CDOs, because tranching has the effect of concentrating the idiosyncratic risk of the reference portfolio in the more junior tranches and concentrating its systematic risk in the more senior tranches. This effect becomes even more pronounced for the senior tranches of CDO-squared products. These issues are very cogently analyzed in Coval, Jurek, and Stafford (2008).

One way to understand why this happens is to see that there are likely to be some defaults in the reference portfolio regardless of the state of the economy. So the amount of loss in the tranches that absorb the first losses

is likely to be as dependent on the idiosyncratic risk arising from exactly which credits are in the portfolio as it is on the state of the economy. But losses will reach the very senior tranches only in situations where the common economic factor suffers a major negative event. A useful analogy would be a put option purchased as protection against a large decline in a stock index; it will pay off only if there is a severe shock to the economy. But just as we saw in Section 11.6.2 that protection buyers tend to strongly bid up the implied volatility on such put options, we can anticipate that senior tranches of CDOs should be priced at steep premiums to expected losses. Closely related points, highlighted by Coval, Jurek, and Stafford (2008), are that senior tranches have very high volatility of returns and very strong sensitivity to model assumptions. And all of these points apply with even more force to senior tranches of CDO-squareds.

As we noted in Section 13.3.3, a major advantage of the Vasicek model is its ability to build intuition concerning the allocation of systematic risk to tranches. Exercise 13.3, using the CDO spreadsheet on the website for this book, allows you to use the Vasicek model to generate measures of systematic risk and volatility for tranches of a CDO. One point that will be made in this exercise is that the reasonableness of correlation inputs to the Vasicek model can be judged by comparing model results to historical default experience, utilizing data such as that presented in Table 13.5.

EXERCISES

13.1 Calculating default rates from bond rates

Using the CreditPricer spreadsheet, begin with the following input:

	Risk-Free Zero-Coupon Rate	Risky Par Rate
1	7.00%	8.00%
2	7.50%	8.60%
3	7.75%	8.90%
4	8.00%	9.20%
5	8.15%	9.40%

1. Solve for the default rates and spreads to the risk-free par curve that corresponds to this case.

2. Change the loss given default to 30 percent and double the default rates. Solve for the risky par bond rates. How does the spread to the risk-free par curve differ from that in the previous step? This shows that it is not just the product of default rate and loss given default that impacts the valuation of risky cash flows.
3. Assume that the company whose risky par rate curve was shown previously also has a five-year bond with a 9 percent coupon that is priced in the market at 98.56. Assuming a constant loss given default irrespective of the time at which default occurs, determine a unique loss given default and a set of default rates from this information. What if the 9 percent coupon five-year bond is selling at 98.46?

13.2 Comparing the jump process credit model to the Merton model

1. Run the **MertonModel** with a stock price of 40, debt per share of 60, equity volatility of 60 percent, and time to maturity of five years. What are the resulting probability of default and loss given default?
2. Run the **JumpProcessCredit** model with the same stock price and equity volatility as you used for the **MertonModel** with a risk-free rate of 5 percent, a loss given default of 60 percent, and a standard deviation of the default barrier of 50 percent. Try different input values for the default barrier level and see what the impact is on the probability of default and the credit spread for a five-year maturity.
3. Prepare an analysis comparing the two models in terms of the impact on probability of default for changes in the stock price and changes in the equity volatility.

13.3 Using the Vasicek model for risk measurement of CDO tranches

1. Set the CDO spreadsheet to run the Vasicek model with Gaussian copula (i.e., set all tail factors and correlation factors to 100.00%). In this exercise we will just be experimenting with default rates, so we will not reduce input loss rates for assumed recoveries.
2. Assume that you have a portfolio of Bb loans. Using results from Table 13.5, set the input loss rate to 9.73%. Experiment with

different input correlation rates to see the impact on the standard deviation and 2.45th percentile losses for the portfolio. Notice that very low input correlation rates produce standard deviations and 2.45th percentile losses for the portfolio that look unrealistically low relative to historical experience, and that very high input correlation rates produce the opposite effect (for example, a 1% input correlation rate produces a portfolio standard deviation of 1.69% and a 2.45th percentile loss of 13.48%, while a 20% input correlation produces a portfolio standard deviation of 8.31% and a 2.45th percentile loss of 32.38%; Table 13.5 shows the historical standard deviation of Bb loan defaults to be 6.50% and maximum loss over any 5 year period to be 23.44%).

3. Through experimentation find an input correlation rate that produces reasonable results relative to the historical Bb loan default standard deviation and maximum loss.
4. Continuing with this example, experiment with different tranche attachment points to find one that will produce expected losses as a percentage of investment in the most senior tranche (the tranche with a 100% detachment point) roughly equal to the historical 0.27% loss rate for Aa loans from Table 13.5. Then compare the standard deviation of losses and 2.45th percentile loss as a percentage of investment for this senior tranche with the historical standard deviation of 0.44% and maximum loss of 1.83% for Aa loans from Table 13.5. You should see that even though the expected losses of the senior tranche match historical losses of Aa loans, the standard deviation and “worst case” losses are considerably higher for the senior tranche than they are for a portfolio of Aa loans. This illustrates the point made in Section 13.4.4 about the impact of tranching on concentration of systematic risk.
5. Further exercises with the CDO spreadsheet could involve experimenting with attachment and detachment points to try to create tranches that match other credit classes in terms of expected loss as a percentage of investment. You can also experiment with the impact of creating fatter tails than the Vasicek model by using input tail factors and correlation factors higher than 100%.

Counterparty Credit Risk

Counterparty credit risk management arising from derivative contracts is an extremely important piece of the management of credit risk for reasons discussed in Section 14.1. Since the first edition of this book was published, the first full-length book treatment of counterparty credit risk, written by one of the leading practitioners in this field, Gregory (2010), has appeared. I will be making frequent reference to this book in what follows and will provide several suggestions for further reading in Gregory that will provide greater detail and examples for points I will raise.

14.1 OVERVIEW

For credit risk, derivatives represent a two-edged sword. On the one hand they have been valuable tools in reducing credit exposure, but on the other hand the use of derivatives leads to the buildup of credit exposure. The hope is that exposure reduction outweighs exposure buildup, but, without careful management, the full potential for credit exposure reduction by derivatives use will not be achieved.

When financial derivatives markets first began to grow in the 1970s, the growth was primarily in currency and interest rate derivatives, and this remains the largest use to the current day (over 85 percent of the notional amount of contracts outstanding, according to figures from Tables 19 and 23A in the Bank for International Settlements' December 2011 Derivatives Statistics). One use of these derivatives was to take on market exposures that could previously be accomplished only by cash instruments, such as loans, bonds, and deposits. As can be seen from Section 10.1.3 and Table 10.2, derivatives minimize the credit exposure and funding requirements entailed by loans, bonds, and deposits.

The management of counterparty credit risk has been accomplished by two very distinct, but related, approaches: the use of derivatives exchanges

and the credit management of over-the-counter (OTC) derivatives that are not traded on exchanges. We will first discuss credit risk management through exchanges in Section 14.2 and the credit risk management for OTC derivatives in Section 14.3. The clear failure of many firms in managing their credit exposure on OTC derivatives (discussed in Section 5.3.1) has led to increasing pressure to move as much counterparty credit risk to exchanges and away from OTC as possible. The potential of and possible problems with this approach were discussed in Section 5.5.7.

To see the significance of credit risk generated by derivatives, consider that U.S. commercial banks had \$281 billion of credit exposure related to derivatives contracts at the end of 2011, about a quarter the size of their \$1,339 billion credit exposure in traditional commercial and industrial loans (figures taken from the Federal Reserve's H8 report). To see the impact of management of credit risk on derivatives credit exposure, as of June 2011 global credit exposure on all over-the-counter derivatives contracts was \$19.5 trillion, compared with \$707 trillion of notional outstandings, and there were another \$83 trillion in notional outstandings on exchange-traded derivatives contracts, on which there should be virtually no credit exposure, as we will see in Section 14.2 (all figures from Tables 19 and 23A in the Bank for International Settlements' December 2011 Derivatives Statistics).

14.2 EXCHANGE-TRADED DERIVATIVES

Counterparty credit risk management of exchange-traded derivatives rests on five key concepts: *novation*, *margining*, *closeout*, *netting*, and *loss mutualization*. The most important of these concepts is novation. As soon as two counterparties (let's call them A and B) agree to a derivative contract traded on an exchange, the contract between the two counterparties is immediately canceled and replaced by two contracts, one between A and the exchange, and the other between the exchange and B; see Gregory (2010, Section 14.1.5) for details.

Neither of the two counterparties needs to have any concern with the credit risk of the other—each has a contractual relationship for delivery on the derivatives contract with the exchange, and the exchange always has very low credit risk because it has the backing of all its members (we'll discuss this further under loss mutualization), because it takes no market risk, and because it carefully controls its credit risk.

To keep the discussion simple in what follows, I will write as if exchanges deal directly with all counterparties. Actually, a typical exchange has two classes of counterparties: exchange members who share in loss

mutualization, and all others. It is only an exchange member who is permitted to be the direct counterparty of the exchange. All other counterparties are actually counterparties of one of the exchange member firms, which places trades with the exchange on behalf of these counterparties. But since exchange members manage their counterparty risk by the exact same method that the exchange handles its counterparty risk, through margining and closeout, a unified description is not too far removed from actual practice. At the end of this section, we discuss the extra detail that is needed to account for the two-tier reality of exchanges and members. Also, in the interest of simplicity, I will always refer to the contracts as being with the exchange, ignoring the possible distinction between the exchange and its affiliated clearinghouse; see Gregory (2010, Section 14.1.3) for details.

The exchange takes no market risk because its only positions arise as a result of novation and hence are always exactly offsetting positions. For example, if A contracts to deliver 100 million dollars to B in exchange for B delivering 70 million euros to A on a certain future date, this contract is replaced by A contracting to deliver 100 million dollars to the exchange for 70 million euros and B contracting to deliver 70 million euros to the exchange for 100 million dollars, both on the same date. So as long as both A and B perform their contractual obligations, the exchange will have no gain or loss, no matter what happens to the dollar/euro exchange rate. Hence, the exchange never bears any market risk.

By contrast, the exchange must be very concerned about counterparty credit risk, since each trade leaves it with credit exposure to both parties of the trade. The exchange manages this credit risk through a very well-defined system of margining, closeout, and netting. The exchange is continuously monitoring the mark-to-market position of every trade, and any mark-to-market losses require a counterparty to immediately pay cash to the exchange to cover the loss (the exchange doesn't keep this cash; it pays it to the counterparty with an offsetting mark-to-market gain). Any time a counterparty fails to provide the cash required to cover a mark-to-market loss, the exchange will declare the counterparty in default and close out all of the counterparty's positions with the exchange. In this closeout, all of the counterparty's positions, whether gains or losses, are netted against one another. The exchange seeks new counterparties to take over these positions. The exchange's losses on these positions are limited to the change from the mark-to-market price the last time the defaulting counterparty posted margin and the price at which a new counterparty is willing to trade. The exchange has three ways in which to cover these losses:

1. First, each counterparty must post with the exchange initial margin at the time it first enters into a trade (this does not need to be in cash;

it can be some high-quality security such as a Treasury bond). Losses in closing out a defaulted position will be charged against this initial margin before any money or securities are returned to the defaulting counterparty.

2. Second, if losses exceed the initial margin, the exchange will sue the defaulting counterparty for the remaining loss. However, recovery may be limited if the defaulting counterparty is actually bankrupt as opposed to just suffering temporary problems in meeting a margin call.
3. Third, any remaining losses are shared among all the members of the exchange. This is the principle of loss mutualization.

In evaluating how much initial margin an exchange should demand to protect itself against the possibility of default, a key factor is to estimate probability distribution of price changes between the last mark-to-market and the transaction with a new counterparty. This depends crucially on the price volatility of the contract, the liquidity of the contract, and the speed with which the closeout mechanism operates. The more liquid a contract (i.e., the more frequently it trades and the larger the size of trading that occurs), the more confidence the exchange can have that the mark-to-market is close to the actual price at which a new trade can be done, and the lower the chance that the forced trading the exchange will do to close out the defaulted position will impact the price.

As we already noted in Section 6.1.1, the management of counterparty credit risk through margining can follow very closely the prescription we have detailed for the market risk of trading: the importance of timely and accurate mark-to-market (Section 6.1.3), value at risk (VaR) (Section 7.1), and stress test (Section 7.2) calculations. In particular, VaR simulations and stress testing should look almost identical to the discussion in Chapter 7. As with trading positions, VaR will focus on losses that might occur under conditions of normal market liquidity, while stress tests will look at losses that might occur over longer periods between closeout and replacement with a new counterparty that result from unusual conditions of market illiquidity.

There are two critical differences between the management of counterparty credit risk on exchanges and the management of the market risk for trading desks that impact VaR calculation methodology. One is that there may be a significant delay between the failure of a counterparty to meet a margin call and the declaration of default (this is called the *grace period*); it may take time for the exchange to confirm that a counterparty truly cannot or is choosing not to meet a margin call, rather than just a delay caused by an operational error or a communication failure. The time that is necessary to make this determination must be built into the VaR calculation, since it

is a time period during which prices may fluctuate. Exchanges try to minimize required initial margin, since this is a key factor in the competition for business, and so will try to minimize the grace period. For example, as pointed out in Gregory (2010, Section 14.1.8), large price movements might trigger intraday margin calls, a practice that is becoming increasingly common and is supported by technology advances. But closing out too quickly may also result in loss of business to competitor exchanges, since it will unduly penalize operational errors.

The second critical difference is that trading desks are experienced in managing market risk positions, and so can be expected to skillfully manage the required closing out of a position. By contrast, exchanges by their nature are not expected to have model risk positions, so closing out a position is not a task they are well positioned for. Exchanges protect themselves by limiting the number of contracts they will trade to a standardized set (e.g., allowing trading for only four settlement dates each year; see Hull 2012, Section 2.2, “Delivery Months”). By utilizing a limited set of standardized contracts, exchanges cultivate liquidity for each contract traded, making marking to market more robust and closeout easier to perform.

Once VaR and stress test computations have been made, an exchange will be in a good position to evaluate the adequacy of initial margin requirements and to estimate the probability that the initial margins will prove insufficient to cover the losses incurred in a closeout. Some of the considerations that will go into evaluating the required size of initial margins are (compare with Gregory 2010, Section 14.1.8):

- The volatility of prices for the particular contracts involved and the length of the grace period, both of which should be direct inputs to the VaR and stress test computations.
- The degree of offset likely between netted positions in different contracts. This also should be an integral part of VaR and stress test computations, but with the same concern for the reliability of historical correlation relationships under stressed market conditions discussed in Sections 7.2.2 and 7.2.3.
- The size of the counterparty’s position relative to the size of trading in the contract. This is a point very similar to that raised in Section 6.1.4 regarding positions that are illiquid due to size. The remedy should be similar to that proposed in Section 6.1.4: simulation of price change between last mark-to-market and completed closeout should be over a longer time period to accommodate the larger position.
- The degree to which a counterparty has financial resources beyond its trading positions. This will impact the likelihood that losses could be recovered through a lawsuit.

- The degree to which a counterparty's losses will tend to be correlated with those of a significant number of the exchange's other counterparties. This might require VaR and stress test calculations that look at the whole universe of counterparties, rather than just one at a time.

The methodology that exchanges use to manage counterparty credit risk through margining and closeout offers both drawbacks and advantages to counterparties. On the negative side is the narrow range of allowed contracts, which limits the degree to which derivatives can be tailored to meet specific needs of a customer. Also on the negative side is the operational complexity of meeting continuous margin calls. On the positive side, the heavy reliance of this approach on controlling credit risk through the actual mechanism of trading reduces reliance on credit evaluation of each customer. This can be very attractive to some customers who might not have the track record needed to withstand a credit review but who have confidence in their ability to manage margin calls. Another positive is that since the exchange has no market position, it has no incentive to hide information about prices at which trades have occurred and the depth of the market. Exchanges typically supply a much greater range and quality of price and market size information than do trading desks that are also holding market positions. Not only do exchanges generally provide complete public information on the sizes and prices of all executed trades, but “in typical exchange-traded markets . . . the best available bid and offer are provided to nearly all market participants nearly instantly” (Duffie, Li, and Lubke 2010). One further negative that must be considered is that exchanges may protect themselves in instances of extreme market volatility by imposing limitations on trading that disadvantage some customers (this point is made forcefully in the section on clearinghouses in Brown (2012, Chapter 10).

A very important positive of the exchange counterparty credit methodology is the ease with which a counterparty can offset a position previously entered into. As time and circumstances change, it is very common to wish to reverse a previous transaction. If your contract is with a private firm, as in an over-the-counter derivative, you must negotiate with this firm to offset the prior position. If your counterparty still wants to keep the position, you have a choice of either offering price concessions to induce your counterparty to offset the position or entering into an offsetting position with a new counterparty, which would offset the market position but leave you with credit exposure to both your original counterparty and the new counterparty. By contrast, the novation feature of exchange-traded derivatives makes offset easy. Since your counterparty on any transaction is the exchange, you can find any new counterparty wanting to enter into an offsetting position

and this will result in the complete cancellation of your original position with the exchange, leaving both you and the exchange with no further credit exposure on the original position or on your new offsetting position. (To make this completely clear, if the original position was between A and B, and later A enters into an offsetting position with C, A will be left with no exposure and the exchange will have offsetting positions with C and B, replacing its original offsetting positions with A and B.)

As a final point, let us account for the actual two-tiered nature of exchanges. As we said toward the beginning of this section, we have been simplifying by writing as if exchanges deal directly with all counterparties. In fact, it is only an exchange member, one who shares in loss mutualization, who is permitted to be the direct counterparty of the exchange. All other counterparties are actually counterparties of one of the exchange member firms, which places trades with the exchange on behalf of these counterparties. When a customer requests a trade through a member, the member is obligated to make that trade on the exchange, so members do not accumulate any market positions with customers. The exchange only needs to manage its credit exposure to its members, while each member needs to manage its credit exposure to its customers. The description we have given thus far, of margining, netting, closeout, VaR, and stress test calculations all apply equally to the exchange's management of its credit exposure to members and to members' management of their credit exposure to customers. If a customer's position requires a margin call by the exchange, it is the member that is obligated to meet the exchange's margin call, and the member in turn will make a margin call to the customer.

From the viewpoint of a customer of a member, there shouldn't be any difference between placing trades through the exchange and the actual placement of trades through a member—the obligation to pay the customer is the exchange's obligation. The exchange will make payments due on mark-to-market increases to the member firm, which is in turn obligated to pass these payments on to the customer. The only potential problem would be if the member does not adequately segregate customer funds from its own funds; in this case, if the member goes bankrupt, the customers could lose on initial margin accounts being kept with the member along with any funds the customer kept in excess of required margin, perhaps as an operational convenience to meet future margin calls. This was considered a remote possibility, given exchange rules and legal requirements for member firms. But the 2011 bankruptcy of MF Global and its failure to segregate customer funds left customers with long delays in access to funds and the definite potential for ultimate loss of part of their margin accounts (see Koutoulas and Roe 2012). It remains to be seen what impact this will have on customer views of the safety of exchange-traded derivatives.

14.3 OVER-THE-COUNTER DERIVATIVES

14.3.1 Overview

Given all the advantages of exchange-traded derivatives, why do customers enter into OTC derivatives, which require far more credit scrutiny, are much more difficult to offset, and are surrounded by far more secrecy concerning prices and market conditions? The answer has to be largely centered on the two main weaknesses of exchange-traded derivatives: lack of customization and the operational intensity of managing margin calls. Firms that want to enter into derivative contracts custom-tailored to a specific need must use OTC derivatives.

An additional motivation for using OTC derivatives is that a counterparty may be seeking an extension of credit in connection with its derivatives trading. Initial margin and daily margin calls require cash or securities that the firm may need for other purposes. Unlike an exchange, the provider of an OTC derivative may be willing to extend credit for an amount that is due in the future under a derivative contract.

While some OTC derivatives contracts are negotiated directly between two firms looking for opposite sides of a trade, the overwhelming majority of OTC derivative contracts involve a derivatives market maker as one of the counterparties to the trade. This reflects both the willingness of derivatives market makers to structure contracts that fit the particular needs of a customer and the nature of market making in providing continuous liquidity by being willing to take either side of a trade at a reasonable price, as we discussed in Section 2.5. Finding a non-market-making firm looking for the opposite side of a trade you want to enter into requires an extensive search.

A market maker in derivatives must therefore have both a sophisticated trading operation with regard to market risk and a very high credit rating. In cases where there have been credit concerns regarding a market-making firm, special arrangements have been made to create a subsidiary that has a higher credit rating than the parent firm that will be the counterparty to all derivatives trades (for details, see Gregory 2010, Section 2.3.1 and Chapter 13).

We can therefore see that in many ways the derivatives market maker plays a very similar role to that of the exchange in managing the credit risk of derivatives. Parties taking opposite market positions have credit exposure to a market maker rather than to one another. But the market maker has more freedom than an exchange in deciding how it wants to manage this credit exposure; the loss mutualization rules of the exchange make it answerable to all of its member firms and constrain its options.

Three primary approaches have been proposed and used for managing counterparty credit risk for OTC derivatives. The earliest approach was to treat the counterparty credit risk on OTC derivatives as much as possible like the traditional credit process for loans. We will discuss this approach in Section 14.3.2. The second approach is to incorporate some of the credit management tools of exchange-traded derivatives to OTC derivatives—closeout, netting, and margining. This approach will be discussed in Sections 14.3.3 and 14.3.4. In 14.3.5, we will discuss the most recent of the approaches, the use of dynamic hedging to manage counterparty credit risk.

14.3.2 The Loan-Equivalent Approach

The earliest approach to the management of counterparty credit risk on OTC derivatives was to incorporate it into the traditional credit process for loans. Since credit risk managers are used to making decisions on the total amount of credit that it is prudent to extend to a given borrower, it is only necessary to calculate the total *loan-equivalent* size of credit extension needed for a given OTC derivative position. The difficulty with this approach is that where a standard loan (other than a line of credit) has a fixed amount that is subject to loss in the event of default, the size of derivative exposure at the time of default depends on the uncertain evolution of market conditions.

The standard solution to this problem has been to set some probability threshold (such as the 99th percentile) and then estimate the near-maximum amount that can be lost in the event of default at this threshold. This near-maximum loss amount is treated as a loan equivalent, and credit risk managers are asked to give approval for this added credit extension to the borrower.

Before discussing the computational aspects of this approach, let us note two major issues:

1. Credit risk management looks not just at total credit exposure but also at the expected recovery in the event of default. While historical experience has been developed for recovery on different classes of credit exposure (see Table 13.4), the relative rarity of default by OTC derivative counterparties has made comparable data difficult to obtain. Some assumption about this recovery rate needs to be made based on some combination of relevant experience and theoretical considerations.
2. Derivatives marketers and traders feel discriminated against by this traditional approach. They point out, with reason, that the actual amount at risk in the event of default would, on statistical grounds, often be less than the near-maximum amount used as a loan equivalent, whereas a

traditional loan will always have the same fixed exposure. Derivatives marketers and traders want to see notions of expected exposure at default supplement or replace the measure of near-maximum exposure at default. However, care must be taken to create a comparable measure to traditional loans. If traditional loan exposure is measured by loan amount, the expected exposure on derivatives must be measured by exposure at default and not be based on expected loss, which differs from expected exposure by the amount of expected recovery in the event of default.

With respect to the second point, there is near-universal agreement that expected exposure at default should be measured and that loan officers in making decisions on credit extensions should look at expected exposure along with near-maximum exposure. There is also near-universal agreement that pricing credit exposure and allocating capital against credit use, as in Section 13.3.4, should be based on expected exposure. More controversial is the proposal by some derivatives marketers and traders that near-maximum exposure should not be considered at all and that only expected exposure should be looked at as a measure of credit risk on OTC derivatives. In my experience, this argument has not gained much traction. Certainly for borrowers with very large exposures, the potential impact of default on the lending firm makes it mandatory for credit officers to consider the near-maximum impact. Even for smaller borrowers, the discipline of looking at near-maximum exposure is a healthy incentive to focus loan officers on the soundness of credit extension decisions.

Turning to the computational aspects of the loan-equivalent approach, there are two basic methodologies to consider: simulation and formulas. Consistent with the basic themes of this book, I advocate the use of simulation as the recommended approach (compare with Sections 1.3 and 6.1.1). Simulation is more accurate than formulas in the calculation of credit exposure on a single derivative, is an absolute necessity for looking at credit exposure of the full set of derivatives for a counterparty or for pricing credit exposure for a portfolio of counterparties, is needed for taking into account correlation between market movements and default probability (so-called wrong-way risk), and is an absolute necessity for taking into account credit mitigation techniques such as netting and margining. We will postpone the discussion of the details of simulation methodology until after the introduction of credit mitigation in Section 14.3.3, allowing for a unified simulation approach.

For all these reasons, calculation of credit exposure through formulas has limited applicability and is relied on only by smaller, less sophisticated firms. However, larger and more sophisticated firms may still utilize

formulas as a quick first approximation to guide initial discussions between derivative traders and loan officers and as an aid to intuition. These approximations are usually based on the reasonable assumption that uncertainty about market variables will grow with the square root of elapsed time, balanced by the decrease in duration of products such as interest rate swaps. For a swap, increasing uncertainty at first dominates, and credit exposure increases, reaches a peak, and then declines through time as the impact of decreasing duration comes to dominate. Gregory (2010, Section 4.2 and Appendix 4A) contains examples of approximation formulas and graphs illustrating typical cases.

For less sophisticated firms attempting to approximate counterparty credit exposure without the use of a full simulation model, portfolio credit risk as calculated in Section 13.3.2 will just have expected loan equivalents representing counterparty exposure as input. Portfolio credit risk computed with this shortcut must be adjusted upward to take into account interactions between credit exposure and market value that would be picked up in a full simulation. This is the so-called *alpha factor* explained in detail in Gregory (2010, Sections 10.4 and 10.5). This exposure increase is present even in the absence of any correlation between default probabilities and market values, simply due to the added volatility of market values contributing to higher tail risk of the credit portfolio.

14.3.3 The Collateralization Approach

The second approach to managing counterparty credit risk on OTC derivatives has been to combine the first approach just described with tools borrowed from the exchanges' management of counterparty credit risk. In particular, a combination of netting and closeout is used to combine derivative positions with a single counterparty, and margining is used to obtain collateral that will offset loss in the event of default. Let's look at these two tools in some more detail.

Netting and closeout are discussed at length in Gregory (2010, Sections 3.4 and 3.5). According to Gregory, "Of all risk mitigation methods, netting has had the greatest impact on the structure of the derivatives markets. Without netting, the current size and liquidity in the derivatives markets would be unlikely to exist. . . . The expansion and greater concentration of derivatives has increased the extent of netting from around 50% in the mid-1990s to close to 100% today." Netting and closeout require a legal agreement between counterparties, most typically under an ISDA Master Agreement (see Gregory 2010, Section 3.4.6), that permits, in the event of a counterparty default, the nondefaulting counterparty to immediately terminate all outstanding derivative contracts between the two counterparties,

determine what is owed on each terminated contract at current market values, and net offsetting amounts owed. It eliminates the possibility of the defaulting counterparty settling contracts on which it owes money at only a recovery fraction of the amount owed, while demanding full payment on contracts on which it is owed money. According to Gregory, “ISDA has obtained legal agreements supporting their Master Agreements in most relevant jurisdictions” (wherever there are doubts about legal enforceability of closeout netting in a jurisdiction, ISDA lobbies for legislative clarity; once clarity has been achieved, ISDA obtains a legal opinion to this effect for the benefit of its members).

Another major advantage of the ISDA Master Agreement is that it has standardized procedures for determining what claims can be made in bankruptcy against a defaulting counterparty. The suggested ISDA language defines the amount that can be claimed as the amount that the nondefaulting party “reasonably determines in good faith to be its total losses and costs” as of the closeout date and states that the nondefaulting party “may (but need not) determine its loss by reference to quotations of relevant rates or prices from one or more leading dealers in the relevant markets.” This language makes clear that the nondefaulting party does not have to enter into a replacement transaction in haste in order to establish a price on which to base its claim in bankruptcy proceedings. Instead, it can utilize market quotations, supplemented by industry-standard models, to establish what the mark-to-market of the transaction was at the time of default, base its bankruptcy claim on that, and exercise its best judgment as to when or whether to actually enter into a replacement transaction.

Margining is discussed at length in Gregory (2010, Sections 3.6, 3.7, and 5.2.1). It works similarly to margining by exchanges, with a call for posting of margin to cover a mark-to-market loss and the failure to post margin constituting a default event that will terminate the trade (and all other trades linked through netting agreements). If OTC derivatives margining worked exactly like exchange margining, it would completely eliminate the advantages of OTC derivatives over exchange-traded derivatives in operational simplicity and credit extension (though still leaving contract customization as an advantage). To retain these advantages, OTC derivatives market makers usually make their margining requirements less burdensome than exchange margining requirements by one or more of the following conditions:

- Margin payments may not be required as often as daily, but may have a less frequent period, such as weekly or monthly.
- Margin payments may be required only once a certain mark-to-market loss threshold has been reached.

- Margin may be allowed to be posted as securities of a specified quality rather than necessarily being cash, though this provision has been losing popularity since events of the 2008 crisis (Gregory 2010, Section 3.6.5).
- Initial margin may not be required.
- More leniency may be permitted in allowing a grace period during which the counterparty has time in which to post margin.

These more lenient margining requirements allow OTC derivatives market makers to accept a greater degree of credit exposure to customers than is normally extended by exchanges.

With this background on netting, closeout, and margining, let's begin to look at the computation of counterparty credit risk exposure by simulation.

There are very strong parallels to the use of simulation and stress testing that can be found in Chapter 7, and much of that material is fully applicable to counterparty credit exposure. As in Chapter 7, we are concerned with the value at which a transaction will actually take place—the replacement value at which a derivative contract can be entered into in the event of default for counterparty credit exposure versus the exit value of an existing transaction in the event of forced liquidation for market risk.

The primary differences between the market risk simulation of Chapter 7 and the simulation of counterparty credit exposure are length of simulation period and the required statistics. Counterparty credit exposure must be calculated over much longer time periods than VaR, since a firm can exit its market exposures over a period of a few days but has a longer contractual commitment to the credit risk on derivatives. While market risk simulations are concerned only with tail risk, counterparty credit exposure simulations need to calculate expected value as well as the tails, as already explained in Section 14.3.2.

For the time being we will assume that the timing of default of the counterparty is independent of the market values of the derivative contracts. We will later drop this assumption in the next section, on wrong-way risk.

Here are some points that must be considered in designing counterparty credit exposure simulations in addition to the points already covered in Chapter 11; for a more detailed description, see Gregory (2010, Chapters 4 and 5), and also compare with Brindle (2000) and Canabarro and Duffie (2003).

- The longer time period that counterparty credit exposure simulation requires necessitates the use of Monte Carlo simulation. With VaR simulation, we can choose between historical simulation and Monte Carlo simulation only because the short time period being simulated means

there are many previous historical periods of the same time length as the period to be simulated.

- Each path of the Monte Carlo simulation determines credit exposure at each possible default time being considered. Calculations along each path take into account not just the values of the derivative contracts but also account for all netting that would occur in the event of default and any margin calls and collateral postings that would have occurred based on the details of the margining agreement with the counterparty.
- Since a single counterparty may have entered into many different types of derivative contracts (equity, interest rate, foreign exchange [FX], credit, commodities, etc.), a full range of market variables must be considered, just as in a VaR calculation, with due care exercised on correlation assumptions between variables.
- As with VaR simulations, full valuation of derivatives along each simulation path may be very resource intensive, and trade-offs will exist between the accuracy of full valuation and the faster turnaround time and lower cost of approximations (compare with the discussion of valuation approximations in Section 7.1.1.2). This is an even greater issue for counterparty credit simulations than for VaR simulations, since each path also requires valuations at many different time periods; see the section on “Computational Considerations” in Brindle (2000) and Gregory (2010, Section 4.1.3). To speed computation, in addition to the approximation measures discussed in Section 7.1.1.2, the number of default times for which valuation is done may be reduced with interpolation utilized for default times in between the ones evaluated. Gregory (2010, Section 4.1.4) discusses possible issues with interpolation between the discrete time points for which calculations are made and measures for reducing interpolation error.
- In counterparty credit exposure simulations, the drift (the expected change in a variable through time) plays a more important role than in VaR calculations. Due to the short time frame of VaR calculations, drift can be assumed to be zero, since volatility will totally dominate drift, particularly in tail calculations. But for counterparty credit exposure, over much longer time periods and where expected value is important along with tail values, drift is very important. As Gregory (2010, Section 4.3.2) notes, “in the long run a strong drift will dominate” since volatility varies with the square root of time whereas the drift scales linearly with time. So attention must be paid to forecasting the drifts of market variables in the Monte Carlo simulation.
- In simulating margining, in addition to all contractual details, assumptions need to be made about delays in the time between a margin call being made and a default for failure to meet the margin call being

declared. As Gregory (2010, Section 5.2.1) explains in detail, the industry standard incorporated into Basel II is to assume a 10-business-day minimum *remargin period* between margin call and default declaration and closing out of positions. This allows time for both operational issues of processing margin requests and delays in detection of non-delivery, and grace periods allowed to permit a counterparty to cure a failure to post margin. As Gregory notes, longer remargin periods may be appropriate for counterparties that may be granted more leniency to maintain good relations or where the nature of the derivatives may require longer periods to resolve disputes over the mark-to-market driving a margin call. Brindle (2000) also notes that in some jurisdictions, statutory *stay periods* may delay the liquidation of collateral, and contractual agreements may stipulate a minimum delay period. The closeout delays assumed in the simulation should be individually tailored to each counterparty.

- When a counterparty agreement allows for noncash collateral, the market value of the collateral should also be simulated along each of the simulation paths, with full consideration of correlation between value of the collateral and value of the derivatives. When the value of the derivatives and of the collateral instrument are positively correlated (e.g., a Treasury bond as collateral and a set of swaps on which the counterparty net pays a fixed rate in the same currency), credit exposure will be greater than if collateral was posted in cash. When the value of the derivatives and of the collateral instrument are negatively correlated (e.g., a Treasury bond as collateral and a set of swaps on which the counterparty net receives a fixed rate in the same currency), credit exposure will be less than if collateral was posted in cash. A worked example of the impact of collateral on credit exposure can be found in Gregory (2010, Section 5.2.5).
- As with VaR, as discussed in Section 7.1.1.2, counterparty credit exposure simulations must account for illiquidity, whether due to infrequent trading or to a large position. Illiquidity must be considered for both the derivative positions and for noncash collateral. Whether due to infrequent trading or to large positions, the basic tool for dealing with illiquidity of derivatives is to lengthen the time assumed between a default event and position closeout. This closely parallels the treatment for illiquidity detailed in Section 6.1.4 and the provision for remargin periods discussed two bullet points previously. Illiquidity will probably have limited impact on counterparty exposure where margining is not used—there is little difference between the price movement in derivative value over, say, two years from now to default and over two years and two weeks, allowing an extra two weeks after default for illiquidity.

But illiquidity can have a major impact on counterparty exposure when margining is used. It could now, for example, double the time from default event to closeout from two weeks to four weeks, increasing exposure by $\sqrt{2}$. Similarly, illiquidity of collateral can be treated by increasing the time period over which the collateral is assumed to be liquidated and hence the uncertainty of the price realized. When illiquidity of a derivative is due to a position with actuarial risk, a separate treatment is needed. This will be discussed next.

- For derivatives with actuarial risk, I strongly favor an approach parallel to that recommended in Section 6.1.2: utilize a liquid proxy in the counterparty exposure simulation but make a separate computation for the residual risk. My argument is that a default by the counterparty will result in the nondefaulting party acquiring and now needing to manage the actuarial risk in the same way it would have needed to manage it if it had created it in a trading position. The reserves that would be needed to manage out of the position, as computed in Section 8.4, will now be a potential cost and hence are an addition to near-maximum credit exposure and need to be accounted for in expected credit cost, multiplied by the proper default probability and loss given default percentage. Another consequence of this argument is that firms should not enter into derivatives positions on transactions for which they lack adequate models and personnel to manage a position that will result from a default. There have been unfortunate examples in which firms have decided to “stand in the middle” between two counterparties on a transaction that they had no experience trading and little understanding of, persuaded that they were “only” taking a counterparty credit risk and not taking any market risk (typically because one of the counterparties was not willing to accept the credit risk of the other and was looking for a counterparty with stronger credit risk). On default of one of the counterparties, these firms found themselves suddenly needing to manage positions they lacked competence to trade.
- Stress testing as a supplement to simulation of counterparty credit exposure plays a smaller role than it does as a supplement to VaR, for reasons similar to those discussed two bullet points previously concerning the minimal impact of illiquidity of positions on counterparty credit exposure. By parallel reasoning, a stress scenario of a temporary period of market illiquidity in normally liquid positions will have little impact on exposure when no margining is employed but may be quite necessary and of significant impact when margining is employed.
- The degree to which netting reduces near-maximum credit exposure is very heavily impacted by correlation assumptions regarding market variables. It is important to make sure that subjective probabilities of

future periods in which correlations that are either very low or very high by historical standards have been given due consideration.

The need for communication of marginal cost of new credit exposures to loan officers discussed in Section 13.3.4 has a parallel requirement for communicating the marginal cost of new counterparty credit exposures to loan officers, traders, and structurers. This is done through the *credit value adjustment* (CVA), a thorough discussion of which can be found in Gregory (2010, Chapter 7). Gregory's discussion of measuring marginal exposure contributions in his Section 4.5 is also relevant. I will limit myself to just a few remarks related to cases where the CVA methodology differs in some respect from the methodology of Section 13.3.4:

- As in the more general case of marginal credit exposures discussed in Section 13.3.4, there is a need for approximations that can be used at the individual credit level. Gregory provides approximation formulas in the appendixes to Chapter 7, but with the important caveat that these work only in the absence of wrong-way risk (i.e., when there is no dependence between default probability and loss given default). When wrong-way risk is present, the techniques of the next section, 14.3.4, need to be used; these are very closely related to the computations in Section 13.3.4 and so would need to utilize approximation techniques covered there, though Gregory's Section 8.3 does provide some approximation formulas specific to CVA for wrong-way risk.
- Many firms have employed an accounting procedure that takes into account the impact on derivative contracts of the default probability of the firm itself (this is termed *bilateral counterparty risk* and is covered by Gregory in Section 7.3). Whatever its virtues as an accounting procedure, it should *never* be utilized in risk management measures such as CVA. From a risk management standpoint, there are no benefits to a firm from its own default, so utilizing it in risk measures would be completely misleading. Even as an accounting procedure, the benefits of this approach are dubious: an attempt to book profits that will fuel short-term bonuses at the potential expense of investor confidence in the firm's reported earnings, as can be seen in the examples in Gregory (2010, 188).

14.3.4 The Collateralization Approach—Wrong-Way Risk

In the previous section on simulation of counterparty credit exposure, we noted that a key assumption in our calculations was independence of counterparty default and market value of the derivatives contracts. For many

counterparties, this is a reasonable assumption. When there is a correlation between default and market value, then computations must be different. A positive correlation between probability of default and market value of the derivatives is known as *wrong-way risk* and increases exposure and CVA measures from what they would have been in the absence of this correlation. A negative correlation between probability of default and the market value of derivatives is known as *right-way risk* and decreases exposures and CVA measures from what they would have been in the absence of this correlation. Gregory (2010, Chapter 8) contains a thorough exposition of wrong-way and right-way risk.

This section addresses how to modify the simulation methodology of the previous section to accommodate this correlation. The short answer is that the default probability of the counterparty must also be simulated along each path, incorporating correlation with the market variables being simulated. We will provide details and examples shortly, but first let us consider some cases in which wrong-way risk is so extreme that simulation should be circumvented and a direct analysis should be made.

Let's start with a trade that has, unfortunately, been proposed all too often by trading desks over the past 15 years. With macabre humor, it is sometimes called an "end of the world" trade. It is a proposal to put on a derivative trade that will provide a payoff only if some really extreme event occurs—let's say 40 percent defaults on a basket of investment-grade corporate loans. No initial margin is being asked of the counterparty providing the protection, and there is no provision for margin calls.

It is easy to see the attractiveness of this trade from the viewpoint of the counterparty providing the protection; it will receive a small annual payment every year, and if the dire circumstances in which it is required to make a payment did occur, it doubts it would still be in business.

It is harder to see why the firm purchasing the protection would want to do the trade. In every case I have encountered, when I asked the trading desk proposing the trade whether they thought there was any chance the counterparty would still be in business if it was required to make a payment, the answer was, "No, but even though this has no financial benefit to the firm, it will provide us relief under such-and-such regulatory capital calculation." My response, as a risk manager, was always: (1) we wouldn't permit trades to be made that cost the firm money with no financial benefit, and (2) even if it appeared to provide regulatory relief, it would be my obligation as someone in a control function to point out to the regulatory authority concerned that it was being gamed. In no way was any modeling required to come to this conclusion.

A less obvious case is one in which no margining is required by a counterparty unless the counterparty receives a ratings downgrade below

a certain level or unless an extremely negative event occurs in the counterparty's stock price or credit spread, in which case a large margin payment is required. In such circumstances, I have always been opposed to giving any credit in counterparty credit exposure calculations for this margining requirement; I would make the calculations assuming no margin requirement at all. My reasoning is that the type of event that triggers the margin call is just the sort of circumstance in which the counterparty will be strapped for cash and will either be forced to default or will appeal to our firm's senior management for relief from the margin call to avoid bankruptcy. Indeed, it was just this type of margining provision that pushed Enron into bankruptcy (see McLean and Elkind 2003, 394–395). So this is a case of wrong-way risk in which there is a high correlation between a required margin payment and a default that prevents it being made.

Another variant on extreme wrong-way risk is an attempt to avoid reliance on margin calls that have a low probability of being fulfilled by converting the counterparty credit risk into a gap market risk. A detailed and instructive worked example of this mechanism is given in Gregory (2010, Section 8.6.4). I will build on Gregory's example in the discussion that follows, but with only a brief sketch of Gregory's details.

In the example, the market-making firm buys or issues a \$100 million credit-linked note (CLN) and enters into a total return swap on the CLN with a hedge fund. The hedge fund posts \$10 million in initial margin and benefits from having a highly leveraged position, receiving a return on the \$100 million note while only needing to invest \$10 million in collateral. The downside for the market maker is that it knows that if the market value of the CLN starts to decline toward \$90 million, it is highly unlikely that the hedge fund will be able to post additional margin, since the hedge fund, under the circumstances that credit spreads have risen high enough to create this size market loss on the CLN, will likely be in trouble due to its high leverage and probable losses on similar trades.

The market maker's trading desk knows it is unlikely to get any credit for margin call provisions due to the extreme wrong-way risk. A possible alternative is to exclude the margin call provision but instead put in a provision that if the value of the CLN gets too close to exhausting the \$10 million initial margin, the market maker has the right to close out the position and sell the CLN. In Gregory's example, a provision is set that if the value of the CLN is at or below \$92.2 million, the position can be closed out. This is supposed to leave the evaluation of the trade entirely to market risk managers since there is no credit risk component remaining. The only losses to the market maker can occur if the gap between the \$92.2 million trigger point and the price at which the CLN can be sold exceeds the \$2.2 million of remaining initial margin. It is the probability of this large market move

occurring that is supposed to be evaluated by standard market risk VaR and stress test methodologies.

I have always been dubious of this type of attempted end run. I think it just replaces one form of wrong-way risk with another form of wrong-way risk: the high correlation between large drops in price of the CLN and large subsequent gap moves. The fundamental flaw in the appeal to VaR and stress test methodologies in evaluating the gap risk is that VaR and stress testing are designed to evaluate the risk of current positions based on current market conditions. For gap risk, we are being asked to evaluate a future position under future market conditions and one that will be triggered by conditions likely to be unfavorable to us. As such, they fall under one of the criteria proposed for actuarial risk in Section 6.1.1, positions that can be liquidated only under restrictive conditions. Hence, they should be evaluated using the tools of Section 8.4, with very conservative reserves to allow for the illiquidity of the position. Subjective judgment by risk management would be required as to the size of gap moves that could occur following the very negative market events that would cause the trigger to be reached.

One more variant of extreme wrong-way risk is the liquidity puts described in Section 5.2.5.2. Here an investment bank was selling an extremely illiquid asset, a super-senior tranche of a CDO, but with the provision that if the firm buying this asset encountered funding difficulties it could sell the asset back to the investment bank at par. This type of transaction should be treated for stress testing purposes as if the asset had not been sold at all—the firm buying the asset would probably run into funding difficulties only in a period of widespread financial distress, exactly the circumstances in which the asset is likely to be worth significantly less and be even harder to find another buyer for. Since the assessment of the potential losses on the asset were that it would lose significant value only in a period of unlikely widespread financial distress, allowing it to be placed back to the investment bank in these circumstances reduces the risk reduction for stress testing purposes of selling the asset to a negligible amount.

We now turn to the details of simulation incorporating correlation between default probabilities and market variables for those instances of wrong-way and right-way risk that do require full calculation.

- Instead of assuming that default occurs independent of market variables, we now directly simulate default probabilities and allow the Monte Carlo simulation to work from these default probabilities to assign defaults to particular paths and time periods. Expected and near-maximum exposure values are computed from only those points at which default has occurred. If those default points are correlated with

market value of the derivatives positions, this will be reflected in the simulation results.

- Correlations between default probabilities and market values will need to be established by a combination of subjective judgments based on economic insight and statistical studies of correlations between market variables and credit spreads as a proxy for default probabilities.
- Much depends on the degree of business diversification of a counterparty. A counterparty with many business lines in different countries and different industries is far less likely to be subject to wrong-way risk than a counterparty with a highly concentrated business.
- One of the most obvious examples of wrong-way risk stems from country risk. A counterparty whose financial health is very dependent on business in a single country is likely to have a high correlation between its default probability and the exchange rate and interest rates of that country. This most frequently impacts long-term FX forwards or cross-currency swaps. As pointed out by Gregory (2010, Section 8.2.3), “another way to look at a cross-currency swap is that it represents a loan collateralized by the opposite currency in the swap. If this currency weakens dramatically, the value of the collateral is strongly diminished.”
- A business whose viability is likely to be strongly impacted by the price of a particular commodity such as oil should show a strong correlation between default probability and the commodity price.
- Correlations between default probabilities of firms based on industry and country have already been discussed in Section 13.3.1. This can have a strong impact if a counterparty is highly correlated with a firm on which it is writing credit protection through a credit default swap (CDS). One of the principal sources of wrong-way risk historically has been the use of CDS counterparties closely related to the firm on which protection is being purchased. The credit portfolio simulations of Section 13.3.2 should be able to capture this. Consider, for example, a loan to Company ABC for which credit protection has been purchased from Company XYZ through a CDS. No loss will occur if ABC defaults and XYZ has not defaulted, since, in this circumstance, XYZ must pay all the costs of the ABC default. If XYZ defaults and ABC has not defaulted, the firm will have a loss (or gain) equal to the replacement cost of the CDS, which is driven by changes in the credit spread for ABC. The simulation calculates this by keeping track of changes in default probabilities and credit spreads for both firms along each simulation path, taking the proper correlation between the default probabilities of the two firms into account, and linking the default probability of ABC to the credit spread of ABC. Gregory (2010, Section 8.4) provides

more detail and examples illustrating wrong-way risk on CDSs, and in Section 8.5 extends this analysis to wrong-way risk on CDOs.

- A significant source of wrong-way risk is counterparties who derive a major portion of their revenues from financial transactions. In such cases, an estimate must be made of how much of the counterparty's trading positions are similar to those on which your firm holds positions with the counterparty. The more similar overall trading positions are to those with your firm, the more likely that default probability has a high correlation with market variables impacting those positions.

While simulation is a requirement for accuracy in measuring wrong-way risk, formulas can be utilized for quick approximations that are useful in gaining intuition and to guide initial discussions between derivative traders and loan officers. Examples of useful formulas and illustrated cases can be found in Gregory (2010, Section 8.3) and Winters (1999).

14.3.5 The Active Management Approach

The third, and newest, approach to managing counterparty credit risk for OTC derivatives involves the active use of purchased credit protection through CDSs (or, equivalently, by short selling of bonds). As such, it shares many of the characteristics of active management of credit portfolios discussed in Section 13.3.4, involving trade-off decisions about when to purchase protection versus when to self-insure by extending credit lines, the communication of internal pricing of new credit extensions based on a combination of the cost to purchase CDS protection and the cost of required capital against self-insurance risk, the active involvement of marketers and traders in making judgments about whether the extension of new credit is worth paying the internal charge, and the management by a central unit of the cost of loan defaults against the revenue accumulated by internal charges for credit extension. The difference between the active management of counterparty credit risk and of portfolio credit risk is that counterparty credit risk active management involves simultaneous management of the cost of credit exposure and the dynamic changes in size of credit exposure due to changes in the market value of counterparty positions. This requires a very specialized skill set that has led most large derivatives dealers to set up specialized business units (counterparty risk groups [CRGs]) for the dynamic management of counterparty credit risk. Gregory (2010, Chapter 12) gives an extended discussion of how this is done.

The centralized unit for managing counterparty exposure will need to create a mechanism for charging trading desks for protection against counterparty risk. This mechanism must follow many of the same criteria as

outlined in Section 13.3.4 in the context of the more general issue of how to charge marketing areas for the extension of credit risk, but with the added complexities of estimating the credit exposure arising from market movements. These charges should create the incentives for trading desks and derivatives structurers to design contracts that minimize credit use. There will be trade-offs between customer desire to minimize the use of devices such as margin calls and the reduction in credit charges that result from such devices. It is the task of traders and structurers to find clever designs that bring the greatest reduction in credit charge for the least amount of customer dissatisfaction.

To the extent this counterparty risk group (CRG) decides to manage counterparty credit risk with the purchase of CDS protection, it requires the use of dynamic hedging techniques originally developed for multiasset exotic derivatives such as quantos. The size of market exposure at any instant is the product of the credit spread of the counterparty and the size of the credit exposure. As we illustrated in Section 12.4.5, this requires dynamic hedging, with a change in derivative value requiring a change in the size of the credit hedge, and a change in the credit spread requiring a change in the size of the derivative hedge. Essentially, this method amounts to replacing the derivative with another counterparty, not all at once on default, but gradually as the original counterparty's credit worsens. Correlation assumptions, driven by wrong-way exposure concerns, will have the intuitively correct effect of increasing the expected cost of the dynamic hedge. The **CrossHedge** spreadsheet gives a detailed example of the dynamic hedging of a counterparty credit position with results shown in Table 12.13.

What the example in Section 12.4.5 illustrates is that, to a good degree of accuracy, the dynamic hedge allows locking in credit protection on the counterparty at the current market credit spread, even though the amount of credit protection will vary over time in a stochastic fashion. This is quite counterintuitive—it would seem that if credit spreads widened at the time that exposure grows you would need to purchase some of the credit protection at higher spreads. But the dynamic hedging approach means that you are always simultaneously hedged against both changes in credit spread and changes in exposure (always with the exception that correlation in price movements between the credit spread and the market exposure caused by wrong-way exposure leads to extra costs). This allows the CRG to be able to price credit exposure at the time of agreeing to the derivatives contract with reasonable confidence. While the example in Section 12.4.5 is written from the point of view of credit protection on a single derivatives contract, the mechanism actually works for covering an entire portfolio of derivatives—essentially, you just substitute the volatility of the whole portfolio for the volatility of the single contract.

In practice, a CRG will choose to use CDS hedging on some exposures and not on others—some counterparties will not have sufficient liquidity in the CDS market to allow the dynamic hedging technique to be used; for other counterparties the credit managers will judge that their view of the credit risk of the name is more favorable than what is priced into the CDS market and they will choose to self-insure for that name, at least for a time. In other cases, mixed approaches will be taken—names that lack a liquid CDS market but whose exposure is at an uncomfortable level for the credit managers will be proxy hedged with a basket of more liquid CDSs on similar names being used to hedge a basket of less liquid names, with the risk having been transformed from outright default risk to the basis risk on default experience of the basket hedge and default experience of the actual basket. The simultaneous dynamic hedging of credit spread (for the proxy basket) and market exposure works in this case as well.

When utilizing dynamic hedging of counterparty credit exposure, a CRG will need to utilize risk measures similar to those we have discussed for dynamic hedging of options in Section 11.4, but with the added complications that exposures to credit and to market variables are being managed simultaneously and that credit risk requires risk measures that include exposure to immediate default. A thorough discussion of the risk measures required can be found in Gregory (2010, Chapter 9).

One issue for CRGs that has been much debated and is highlighted by Gregory (2010) in Section 12.4.4 is whether the CRG should engage in dynamic hedging of the market exposure of a derivatives book in a case where it is completely self-insuring the credit risk for that counterparty. Unlike the dynamic cross-hedging examples just given, there is no cost of a CDS position that is being offset by the market exposure hedge. All that is being hedged is an accounting entry of the mark-to-market of the self-insurance strategy. The economic value of paying money to hedge accounting entries is regarded with extreme suspicion by many risk management practitioners, myself included. But if there is some form of active hedging in the CDS market, even if it is only against a basket of names that provide a liquid proxy, then I would find dynamic hedging of market exposure to be quite reasonable.

In taking over management of the counterparty credit risk of derivatives, the CRG must be prepared to manage all aspects of a counterparty default (see Gregory 2010, Section 12.2.6). This includes the settlement process on any CDS protection that has been purchased (which may involve delivery squeezes, as discussed in Section 13.1.1.2), the legal process for recovery of amounts owed, and responsibility for the liquidity costs of replacing defaulted contracts. The CRG must factor all of these possible costs into its pricing of default insurance to the firm's trading desks.

There are other strategies that a CRG can pursue in providing protection. It might, for example, contact a counterparty with which the firm has a large outstanding exposure and seek to negotiate a reduction in exposure. This could be especially attractive if deterioration in this counterparty's credit outlook causes particular concern to the firm's credit risk managers. Reduction in exposure could come in several different forms: a one-time posting of margin or renegotiating the terms of existing contracts to provide for tighter terms on posting of margin. Of course, posting margin or tightening margin requirements is costly to the counterparty, so some concession must be offered as inducement—probably as a renegotiation of the financial terms of the derivative contracts to make strikes or spreads more favorable to the counterparty. The CRG would need to compensate the relevant trading desk for any such pricing concessions and must judge whether this up-front cost is worth the reduction in credit risk.

Another strategy that a CRG could pursue in reducing exposure to a counterparty is to offer the counterparty a mutual reduction in exposure—reducing the counterparty's credit exposure to the firm by changing the financial terms on some derivative contracts on which the firm owes money to the counterparty in exchange for reducing the firm's credit exposure to the counterparty by changing the financial terms on some derivative contracts on which the counterparty owes money to the firm. These changes in financial terms can be done in such a way as to leave the net amount owed by one party to the other unchanged, but with lower gross amounts owed. While netting and closeout master agreements accomplish much the same thing, actual reduction in gross amounts owed reduces the amounts that will be in contention in litigation that follows a default, and thus offers positive benefits.

A greater impact on exposures could be achieved by moving beyond bilateral negotiations for changed financial terms to multilateral negotiations in which a counterparty's exposure to one firm is reduced in exchange for a reduction in another firm's exposure to the counterparty. This results in actual reduction in credit exposure, not just the reduction of litigation risk of the bilateral negotiation of changed financial terms discussed previously. Here's a simple illustration. Suppose Bank A currently is owed \$50 million on an interest rate swap by Counterparty B, and Bank C currently owes Counterparty B \$50 million on an FX forward. If Counterparty B is willing to renegotiate the financial terms on these two contracts, it would not have to make any payments, since the \$50 million it would owe to Bank A for the renegotiation would be offset by the \$50 million it is owed by Bank C. Bank C would owe a \$50 million payment to Bank A, but Bank A would offer Bank C some discount on this as an inducement to lowering Bank A's credit exposure to Counterparty B and to compensate Bank C for losing the

cushion it had against having a credit exposure to B. In summation, Bank A benefits from reduced credit exposure but may have to pay something for it, Counterparty B is not impacted and in fact might gain slightly by reduced credit exposure to Bank C (though it may ask for some payment from Bank A for its cooperation), and Bank C will benefit to the extent it receives a payment from Bank A. Other creditors of Counterparty B are potentially disadvantaged, since in a default they would no longer have a claim on the amount owed to B by Bank A, but they have no standing in the transaction as long as B is a going concern.

Variants of this last transaction have been introduced as a way for derivatives market makers to lower credit usage on derivatives transactions between market makers, and thereby free up credit lines. For example, several market makers get together and engage in *trade compression*, in which the market makers identify a set of derivative transactions that can be canceled and replaced by another set of derivative transactions, leaving market exposures close to unchanged but with a significant decrease in credit exposures. In addition to canceling trades that offset one another in market exposure, slight differences in contract detail that have little impact on market exposure can be eliminated to increase possibilities for contract cancellation. Some vendors now offer analytical services for developing proposed replacements that optimize the reduction in credit exposure that can be achieved by trade compression. Vause (2010) has a thorough discussion of trade compression and similar counterparty credit reduction techniques with examples. ISDA (2012) provides a detailed exposition of compression in the important case of interest rate swaps and illustrates the trade-off between a firm's tolerance for small changes in interest rate exposure and the degree of compression that can be accomplished.

Generally speaking, having a derivatives position with a counterparty that is marked to market in your favor gives rise to credit exposure, but there is no offsetting credit benefit from having derivatives positions with a counterparty that is marked to market against you. Many CRGs have been searching for ways to achieve a more symmetrical position. We have just seen (in the next-to-previous paragraph) an example in which a firm can benefit from the credit consequences of a mark-to-market against it, since Bank C would be paid by Bank A to use its negative exposure to offset A's positive exposure to Counterparty B. But this captures only part of the value of the exposure. A strategy that has been proposed for capturing the full value of the exposure is to purchase a bond of the counterparty that you net owe money to on derivative contracts with a maturity close to that of your derivative positions. Let's consider an example to see how this might work.

Let's say you net owe \$100 million in derivatives marked to market to a counterparty in a weak credit condition. Say you can purchase \$100 million

face value of its bonds for \$90 million owing to its poor credit outlook. If the counterparty does not default, then you gain \$10 million from the bond that you purchased at \$90 million maturing at \$100 million. If it does default, you can use the bond you own as an offset in bankruptcy proceedings to the \$900 you owe the counterparty on the derivatives. So you have been able to use the amount you owe on your derivatives contracts to purchase free default protection on the bonds. The CRG would, of course, need to dynamically manage the amount of bond it holds to match changes in the derivatives market exposure in just the same way it dynamically manages the amount of CDS protection it buys when it is owed money on the derivatives position. The risk of this strategy is that a bankruptcy court could possibly object to offsetting the derivatives position and the bond holding.

Finally, one option for a CRG would be to just purchase complete protection against the counterparty credit risk on a particular derivatives trade through a *contingent credit default swap* (CCDS). This is a CDS that in the event of default pays the amount that has been lost on the referenced derivatives trade. So, in effect, the CRG is turning the management of credit risk on this trade over to another firm. The firm selling the CCDS will have all of the issues of managing risk on this trade that we have discussed throughout this chapter and will need to be paid accordingly. There are many negatives arguing against the use of a CCDS, such as the mismatch between the amount of protection purchased and the amount of protection actually needed, since buying protection on a single transaction cannot take reduction in exposure through netting and margining into account. The CCDS is therefore probably a solution for only very large single transactions that are unlikely to have much offset against them. A thorough discussion of CCDSs can be found in Gregory (2010, Section 9.8.2).

References

A copy of this references list will be maintained on the website for this book, with ongoing updates to the Internet sites referenced.

- Acharya, Viral, Matthew Richardson, Stijn van Nieuwerburgh, and Lawrence White. 2011. *Guaranteed to Fail*. Princeton, NJ: Princeton University Press.
- Agrawal, Deepak, Navneet Arora, and Jeffrey Bohn. 2004. "Parsimony in Practice, and EDF-Based Model of Credit Spreads." Moody's KMV. www.moodysanalytics.com/~media/Insight/Quantitative-Research/Credit-Valuation/04-29-04-Parsimony-in-Practice.ashx.
- Allen, Peter, Stephen Einchcomb, and Nicholas Granger. 2006. "Variance Swaps." *J.P. Morgan Securities European Equity Derivatives Research Investment Strategies* no. 28.
- Allen, Steven, and Otello Padovani. 2002. "Risk Management Using Quasi-Static Hedging." *Economic Notes* 31 (2): 277–336.
- Almgren, Robert, and Neil Chriss. 2001. "Optimal Execution of Portfolio Transactions." *Journal of Risk* 3 (2): 5–39.
- Altman, Edward, Neil Fargher, and Egon Kalotay. 2010. "A Simple Empirical Model of Equity-Implied Probabilities of Default." <http://pages.stern.nyu.edu/~ealtman/DefaultModel2010JFI.pdf>.
- Altman, Edward, and Egon Kalotay. 2010. "A Flexible Approach to Modeling Ultimate Recoveries on Defaulted Loans and Bonds." http://pages.stern.nyu.edu/~ealtman/FlexibleRecovery_v1.1.pdf.
- Altman, Edward, Andrea Resti, and Andrea Sironi. 2001. "Analyzing and Explaining Default Recovery Rates." www.defaultrisk.com/pp_recov_28.htm.
- Amato, Jeffery, and Eli Remolona. 2003. "The Credit Spread Puzzle." *BIS Quarterly Review* (December). www.bis.org/publ/qtrpdf/r_qt0312e.pdf.
- Andersen, Leif, and Jesper Andreasen. 2000. "Static Barriers." *Risk* 9:120–122.
- Andersen, Leif, and Jesper Andreasen. 2001. "Factor Dependence of Bermudan Swaptions: Fact or Fiction?" *Journal of Financial Economics* 62 (1): 3–37.
- Anderson, Leif, Jakob Sidenius, and Susanta Basu. 2003. "All Your Hedges in One Basket." *Risk* 11:67–72. www.risk.net/data/Pay_per_view/risk/technical/2003/1103_credit_der.pdf.
- Araten, Michel, and Michael Jacobs. 2001. "Loan Equivalents for Revolving Credits and Advised Lines." *RMA Journal* 5:34–39. http://michaeljacobsjr.com/Jacobs_Araten_RMA_LEQ_Pub_April_2001.pdf.

- Artzner, Philippe, Freddy Delbaen, Jean-Marc Eber, and David Heath. 1997. "Thinking Coherently." *Risk* 11:68–71. A longer version is available at www.math.ethz.ch/~delbaen/ftp/preprints/CoherentMF.pdf.
- Ashcroft, Adam, and Til Schuermann. 2008. "Understanding the Securitization of Subprime Mortgage Credit." Federal Reserve Bank of New York Staff Reports no. 318. www.newyorkfed.org/research/staff_reports/sr318.pdf.
- Asiaweek*. 1996. "Perils of Profit." July 5. www-cgi.cnn.com/ASIANOW/asiaweek/96/0705/ed2.html.
- Bahar, Reza, and Krishan Nagpal. 2000. "Modeling the Dynamics of Rating Transition." *Credit* 3:57–63.
- Bai, Jennie, and Pierre Collin-Dufresne. 2011. "The CDS-Bond Basis during the Financial Crisis of 2007–2008." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1785756http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1785756.
- Basel Committee on Banking Supervision. 2009a. "Principles for Sound Stress Testing Practices and Supervision." www.bis.org/publ/bcbs155.pdf.
- Basel Committee on Banking Supervision. 2009b. "Supervisory Guidance for Assessing Banks' Financial Instrument Fair Value Practices." www.bis.org/publ/bcbs153.pdf.
- Baxter, Martin, and Andrew Rennie. 1996. *Financial Calculus*. Cambridge: Cambridge University Press.
- Bennett, Oliver. 2001. "Splitting Headaches." *Risk* 7:36–37.
- Bluhm, Christian, Ludger Overbeck, and Christoph Wagner. 2002. *An Introduction to Credit Risk Modeling*. Boca Raton, FL: CRC Press.
- Bodie, Zvi, Alex Kane, and Alan Marcus. 2009. *Investments*. 8th ed. New York: McGraw-Hill.
- Bohn, Jeffrey, and Roger Stein. 2009. *Active Credit Portfolio Management in Practice*. Hoboken, NJ: John Wiley & Sons.
- Bookstaber, Richard. 2007. *A Demon of Our Own Design*. Hoboken, NJ: John Wiley & Sons.
- Bouchet, Michel, Ephraim Clark, and Bertrand Gros Lambert. 2003. *Country Risk Assessment*. Chichester, England: John Wiley & Sons.
- Brealey, Richard, Stewart Myers, and Franklin Allen. 2011. *Principles of Corporate Finance*. 10th ed. New York: McGraw-Hill.
- Breuer, Thomas, and Imre Csiszar. 2010. "Stress Tests: From Arts to Science." https://homepages.fhv.at/tb/cms/?Working_Papers.
- Breuer, Thomas, and Gerald Krenn. 2000. "Identifying Stress Test Scenarios." <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.27.118>.
- Brindle, Andy. 2000. "Calculating with Counterparties." *Risk* 1:100–103. www.risk.net/data/Pay_per_view/risk/technical/2000/risk_0100_jpmorgan.pdf.
- Broadie, Mark, Paul Glasserman, and Gautam Jain. 1997. "Enhanced Monte Carlo Estimates of American Option Prices." *Journal of Derivatives* 5 (1): 25–44. A closely related article can be found at www.axelkind.com/teaching/financial_derivatives/papers/broadie_glasserman_1997_jedc.pdf.
- Broom, Giles. 2011. "UBS in 'Disarray.'" *Bloomberg News*. September 26. www.bloomberg.com/news/2011-09-25/ubs-in-disarray-as-gruebel-quits-ermotti-named-interim-chief.html.

- Brown, Aaron. 2012. *Red-Blooded Risk*. Hoboken, NJ: John Wiley & Sons.
- Bruck, Connie. 1988. *The Predators' Ball*. New York: Penguin.
- Brunnermeier, Markus. 2009. "Deciphering the Liquidity and Credit Crunch 2007–2009." *Journal of Economic Perspectives* 1:77–100. www.princeton.edu/~markus/research/papers/liquidity_credit_crunch.pdf.
- Bucay, Nisso, and Dan Rosen. 2001. "Applying Portfolio Credit Risk Models to Retail Portfolios." *Journal of Risk Finance* 3:35–61. www.financerisks.com/filedati/WP/ALGO_PAPER/ch5_retail.pdf.
- Burghardt, Galen, and Gerald Hanweck. 1993. "Calendar-Adjusted Volatilities." In *Volatility: New Estimation Techniques for Pricing Derivatives*, edited by Robert Jarrow. London: Risk Books.
- Burghardt, Galen, and Susan Kirshner. 1994. "One Good Turn." *Risk* 11:44–54.
- Canabarro, Eduardo, and Darrell Duffie. 2003. "Measuring and Marking Counterparty Risk." In *Asset/Liability Management of Financial Institutions*, edited by Leo Tilman, Chapter 9. London: Euromoney Books. www.darrellduffie.com/uploads/surveys/DuffieCanabarro2004.pdf.
- Carr, Peter. 2001. "Closed Form Barrier Option Valuation with Smiles." www.math.nyu.edu/research/carrp/papers/pdf/closed4.pdf.
- Carr, Peter, and Andrew Chou. 1996. "Breaking Barriers." *Risk* 9:139–145. www.math.nyu.edu/research/carrp/papers/pdf/breakingbarriers.pdf.
- Carr, Peter, and Andrew Chou. 1997. "Hedging Complex Barrier Options." www.math.nyu.edu/research/carrp/papers/pdf/multipl3.pdf.
- Carr, Peter, Katrina Ellis, and Vishal Gupta. 1998. "Static Hedging of Exotic Options." *Journal of Finance* 53 (3): 1165–1190.
- Carr, Peter, and Dilip Madan. 2002. "Towards a Theory of Volatility Trading." www.math.nyu.edu/research/carrp/papers/pdf/twrdsfig.pdf.
- Cass, Dwight. 2000. "The Devil's in the Details." *Risk* 8:26–28.
- Chew, Lillian. 1996. *Managing Derivatives Risks*. New York: John Wiley & Sons.
- Chriss, Neil. 1997. *Black-Scholes and Beyond*. Chicago: Irwin.
- Clementi, Gian Luca, Thomas Cooley, Matthew Richardson, and Ingo Walter. 2009. "Rethinking Compensation in Financial Firms." In *Restoring Financial Stability*, edited by Viral Acharya and Matthew Richardson. Hoboken, NJ: John Wiley & Sons.
- Clewlow, Les, and Chris Strickland. 1998. *Implementing Derivatives Models*. New York: John Wiley & Sons.
- Cochrane, John. 2001. *Asset Pricing*. Princeton, NJ: Princeton University Press.
- Cordell, Larry, Yilin Huang, and Meredith Williams. 2012. "Collateral Damage: Sizing and Assessing the Subprime CDO Crisis." Working Paper No. 11-30/R, Research Department, Federal Reserve Bank of Philadelphia. www.philadelphiafed.org/research-and-data/publications/working-papers/2011/wp11-30.pdf.
- Counterparty Risk Management Policy Group. 1999. "Improving Counterparty Risk Management Practices." www.defaultrisk.com/pp_other_08.htm.
- Coval, Joshua, Jakub Jurek, and Erik Stafford. 2008. "The Economics of Structured Finance." Harvard Business School Working Paper no. 09-060. www.hbs.edu/research/pdf/09-060.pdf.

- CreditGrades. 2002. "CreditGrades Technical Document." RiskMetrics Group. www.msci.com/resources/research/technical_documentation/CGtechdoc.pdf.
- Crosbie, Peter, and Jeff Bohn. 2003. "Modeling Default Risk." Moody's KMV. www.moodysanalytics.com/~media/Insight/Quantitative-Research/Default-and-Recovery/03-18-12-Modeling-Default-Risk.ashx.
- Crouhy, Michel, Dan Galai, and Robert Mark. 2001. *Risk Management*. New York: McGraw-Hill.
- Culp, Christopher. 2001. *The Risk Management Process*. New York: John Wiley & Sons.
- Culp, Christopher, and Merton Miller. 1995a. "Hedging in the Theory of Corporate Finance: A Reply to Our Critics." *Journal of Applied Corporate Finance* (Spring): 121–127.
- Culp, Christopher, and Merton Miller. 1995b. "Metallgesellschaft and the Economics of Synthetic Storage." *Journal of Applied Corporate Finance* (Winter): 62–75. www.rmcsinc.com/articles/JACF74.pdf.
- Dash, Eric, and Julie Creswell. 2008. "Citigroup Saw No Red Flags Even as It Made Bolder Bets." *New York Times*, November 22.
- Davidson, Andrew. 2007. "Six Degrees of Separation." www.securitization.net/pdf/content/ADC_SixDegrees_1Aug07.pdf.
- Davidson, Andrew, Anthony Sanders, Lan-Ling Wolfe, and Anne Ching. 2003. *Securitization*. Hoboken, NJ: John Wiley & Sons.
- Davies, Rob. 2008. "Genius or Blunder?" *Risk*. 3:22–27. www.risk.net/risk-magazine/feature/1498128/genius-blunder.
- Dembo, Ron. 1994. "Hedging in Markets That Gap." In *Handbook of Derivatives and Synthetics*, edited by Robert A. Klein and Jess Lederman. Chicago: Probus.
- Demeterfli, Kresimir, Emanuel Derman, Michael Kamal, and Joseph Zou. 1999. "Guide to Variance Swaps." *Risk* 6:54–59. A longer version is available at www.ederman.com/new/docs/gs-volatility_swaps.pdf.
- Derivative Strategies. 1998. "What Really Happened at UBS?" October. www.derivativesstrategy.com/magazine/archive/1998/1098fea1.asp.
- Derman, Emanuel. 1999. "Regimes of Volatility." *Risk* 4:55–59. A longer version is available at www.ederman.com/new/docs/risk-regimes_of_volatility.pdf.
- Derman, Emanuel. 2001. "Markets and Models." *Risk* 7:48–50. My references are primarily to the longer version available at www.ederman.com/new/docs/risk-price_verification.pdf.
- Derman, Emanuel, Deniz Ergener, and Iraj Kani. 1995. "Static Options Replication." *Journal of Derivatives* 2 (4): 78–95.
- Derman, Emanuel, and Iraj Kani. 1994. "Riding on a Smile." *Risk* 2:32–39.
- Derman, Emanuel, and Iraj Kani. 1998. "Stochastic Implied Trees." *International Journal of Theoretical and Applied Finance* 1 (1): 61–110. A longer version is available at www.ederman.com/new/docs/gs-options_replication.pdf.
- Derman, Emanuel, and Joseph Zou. 2001. "A Fair Value for the Skew." *Risk* 1:111–114.
- de Servigny, Arnaud, V. Peretyatkin, W. Perraudin, and Olivier Renault. 2003. "Portfolio Risk Tracker: Description of the Methodology." Standard & Poor's Risk Solutions.

- de Servigny, Arnaud, and Olivier Renault. 2004. *Measuring and Managing Credit Risk*. New York: McGraw-Hill.
- DeWit, Jan. 2006. "Exploring the CDS-Bond Basis." National Bank of Belgium working paper. www.nbb.be/doc/ts/publications/wp/wp104En.pdf.
- Diebold, Francis, Til Schuermann, and John Stroughair. 2000. "Pitfalls and Opportunities in the Use of Extreme Value Theory in Risk Management." In *Extremes and Integrated Risk Management*, edited by Paul Embrechts. London: Risk Books. www.ssc.upenn.edu/~fdiebold/papers/paper21/dss-f.pdf.
- Dixit, Avinash, and Robert Pindyck. 1994. *Investment under Uncertainty*. Princeton, NJ: Princeton University Press.
- Dowd, Kevin. 2005. *Measuring Market Risk*. 2nd ed. Hoboken, NJ: John Wiley & Sons.
- Duc, Francois, and Yann Schorderet. 2008. *Market Risk Management for Hedge Funds*. Hoboken, NJ: John Wiley & Sons.
- Duffie, Darrell. 2011. *How Big Banks Fail*. Princeton, NJ: Princeton University Press.
- Duffie, Darrell, Andreas Eckner, Guillaume Horel, and Leandro Saita. 2009. "Frailty Correlated Defaults." *Journal of Finance*. 5:2089–2123. www.darrellduffie.com/uploads/pubs/DuffieEcknerHorelSaita2009.pdf.
- Duffie, Darrell, Ada Li, and Theo Lubke. 2010. "Policy Perspectives on OTC Derivatives Market Structure." Federal Reserve Bank of New York Staff Reports no. 424. www.newyorkfed.org/research/staff_reports/sr424.pdf.
- Duffie, Darrell, and Kenneth Singleton. 2003. *Credit Risk*. Princeton, NJ: Princeton University Press.
- Dunbar, Nicholas, and Elisa Martinuzzi. 2012. "Goldman Secret Greek Loan Shows Two Sinners as Client Unravels." *Bloomberg News*, March 5. www.bloomberg.com/news/2012-03-06/goldman-secret-greece-loan-shows-two-sinners-as-client-unravels.html.
- Dwyer, Paula. 1996. "Sumitomo's Descent into the Abyss." *Business Week*, July 1. www.businessweek.com/1996/27/b348241.htm.
- Eichenwald, Kurt. 1995. *Serpent on the Rock*. New York: HarperBusiness.
- Embrechts, Paul. 2000. "Extreme Value Theory: Potential and Limitations as an Integrated Risk Management Tool." *Derivatives Use, Trading, and Regulation* 6:449–456. www.math.ethz.ch/~baltes/ftp/evtpot.pdf.
- Fabozzi, Frank, Sergio Focardi, and Petter Kolm. 2006. *Financial Modeling of the Equity Market*. Hoboken, NJ: John Wiley & Sons.
- Fabozzi, Frank, Sergio Focardi, and Petter Kolm. 2010. *Quantitative Equity Investing: Techniques and Strategies*. Hoboken, NJ: John Wiley & Sons.
- Falloon, William. 1998. "The Devil in the Documentation." *Risk* 11:32–35.
- Fay, Stephen. 1996. *The Collapse of Barings*. London: Richard Cohen Books.
- Federal Reserve Board. 2009. "The Supervisory Capital Assessment Program: Design and Implementation." www.federalreserve.gov/bankinforeg/bcreg20090424a1.pdf.
- Federal Reserve Board. 2011. "Supervisory Guidance on Model Risk Management." www.federalreserve.gov/bankinforeg/srletters/sr1107a1.pdf.
- Financial Crisis Inquiry Report. 2011. New York: PublicAffairs. Downloadable at www.gpo.gov/fdsys/pkg/GPO-FCIC/pdf/GPO-FCIC.pdf.

- Financial Stability Board. 2010. "Implementing OTC Derivatives Markets Reforms." www.financialstabilityboard.org/publications/r_101025.pdf.
- Financial Stability Forum. 2008. "Enhancing Market and Institutional Resilience." www.financialstabilityboard.org/publications/r_0804.pdf.
- Financial Stability Forum. 2009a. "Addressing Procyclicality in the Financial System." www.financialstabilityboard.org/publications/r_0904a.pdf.
- Financial Stability Forum. 2009b. "Principles for Sound Compensation Practices." www.financialstabilityboard.org/publications/r_0904b.pdf.
- Garfield, Andrew. 1998. "Peter Young Charged with Morgan Grenfell Fraud." *The Independent*. October 20. www.independent.co.uk/news/business/peter-young-charged-with-morgan-grenfell-fraud-1179481.html.
- Gatheral, Jim. 2006. *The Volatility Surface*. Hoboken, NJ: John Wiley & Sons.
- Giesecke, Kay, and Alexander Shkolnik. 2011. "Importance Sampling for Event Timing Models." www.stanford.edu/dept/MSandE/cgi-bin/people/faculty/giesecke/pdfs/is.pdf.
- Gillen, David, Yoolim Lee, and Bill Austin. 1999. "J.P. Morgan's Korean Debacle." *Bloomberg Magazine* 3:24–30.
- Glasserman, Paul. 2004. *Monte Carlo Methods in Financial Engineering*. New York: Springer-Verlag.
- Glasserman, Paul, and Jingyi Li. 2005. "Importance Sampling for Portfolio Credit Risk." *Management Science* 11:1643–1656. www2.gsb.columbia.edu/faculty/pglasserman/Other/is_credit.pdf.
- Global Legal Group. 2011. "The International Comparative Legal Guide to Securitization 2011." www.iclg.co.uk/index.php?area=4&kh_publications_id=194.
- Greenlaw, David, Jan Hatzius, Anil Kashyap, and Hyun Song Shin. 2008. "Leveraged Losses: Lessons from the Mortgage Market Meltdown." U.S. Monetary Policy Forum Report no. 2. <http://research.chicagobooth.edu/igm/events/docs/MPFReport-final.pdf>.
- Gregory, Jon. 2010. *Counterparty Credit Risk*. Chichester, UK: John Wiley & Sons.
- Group of Thirty. 1993. "Global Derivatives Study Group: Practices and Principles."
- Group of Thirty. 2009. "Financial Reform: A Framework for Financial Stability." www.group30.org/images/PDF/Financial_Reform-A_Framework_for_Financial_Stability.pdf.
- Gumerlock, Robert. 1999. "The Future of Risk." *Euromoney* 6:112–114.
- Gupta, Ajay. 1997. "On Neutral Ground." *Risk* 7:37–41.
- Gupton, Greg, Christopher Finger, and Mickey Bhatia. 1997. "Creditmetrics-Technical Document." www.defaultrisk.com/pp_model_20.htm.
- Gupton, Greg, and Roger Stein. 2005. "Losscalc V2: Dynamic Prediction of LGD." Moody's KMV. January. www.defaultrisk.com/_pdf6j4/LCv2_DynamicPredictionOfLGD_fixed.pdf.
- Hammond, John, Ralph Keeney, and Howard Raiffa. 1999. *Smart Choices*. Boston, MA: Harvard Business School Press.
- Hansell, Saul. 1997. "Joseph Jett: A Scoundrel or a Scapegoat?" *New York Times*, April 6.
- Harris, Larry. 2003. *Trading and Exchanges: Market Microstructure for Practitioners*. New York: Oxford University Press.

- Hasanhodzic, Jasmina, and Andrew Lo. 2007. "Can Hedge Fund Returns Be Replicated? The Linear Case." *Journal of Investment Management* 2:5–45. http://jasminah.com/Papers/hasanhodzic_lo07b.pdf.
- Helwege, Jean, Samuel Maurer, Asani Sarkar, and Yuan Wang. 2009. "Credit Default Swap Auctions." Federal Reserve Bank of New York Staff Report no. 372. www.newyorkfed.org/research/staff_reports/sr372.pdf.
- Henderson, Schuyler. 1998. "Credit Derivatives: Selected Documentation Issues." *Journal of International Banking and Financial Law* 9.
- Heston, Steven. 1993. "A Closed-Form Solution for Options with Stochastic Volatility." *Review of Financial Studies* 6 (2): 327–343. www.javaquant.net/papers/Heston-original.pdf.
- Holland, Kelley, and Linda Himelstein. 1995. "The Bankers Trust Tapes." *Business Week*, October 16. www.businessweek.com/1995/42/b34461.htm.
- Holton, Glyn. 2003. *Value-at-Risk*. San Diego, CA: Academic Press.
- Huertas, Thomas. 2011. *Crisis: Cause, Containment and Cure*. 2nd ed. Basingstoke, UK: Palgrave Macmillan.
- Hull, John. 2000. *Options, Futures, and Other Derivatives*. 4th ed. Upper Saddle River, NJ: Prentice Hall.
- Hull, John. 2002. *Options, Futures, and Other Derivatives*. 5th ed. Upper Saddle River, NJ: Prentice Hall.
- Hull, John. 2009. "The Credit Crunch of 2007." www.rotman.utoronto.ca/~hull/downloadablepublications/CreditCrunch.pdf.
- Hull, John. 2012. *Options, Futures, and Other Derivatives*. 8th ed. Upper Saddle River, NJ: Prentice Hall.
- Hull, John, Mirela Predescu, and Alan White. 2004. "The Relationship between Credit Default Swap Spreads, Bond Yields, and Credit Rating Announcements." *Journal of Banking and Finance* (November): 2789–2811. www.rotman.utoronto.ca/~hull/DownloadablePublications/HPWPaperonCDSSpreads.pdf.
- Hull, John, and Wulin Suo. 2001. "A Methodology for Assessing Model Risk and Its Application to the Implied Volatility Function Model." www.rotman.utoronto.ca/~hull/DownloadablePublications/HullSuoPaper.pdf.
- Hull, John, and Alan White. 1994. "Numerical Procedures for Implementing Term Structure Models II: Two-Factor Models." *Journal of Derivatives* 2 (2): 37–48.
- Hull, John, and Alan White. 2004. "Valuation of a CDO and an n th to Default CDS without Monte Carlo Simulation." *Journal of Derivatives* 2:8–23. www.rotman.utoronto.ca/~hull/DownloadablePublications/HullWhiteCDOPaper.pdf.
- Ineichen, Alexander. 2003. *Absolute Returns*. Hoboken, NJ: John Wiley & Sons.
- International Comparative Legal Guide to Securitization. 2011. London: Global Legal Group.
- International Monetary Fund. 2009. *Global Financial Stability Report*. April. www.imf.org/External/Pubs/FT/GFSR/2009/01/pdf/text.pdf.
- ISDA. 2012. "Interest Rate Swaps Compression: A Progress Report." February 23. www2.isda.org/functional-areas/research/studies/.
- Jackel, Peter. 2002. *Monte Carlo Methods in Finance*. Chichester, UK: John Wiley & Sons.

- Jackwerth, Jens, and Mark Rubinstein. 1996. "Recovering Probability Distributions from Option Prices." *Journal of Finance* 51 (5): 1611–1631. www.er.ethz.ch/teaching/Risk_neutral_Proba_Nonparametric_1996.pdf.
- Jacobs, Michael, and Ahmet Karagozoglu. 2011. "Modeling Ultimate Loss Given Default on Corporate Debt." *Journal of Fixed Income* 1:6–20.
- Jameson, Rob. 1998a. "Operational Risk: Getting the Measure of the Beast." *Risk* 11:38–41.
- Jameson, Rob. 1998b. "Operational Risk: Playing the Name Game." *Risk* 10:38–42.
- Jett, Joseph, with Sabra Chartrand. 1999. *Black and White on Wall Street*. New York: William Morrow.
- Jewson, Stephen, Anders Brix, and Christine Ziehmann. 2005. *Weather Derivatives Valuation*. Cambridge: Cambridge University Press.
- Jorion, Philippe. 2007. *Value at Risk*. 3rd ed. New York: McGraw-Hill.
- Kagel, J. H., and Alvin Roth. 1995. *Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.
- Kealhofer, Stephen, and Jeffrey Bohn. 2001. "Portfolio Management of Default Risk." Moody's KMV. www.moodysanalytics.com/~media/Insight/Quantitative-Research/Portfolio-Modeling/93-15-11-Portfolio-Management-of-Default-Risk.ashx.
- Khuong-Huu, Philippe. 1999. "Swaptions on a Smile." *Risk* 8:107–111. www.risk.net/data/Pay_per_view/risk/technical/1999/risk_0899_jpmorgan.pdf.
- Kim, Jongwoo, and Christopher Finger. 2000. "A Stress Test to Incorporate Correlation Breakdown." *Journal of Risk*. www.risk.net/digital_assets/4867/v2n3a1.pdf.
- Kooi, Mari. 1996. "Analyzing Sumitomo." <http://riskinstitute.ch/134800.htm>.
- Kotowitz, Y. 1989. "Moral Hazard." In *The New Palgrave: Allocation, Information, and Markets*, edited by John Eatwell, Murray Milgate, and Peter Newman. New York: W.W. Norton.
- Koutoulas, James, and John Roe. 2012. "White Paper: Background, Impacts & Solutions to MF Global's Demise." www.financialsense.com/node/725.
- Laurent, Jean-Paul, and Jon Gregory. 2003. "Basket Default Swaps, CDO's and Factor Copulas." http://laurent.jeanpaul.free.fr/basket_cdo_factor_copula_2003.pdf.
- Lee, Roger. 2001. "Local Volatilities under Stochastic Volatility." *International Journal of Theoretical and Applied Finance* 4:45–89.
- Leeson, Nick, with Edward Whitley. 1996. *Rogue Trader*. Boston: Little, Brown.
- Leibowitz, Martin, and Alfred Weinberger. 1981. "The Uses of Contingent Immunization." *Journal of Portfolio Management* 1:51–55.
- Leibowitz, Martin, and Alfred Weinberger. 1982. "Contingent Immunization—Part I: Risk Control Procedures." *Financial Analysts Journal* 38 (6): 17–31.
- Leibowitz, Martin, and Alfred Weinberger. 1983. "Contingent Immunization—Part II: Problem Areas." *Financial Analysts Journal* 39 (1): 35–50.
- Lewis, Michael. 2011. *The Big Short*. New York: W.W. Norton.
- Li, David. 2000. "On Default Correlation: A Copula Function Approach." RiskMetrics Working Paper no. 99-07. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=187289.

- Litterman, Robert. 1997a. "Hot Spots and Hedges, Part 1." *Risk* 3:42–45.
- Litterman, Robert. 1997b. "Hot Spots and Hedges, Part 2." *Risk* 5:38–42. <http://home.ku.edu.tr/~cdemiroglu/Teaching/Risk/hotspots.pdf>.
- Litterman, Robert, and Jose Scheinkman. 1988. "Common Factors Affecting Bond Returns." *Journal of Fixed Income* 1 (1): 54–61.
- Lo, Andrew. 2008. "Hedge Funds, Systemic Risk, and the Financial Crisis of 2007–2008." Testimony for the U.S. House of Representatives Committee on Oversight and Government Reform November 13 hearing on hedge funds. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1301217.
- Lo, Andrew. 2012. "Reading about the Financial Crisis: A 21-Book Review." *Journal of Economic Literature* 1:151–178. www.argentumlux.org/documents/JEL_6.pdf.
- Lowenstein, Roger. 2000. *When Genius Failed*. New York: Random House.
- Lowenstein, Roger. 2004. *Origins of the Crash*. New York: Penguin.
- Lowenstein, Roger. 2010. *The End of Wall Street*. New York: Penguin.
- Ludwig, Eugene. 2002. "Report on Allied Irish Bank." www.aibgroup.com/servlet/ContentServer?pagename=AIB_Investor_Relations/AIB_Download/aib_d_download&c=AIB_Download&cid=1015597173380&channel=IRHP.
- Madan, Dilip, Peter Carr, and Eric Chang. 1998. "The Variance Gamma Process and Option Pricing." www.math.nyu.edu/research/carrp/papers/pdf/VGEFRpub.pdf.
- Malcolm, Fraser, Pawan Sharma, and Joseph Tanega. 1999. *Derivatives: Optimal Risk Control*. London: Prentice Hall.
- Matytsin, Andrew. 1999. "Modeling Volatility and Volatility Derivatives." www.math.columbia.edu/~smirnov/Matytsin.pdf.
- Mayer, Martin. 1995. "Joe Jett: Did the Computer Make Him Do It?" *Institutional Investor* 3:7–11.
- Mayer, Martin. 1997. *The Bankers: The Next Generation*. New York: Truman Talley Books.
- McDonald, Robert. 2006. *Derivatives Markets*. Boston: Pearson Education.
- McKay, Peter. 1999. "Merrill Lynch to Pay Fines of \$25 Million for Scandal." *Dow Jones Newswires*, July 1.
- McLean, Bethany, and Peter Elkind. 2003. *The Smartest Guys in the Room*. New York: Penguin.
- McLean, Bethany, and Joe Nocera. 2010. *All the Devils Are Here*. New York: Penguin.
- McNeil, Alexander. 2000. "Extreme Value Theory for Risk Managers." In *Extremes and Integrated Risk Management*, edited by Paul Embrechts. London: Risk Books.
- Mello, Antonio, and John Parsons. 1995. "Maturity Structure of a Hedge Matters." *Journal of Applied Corporate Finance* (Spring): 106–120.
- Meucci, Attilio, 2005. *Risk and Asset Allocation*. Berlin: Springer-Verlag.
- Moody's. 2010. "Municipal Bond Defaults and Recoveries, 1970–2009." www.naic.org/documents/committees_e_capad_vos_c1_factor_review_sg_related_docs_moodys_us_municipal_bonds.pdf.

- Moody's. 2011a. "Corporate Default and Recovery Rates, 1920–2010." www.naic.org/documents/committees_e_capad_vos_c1_factor_review_sg_related_docs_moodys_corporate_default.pdf.
- Moody's. 2011b. "Sovereign Default and Recovery Rates, 1983–2010." www.naic.org/documents/committees_e_capad_vos_c1_factor_review_sg_related_docs_moodys_sovereign_default.pdf.
- Moosa, Imad. 2007. *Operational Risk Management*. New York: Palgrave Macmillan.
- Moosa, Imad. 2008. *Quantification of Operational Risk under Basel II*. New York: Palgrave Macmillan.
- Morini, Massimo. 2011. *Understanding and Managing Model Risk*. Chichester, UK: John Wiley & Sons.
- Neuberger, Anthony. 1996. "The Log Contract and Other Power Options." In *The Handbook of Exotic Options*, edited by Israel Nelken. Chicago: Irwin.
- Norris, Floyd. 2009. "A Lack of Rigor Costs MBIA." *New York Times*, November 12.
- O'Brien, Timothy L. 1999. "Taking the Danger out of Risk." *New York Times*, January 28.
- O'Kane, Dominic. 2008. *Modelling Single-Name and Multi-Name Credit Derivatives*. Chichester, UK: John Wiley & Sons.
- O'Kane, Dominic, and Robert McAdie. 2001. "Trading the Default Swap Basis." *Risk* 9, Lehman Brothers–sponsored article. www.scribd.com/dtayeconsulting/d/26559374-Explaining-the-Basis-Cash-Versus-Default-Swaps.
- Oyama, Tsuyoshi. 2010. *Post-Crisis Risk Management*. Singapore: John Wiley & Sons.
- Perold, Andre. 1998. "Capital Allocation in Financial Firms." Harvard Business School working paper.
- Perold, Andre. 1999a. Harvard Business School Case Study (9-200-007) of Long-Term Capital Management (A).
- Perold, Andre. 1999b. Harvard Business School Case Study (9-200-009) of Long-Term Capital Management (C).
- Pirrong, Craig. 2011. "The Economics of Central Clearing: Theory and Practice." ISDA Discussion Paper Series no. 1. Available at www2.isda.org/functional-areas/research/discussion-papers/.
- PricewaterhouseCoopers. 2011. *A Practitioner's Guide to Basel III and Beyond*. London: Thomson Reuters.
- Rajan, Raghuram. 2010. *Fault Lines*. Princeton, NJ: Princeton University Press.
- Ramberg, John, Edward Dudewicz, Pandu Tadikamalla, and Edward Mykytka. 1979. "A Probability Distribution and Its Use in Fitting Data." *Technometrics* 2:201–214.
- Rawnsley, Judith. 1995. *Total Risk*. New York: HarperCollins.
- Rebonato, Riccardo. 1998. *Interest-Rate Options Models*. 2nd ed. New York: John Wiley & Sons.
- Rebonato, Riccardo. 2002. *Modern Pricing of Interest Rate Derivatives*. Princeton, NJ: Princeton University Press.
- Rebonato, Riccardo. 2003. "Theory and Practice of Model Risk Management." In *Modern Risk Management: A History*, 223–248. London: Risk Books. Available on the web through CiteSeerX.

- Rebonato, Riccardo. 2004. *Volatility and Correlation*. 2nd ed. Hoboken, NJ: John Wiley & Sons.
- Rebonato, Riccardo. 2007. *Plight of the Fortune Tellers*. Princeton, NJ: Princeton University Press.
- Reiner, Eric. 1992. "Quanto Mechanics." *Risk* 3:59–63.
- Richardson, Matthew, and Lawrence White. 2009. "The Ratings Agencies." In *Restoring Financial Stability*, edited by Viral Acharya and Matthew Richardson. Hoboken, NJ: John Wiley & Sons.
- Risk Management Association. 2000. *Credit Risk Capital for Retail Products*. www.johningo.com/pdfdocs/RMARetailCreditCapit614.pdf
- RiskMetrics Group. 1996. *RiskMetrics Technical Document*. 4th ed. www.msici.com/resources/research/technical_documentation/td4e.pdf.
- Roubini, Nouriel, and Stephen Mihm. 2010. *Crisis Economics*. New York: Penguin.
- Rullière, Didier, Diana Dorobantu, and Areski Cousin. 2010. "An Extension of Davis and Lo's Contagion Model." <http://arxiv.org/pdf/0904.1653v2.pdf>.
- Salmon, Felix. 2009. "Recipe for Disaster: The Formula That Killed Wall Street." *Wired*, February 23. www.wired.com/techbiz/it/magazine/17-03/wp_quant?currentPage=all.
- Saunders, Anthony, and Linda Allen. 2010. *Credit Risk Measurement*. 3rd ed. Hoboken, NJ: John Wiley & Sons.
- Saunders, Anthony, Roy Smith, and Ingo Walter. 2009. "Enhanced Regulation of Large, Complex Financial Institutions." In *Restoring Financial Stability*, edited by Viral Acharya and Matthew Richardson. Hoboken, NJ: John Wiley & Sons.
- Schachter, Barry. 2001. "How Well Can Stress Tests Complement VaR?" www.gloriamundi.org/UploadFile/2010-2/stressandevt1.pdf.
- Schonbucher, Philipp. 2003. *Credit Derivatives Pricing Models*. Chichester, UK: John Wiley & Sons.
- Schutz, Dirk. 2000. *The Fall of UBS*. New York: Pyramid Media Group.
- Senior Supervisors Group. 2008. "Observations on the Risk Management Practices during the Recent Market Turbulence." March 6. www.newyorkfed.org/newsevents/news/banking/2008/SSG_Risk_Mgt_doc_final.pdf.
- Sharpe, William. 1992. "Asset Allocation: Management Style and Performance Measurement." *Journal of Portfolio Management* 2:7–19. www.stanford.edu/~wfsarpe/art/sa/sa.htm.
- Shaw, Julian. 1997. "Beyond VaR and Stress Testing." In *VaR – Understanding and Applying Value-at-Risk*, 211–224. London: Risk Publications.
- Shiller, Robert. 2003. *The New Financial Order*. Princeton, NJ: Princeton University Press.
- Shiller, Robert. 2008. *The Subprime Solution*. Princeton, NJ: Princeton University Press.
- Shirreff, David. 1998. "Another Fine Mess at UBS." *Euromoney* 11:41–43.
- Shirreff, David. 2000. "Lessons from the Collapse of Hedge Fund, Long-Term Capital Management." http://elsa.berkeley.edu/users/webfac/craigne/e137_f03/137lessons.pdf.
- Sifakis, Carl. 1982. *Encyclopedia of American Crime*. New York: Facts on File.
- Smithson, Charles. 2000. "What's New in the Options Markets?" *Risk* 5:54–55.

- Smithson, Charles, and Neal Pearson. 2008. "Valuing CDOs of ABSs." *Risk* 3:84–87. www.rutterassociates.com/pdf/Smithson&Pearson_ValuingTranchesOfCDOs-OnABS_RISK_Mar2008.pdf.
- Société Générale Mission Green Report. 2008. [www.sp.socgen.com/sdp/sdp.nsf/V3ID/8D7F118212725EC6C1257452005AA90E/\\$file/report%20part%203.pdf](http://www.sp.socgen.com/sdp/sdp.nsf/V3ID/8D7F118212725EC6C1257452005AA90E/$file/report%20part%203.pdf).
- Société Générale Special Committee of the Board of Directors Report to Shareholders. 2008. [www.sp.socgen.com/sdp/sdp.nsf/V3ID/C54C92B634428055C1257452005A7F61/\\$file/report%20part%201.pdf](http://www.sp.socgen.com/sdp/sdp.nsf/V3ID/C54C92B634428055C1257452005A7F61/$file/report%20part%201.pdf).
- Sorkin, Andrew Ross. 2010. *Too Big to Fail*. New York: Penguin.
- Squam Lake Group. 2010. *Squam Lake Report: Fixing the Financial System*. Princeton, NJ: Princeton University Press.
- Standard & Poor's. 1996. "April 1 Credit Week."
- Stiglitz, Joseph. 2010. *Freefall*. New York: W.W. Norton.
- Stigum, Marcia. 1989. *The Repo and Reverse Markets*. Homewood, IL: Dow Jones-Irwin.
- Swensen, David. 2000. *Pioneering Portfolio Management*. New York: Free Press.
- Taleb, Nassim. 1997. *Dynamic Hedging*. New York: John Wiley & Sons.
- Taleb, Nassim. 2010. *The Black Swan*. New York: Random House.
- Tett, Gillian. 2009. *Fool's Gold*. New York: Free Press.
- Thaler, Richard. 1992. *The Winner's Curse*. New York: Free Press.
- Tsiveriotis, Kostas, and Chris Fernandes. 1998. "Valuing Convertible Bonds with Credit Risk." *Journal of Fixed Income* 8:95–102.
- Tuckman, Bruce, and Angel Serrat. 2012. *Fixed Income Securities*. 3rd ed. Hoboken, NJ: John Wiley & Sons.
- Turner Review. 2009. "A Regulatory Response to the Global Banking Crisis." Financial Services Authority. www.fsa.gov.uk/pubs/other/turner_review.pdf.
- UBS. 2008. "Shareholder Report on UBS's Write-Downs." Available at https://www.ubs.com/global/en/about_ubs/investor_relations/share_information/shareholderreport.html.
- Varian, Hal. 1987. "The Arbitrage Principle in Financial Economics." *Journal of Economic Perspectives* 1 (2): 55–72. <http://ideas.repec.org/a/aea/jecper/v1y1987i2p55-72.html>.
- Vause, Nicholas. 2010. "Counterparty Risk and Contract Volumes in the Credit Default Swap Market." *BIS Quarterly Review* (December): 59–69. <http://finance.eller.arizona.edu/lam/fixi/creditmod/Vause.pdf>.
- Wang, Jin. 2001. "Generating Daily Changes in Market Variables Using a Multivariate Mixture of Normal Distributions." *Proceedings of the 33rd Winter Simulation Conference*. Washington, DC: IEEE Computer Society.
- Weiss, Gary. 1994. "What Lynch Left Out." *BusinessWeek*, August 22. www.businessweek.com/stories/1994-08-21/what-lynch-left-out.
- Whaley, Elizabeth, and Paul Wilmott. 1994. "Hedge with an Edge." *Risk* 10:82–85.
- Williams, Jeffrey. 1986. *The Economic Function of Futures Markets*. Cambridge: Cambridge University Press.
- Wilson, Charles. 1989. "Adverse Selection." In *The New Palgrave: Allocation, Information, and Markets*, edited by John Eatwell, Murray Milgate, and Peter Newman. New York: W.W. Norton.

- Wilson, Harry. 2011. "UBS Rogue Trader: Investigations Focus on Fictitious Hedges." *The Telegraph*, September 16. www.telegraph.co.uk/finance/financial-crime/8772540/Fictitious-hedges-see-UBS-rogue-trader-losses-climb-to-2.3bn.html.
- Wilson, Thomas. 1997. "Portfolio Credit Risk." *Risk* 9:111–117 and 10:56–61. Shorter version available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1028756.
- Wilson, Thomas. 1998. "Value at Risk." In *Risk Management and Analysis*, Vol. 1, edited by Carol Alexander, 61–124. New York: John Wiley & Sons.
- Winters, Bill. 1999. "Wrong-Way Exposure." *Risk* 7:52–55. www.risk.net/data/Paper_view/risk/technical/1999/risk_0799_jpmorgan.pdf.
- Wolfe, Eric. 2001. "NatWest Markets: E-Risk." www.riskmania.com/pdsdata/NatWestCaseStudy-erisk.pdf.
- Zandi, Mark. 2009. *Financial Shock*. Updated ed. Upper Saddle River, NJ: FT Press.

About the Companion Website

This book has an associated website (www.wiley.com/go/frm2e) containing Microsoft Excel spreadsheets that can be used to experiment with many of the concepts covered in the text. Most of the book's exercises are built around these calculators. Full documentation of the spreadsheets is contained in an accompanying Word document on the website. This appendix briefly describes the spreadsheets that are available. They are listed in the order you will encounter them in the text.

I have chosen to build all of these calculators in Excel with minimal use of user-defined functions for two reasons:

1. By using Excel rather than a programming language, I am hoping to maximize the number of readers who will be able to follow the calculations.
2. By minimizing user-defined functions, I am making the machinery of the computations as visible as possible.

These calculators have all been built specifically to illustrate the material of this book (and the course I teach on which the book is based). They are *not* designed to be used to actually manage risk positions. Specifically, they don't include the sort of detail, such as day count conventions, that is important in a trading environment. This sort of detail can be distracting when trying to learn broad concepts. For similar reasons, I have often chosen simple alternatives over more complex ones to illustrate a point. For example, I have chosen to represent volatility smile and skew through a simple formula that favors the ease of seeing the approximate impact of changes in input variables over the accuracy of a more complex representation.

Using such calculators for actual trading would require programs that are easily scalable; that is, they can readily accommodate adding a larger number of positions. I have deliberately sacrificed scalability for the ease of handling a small number of positions. Scalability nearly always requires the use of a programming language as opposed to a primarily spreadsheet-based approach. For readers who want to pursue building more robust calculators, and for teachers who want to assign exercises involving the building of scalable versions of some of these calculators, these spreadsheets should be able

to serve as good sources for parallel tests of computations, particularly since Excel gives an immediate display of all the numerical results of the intermediate stages of the calculations.

The spreadsheets, in the order of the corresponding material in the text, are as follows.

- The **MixtureOfNormals** spreadsheet produces series of random variables displaying fat tails and clustering of large moves by mixing together two normally distributed series. It is utilized for exercises in Sections 1.3 and 7.1.1.
- The **WinnersCurse** spreadsheet illustrates the mechanism of the winner's curse in auction situations, as explained in Section 2.4.
- The **VaR** spreadsheet computes VaR using three different methods—historical simulation, Monte Carlo simulation, and variance covariance. It enables the user to compare results obtained through the three methods and explore possible modifications. This is discussed in Section 7.1 and is used in Exercises 7.1 and 7.3.
- The **EVT** spreadsheet uses the extreme value theory formulas from the box “Key Results from EVT” in Chapter 7 to calculate VaR and short-fall VaR for selected percentiles.
- The **Rates** spreadsheet can be used either to value and compute risk statistics for a portfolio of linear instruments (such as forwards, swaps, and bonds) based on an input set of forward rates or to determine a set of forward rates that achieve an optimum fit with a given set of prices for a portfolio of linear instruments while maximizing the smoothness of the forward rates selected. This is discussed in Sections 10.2.1 and 10.4.
- The **Bootstrap** spreadsheet produces a comparison between the bootstrap and optimal fitting methodologies for extracting forward rates from an observed set of swap rates. This spreadsheet was used to produce Figure 10.1 in Section 10.2.1.
- The **RateData** spreadsheet contains a historical time series of U.S. interest rate data. It is used in Exercises 10.1 and 10.2.
- The **NastyPath** spreadsheet is an illustration of the size of losses that can be incurred when dynamically delta hedging an option. The example follows the dynamic delta hedging of a purchased call option over the 30 days of its life. This is discussed in the example in Section 11.2.
- The **PriceVolMatrix** spreadsheet computes the price-volatility matrix and volatility surface exposure for a small portfolio of vanilla European-style options. It illustrates the material discussed in Section 11.4.
- The **PriceVolMatrixCycle** spreadsheet is a particular run of the **PriceVolMatrix** spreadsheet that has been used to produce Table 11.5.

- The **VolCurve** spreadsheet fits a forward volatility curve to observed options prices. This spreadsheet is designed for European options other than interest rate caps and floors. This is discussed in Section 11.6.1.
- The **CapFit** spreadsheet fits a forward volatility curve to observed options prices for interest rate caps. Since caps are baskets of options, with each option within the basket termed a *caplet*, the spreadsheet needs to break each cap apart into its constituent caplets and price each one individually. This is discussed in Section 11.6.1.
- The **VolSurfaceStrike** spreadsheet interpolates implied option volatilities by strike for a given tenor, utilizing the methods discussed in Section 11.6.2. The interpolation can be performed in two modes:
 1. Implied volatilities are input for enough strikes to allow for reasonable interpolation.
 2. Implied volatilities are input for only three strikes.
- The **OptionRoll** spreadsheet is a variant of the **PriceVolMatrix** spreadsheet. It differs in the form of the optimization, which is set up to calculate a hedge that will minimize a future roll cost. It illustrates the material discussed in Section 11.6.3.
- The **OptionMC** spreadsheet calculates a single path of a Monte Carlo simulation of the delta hedging of a vanilla European-style call option position. It is designed to help you check your work for the Monte Carlo simulation exercise in Chapter 11 (Exercise 11.2).
- The **OptionMC1000** spreadsheet used in Exercise 11.2 is identical to the **OptionMC** spreadsheet except that it is set up for 1,000 time steps instead of 20 time steps.
- The **OptionMCHedged** spreadsheet used in Exercise 11.2 is a variant on the **OptionMC** spreadsheet. It calculates a single path of a Monte Carlo simulation of the delta hedging of the European-style call option hedged by two other call options with the same terms but different strike prices.
- The **OptionMCHedged1000** spreadsheet used in Exercise 11.2 is identical to the **OptionMCHedged** spreadsheet except that it is set up for 1,000 time steps instead of 20 time steps.
- The **BasketHedge** spreadsheet calculates and prices a piecewise-linear hedge using forwards and plain-vanilla European options for any exotic derivative whose payoffs are nonlinear functions of the price of a single underlying asset at one particular point in time. The spreadsheet consists of a **Main** worksheet that can be used for any payoff function and other worksheets that contain illustrations of how the **Main** worksheet can be used to hedge particular payoff functions. The particular functions

illustrated are a single-asset quanto, a log contract, interest rate convexity, and a compound option. This is discussed in Section 12.1.

- The **BinaryMC** spreadsheet provides a Monte Carlo simulation of binary options using the method discussed in Section 12.1.4.
- The **ForwardStartOption** spreadsheet is a slight variant on the **PriceVolMatrix** spreadsheet that can be used for the risk management of forward-starting options using the method discussed in Section 12.2.
- The **CarrBarrier** spreadsheet compares the pricing of barrier options using Carr's static hedging replication with those computed using standard analytic formulas. The cost of unwinding the static hedge is also calculated. This is discussed in Section 12.3.3.
- The **CarrBarrierMC** spreadsheet provides a Monte Carlo simulation of barrier options using Carr's static hedging replication, as discussed in Section 12.3.3.
- The **OptBarrier** spreadsheet illustrates the use of optimization to find a hedge for a down-and-out call barrier option, as discussed in Section 12.3.3.
- The **DermanErgenerKani** spreadsheet used in Exercise 12.6 calculates the pricing of knock-out barrier options using the Derman-Ergener-Kani static hedging replication. The cost of unwinding the static hedge is also calculated. It illustrates the material discussed in Section 12.3.3.
- The **DermanErgenerKani20** spreadsheet also calculates the pricing of knock-out barrier options using the Derman-Ergener-Kani static hedging replication. It displays intermediate results more explicitly than the **DermanErgenerKani** spreadsheet, but is less flexible for expansion to a larger number of time steps.
- The **DermanErgenerKaniDoubleBarrier** spreadsheet calculates the pricing of double barrier knock-out barrier options using the Derman-Ergener-Kani static hedging replication. The cost of unwinding the static hedge is also calculated. This is discussed in Section 12.3.5.
- The **DermanErgenerKaniPartialBarrier** spreadsheet calculates the pricing of partial barrier knock-out barrier options using the Derman-Ergener-Kani static hedging replication. The cost of unwinding the static hedge is also calculated. This is discussed in Section 12.3.5.
- The **BasketOption** spreadsheet computes an approximate value for the volatility to be used to price an option on a basket of assets and also computes the sensitivity of this volatility to changes in the volatility of the underlying asset and in the correlation between assets. This is discussed in Section 12.4.1.
- The **CrossHedge** spreadsheet simulates the hedging of a quanto that pays the product of two asset prices. The hedge is simulated using two different assumptions: if the asset price moves are completely uncorrelated

and if the asset price moves are completely correlated. This is discussed in Section 12.4.5.

- The **AmericanOption** spreadsheet calculates risk statistics for the early exercise value of American call options, as discussed in Section 12.5.1.
- The **TermStructure** spreadsheet illustrates the difficulties involved in pricing products that are dependent on yield-curve shape. It shows that different combinations of input parameters that result in the identical pricing of European caps/floors and swaptions can lead to very different pricings of these products. This is discussed in Section 12.5.2.
- The **Swaptions** spreadsheet calculates current swaption volatilities from current forward rate agreement (FRA) levels, forward FRA volatilities, and correlations between FRAs. Using the Solver, it can find forward FRA volatilities that will reproduce observed current swaption volatilities, as discussed in Section 12.5.3.
- The **CreditPricer** spreadsheet translates between par yields and default rates for risky bonds and also prices risky bonds based on the derived default rates, as discussed in Section 13.1.
- The **MertonModel** spreadsheet calculates default probabilities and the distance to default using the simplified model documented in Section 13.2.4.
- The **JumpProcessCredit** spreadsheet calculates default probabilities and credit spreads using the jump process model discussed in Section 13.2.4.1.
- The **CDO** spreadsheet calculates default probabilities for tranches of CDOs utilizing a Vasicek model with the large homogeneous portfolio (LHP) assumption, as discussed in Section 13.3.3.

- ACA Insurance, 107
- Accounting arbitrage, 107
- Accounting risk:
 defined, 30
 as form of reputational risk, 42
- Accrual swaps, 378
- Acharya, Viral, 84
- Ackerlof, George, 20
- Active management approach, 526–531
- Actuarial risk management:
 actuarial risk, defined, 2
 comparison with financial risk management, 2–4
 financial risk in estimating, 2
 liquid proxies in, 520
 moral hazard in, 13–14
 for positions that are born illiquid, 143–144
- Adjustable-rate mortgages (ARMs), 86–87. *See also* Financial crisis of 2007–2008
- Adoboli, Kweku, 67
- Adverse selection, 19–21
 controlling, 22
 defined, 20
 information asymmetry and, 20–21
 legal risk and, 37
- Agrawal, Deepak, 484
- AIG (American International Group), 86, 93, 97, 102, 107, 108, 114
- Allen, Franklin, 141
- Allen, Linda, 464, 481
- Allen, Peter, 331
- Allfirst First Maryland Bancorp. *See* Allied Irish Bank (AIB) case
- Allied Irish Bank (AIB) case, 31, 51, 57–59, 64, 65
- detection of unauthorized financial positions, 58
- development of unauthorized position, 57
- failure to detect unauthorized positions, 57–58
- further reading, 59
- incident, 57
- lessons to be learned, 59
- result, 57
- Almgren, Robert, 255–256
- Alpha factor, 515
- Altman, Edward, 464–465, 467, 479
- Amato, Jeffery, 84, 465
- American International Group (AIG), 86, 93, 97, 102, 107, 108, 114
- American options:
 defined, 426
 difficulty in valuing, 409
 European options versus, 426
 hedging, 435
 intensity of use, 363
- AmericanOption spreadsheet, 427–428, 551
- Analysis of overrides, 233–234
- Analysis of revenue, 156–157
- Andersen, Leif, 404, 434, 490
- Andreassen, Jesper, 404, 434
- Andrews, Edmund, 87
- Araten, Michel, 469
- Arbitrage:
 accounting, 107
 Allied Irish Bank (AIB) case, 31, 51, 57–59, 64, 65
 arbitrage theory in decomposing risk, 142

- Arbitrage (*Continued*)
 Barings Bank case, 31, 51, 55–57, 64, 66
 cash-and-carry, 259–260, 300–301, 302
 internal, 167
 no-arbitrage principle, 239–241
Arbitrage pricing theory (APT), 141
Arbitrageurs, 26–27
Armitstead, Louise, 132
Arora, Navneet, 484
Arthur Andersen, 12–13, 79–80
Artzner, Philippe, 188
Ashcroft, Adam, 85, 90
Asian credit crisis of 1997, 187, 206
Asian options, 413
Askin, David, 66–67
Askin Capital Management, 66–67
Asset-backed securities, 278–279. *See also* Credit default swaps (CDS)
Asset liquidity risk. *See also* Funding liquidity risk
 basis risks versus, 3–4, 255–256, 289
 defined, 3
 funding liquidity risk versus, 42
 positions that achieve illiquidity, 143
 positions that are born illiquid, 143–144
 in risk measurement, 142–147
Asset-or-nothing options, 371
Asset swaps:
 credit default swaps (CDS) versus, 447, 448
 in credit risk management, 447
AT&T, 263
Auctions, winner's curse and, 22–24
Austin, Bill, 80

Background checks, 33, 87–88
Back office:
 defined, 10
 fraud risk and, 32–35
Back-testing, in value at risk (VaR)
 analysis, 191, 233
Backwardation, 303
Bahar, Reza, 461

Bai, Jennie, 456
Bankers Trust (BT) case, 41, 77–79, 367
Bank for International Settlements (BIS), 101, 115, 122, 505–506
Bank of America, 274–275
Bank of England, 126
Bank of New York, 36
Bankruptcy:
 impact of bankruptcy law, 452–453
 legal risk and, 39–40
 need for more orderly bankruptcy proceedings, 131
 skew pattern in equity markets and, 350
 “too big to fail” mentality and, 11, 72, 105–106, 114, 121, 124
Banziger, Hugo, 62
Barings Bank case, 31, 51, 55–57, 64, 66
 detection of unauthorized positions, 56
 development of unauthorized positions, 55
 failure to detect unauthorized positions, 55–56
 further reading, 57
 incident, 55
 lessons to be learned, 56
 result, 55
Barrier options, 381–404
 barriers, defined, 382
 Carr hedge, 382, 388, 391–401
 defined, 382
 Derman-Ergener-Kani hedge, 387–391, 403–404
 double, 382
 drift, 382–383
 dynamic hedging models, 385–387
 knock-in (down and in), 382
 knock-out (down and out/up and out), 382
 ladder options, 402–403
 lookback options, 402
 partial-time, 382
 put-call symmetry, 391–392
 with rebates, 402
 standard analytic models, 383–384

- static hedging models, 387–391
- value of barrier based on analytic formula, 384
- Base correlation, 498
- Basel Committee on Banking Supervision, 29, 46–47, 130, 193, 196, 209, 219
- Basis risk:
 - CDS-bond, 454–456
 - defined, 3, 255–256
 - liquidity risk versus, 3–4, 255–256, 289
- Basis swaps, 298–299
- BasketHedge spreadsheet, 364, 368–370, 381, 439, 549–550
- BasketOption spreadsheet, 412–413, 441, 550
- Basu, Susanta, 490
- Baxter, Martin, 418
- Bear Stearns, 95, 101, 106
- Bennett, Oliver, 450
- Bermudan options, 426
 - hedging, 435, 436
 - intensity of use, 363
- Bermudan swaptions, 432, 433–434
- Bet options. *See* Binary options
- Bhatia, Mickey, 461, 465, 466, 481
- Bilateral counterparty risk, 521
- Binary credit default swaps, 449–450
- BinaryMC spreadsheet, 376, 440, 550
- Binary options, 371–377
 - asset-or-nothing option, 371
 - cash-or-nothing option, 371
- Binomial tree model, 425–426
- Black-Derman-Toy model, 434, 435
- Black-Karasinski model, 434, 435
- Black-Scholes option pricing model, 98, 144, 157–158
 - for exotic options, 359–361, 383, 385, 387, 388, 405, 427
 - model risk and, 210–211, 220
 - for vanilla options, 311–324, 331, 335, 344–345, 349–350, 355–356
- Black Swan, The* (Taleb), 138
- Bleed (Taleb), 343–344
- Bluhm, Christian, 490
- Bodie, Zvi, 141
- Bohn, Jeffrey, 464–467, 475–479, 481–484, 492
- Bonds:
 - CDS-bond basis risk, 454–456
 - in credit risk management, 447
 - market for, 447
- Book running. *See* Market makers/
market making
- Bookstaber, Richard, 95, 132
- Bootstrapping, 287–288, 290
- Bootstrap spreadsheet, 289, 548
- Borrowing costs, 299–304
 - forward prices and, 303–304
 - nature of borrowing demand, 299–300
 - possibility of cash-and-carry arbitrage, 300–301
 - seasonality, 302
 - variability of storage costs, 301
- Bouchet, Michel, 486
- Brace-Gatarek-Museila (BGM) model, 425
- Brealey, Richard, 141
- Breeden-Litzenberger theorem, 364
- Breuer, Thomas, 200, 486
- Brindle, Andy, 517–519
- British Bankers Association, 268–269, 297
- Brix, Anders, 144
- Broadie, Mark, 409
- Broom, Giles, 67
- Brown, Aaron, 27, 132, 265, 510
- Bruck, Connie, 237
- Brunnermeier, Markus, 85–88, 90, 111
- Bucay, Nisso, 485
- Burghardt, Galen, 302, 347
- Burnout, 424
- Business Week*, 87
- Cabiallavetta, Mathis, 61
- Calendar spread, 336, 337
- Call spreads, 331–334, 372
- Canabarro, Eduardo, 517
- Cancel-and-correct activity, 62–63, 65
- CapFit spreadsheet, 347, 549

- Capital asset pricing model (CAPM), 141–142, 484
- Capital structure, leverage in, 477
- Caps/caplets, 347, 430, 432
- Carr, Peter, 245, 364, 382, 386, 388, 391, 394, 404
- CarrBarrierMC spreadsheet, 396, 440, 550
- CarrBarrier spreadsheet, 394–396, 440, 550
- Carr hedge:
 - advantages, 393–394
 - broader applications, 404
 - comparison with other static hedging models, 387–390
 - deriving, 393
 - development of, 382
 - key points, 391, 394–396
 - put-call symmetry and, 391–392
 - static hedge, 383–384, 397–401
- Carty, Lea, 467
- Cash-and-carry arbitrage, 259–260, 300–301, 302
- Cash-or-nothing options, 371
- Cash settle, 274–275
- Cass, Dwight, 450
- CDO spreadsheet, 496, 551
- CDX index, 493
- Central counterparty clearinghouse (CCP), 127
- Chain letter frauds, 17
- Chang, Eric, 386
- Chase Manhattan Bank, 105, 123, 141, 202, 205–206, 249, 319, 469
- Chase Manhattan Bank/Drysdale Securities case, 45, 51, 52–53
- detection of unauthorized positions, 53
- development of unauthorized positions, 52–53
- failure to detect unauthorized positions, 53
- further reading, 53
- incident, 52
- lessons learned, 53
- result, 52
- Chew, Lillian, 51, 57, 77
- Chief financial officer, funding liquidity risk control and, 43
- Ching, Anne, 424
- Cholesky decomposition method, 178
- Chou, Andrew, 394, 404
- Chriss, Neil, 22, 255–256, 386
- Citigroup, 40, 79–80, 94–98, 101, 106
- Clementi, Gian Luca, 121
- Clewlow, Les, 386, 426
- Cliquet options, 378–379
- Closeouts:
 - exchange-traded derivatives, 506, 507, 508, 509, 510, 511
 - over-the-counter derivatives, 513, 515, 516, 517, 519, 520, 529
- Cochrane, John, 141
- Collateral:
 - Chase Manhattan Bank/Drysdale Securities case, 45, 51, 52–53
 - continuous collateral calls on futures contracts, 273
 - nondeliberate incorrect information and, 35
 - Société Générale case, 31, 61–66, 67
- Collateralization approach, 515–526
- ISDA Master Agreement, 515–516
- wrong-way risk, 521–526
- Collateralized debt obligations (CDOs). *See also* Credit default swaps (CDS); Financial crisis of 2007–2008
- CDO creators in financial crisis of 2007–2008, 88–89, 111, 116–117
- computational approximations, 226
- credit risk management and, 445–446
- default basket, 495
- equity tranches, 89, 92–93, 94–95, 494–495
- faulty models in financial crisis of 2007–2008, 98–99
- illiquidity of, 143–144
- mezzanine tranches, 94, 102, 104–105, 494
- multiname credit derivatives, 493–501

- risk management and reporting
 - for portfolio credit exposures, 490–492
 - senior tranches, 494
 - super-senior tranches, 94–102, 104–105, 106–107, 113, 494
- Collin-Dufresne, Pierre, 456
- Commercial paper:
 - in credit contagion of 2007–2008, 109
 - estimating amount owed at default, 468–469
- Commodities:
 - broad definition, 253–254
 - financial, 254
 - physical, 254, 259–260, 301–304
- Component VaR, 204–205
- Compound options, 379–381
- Compound worksheet, 381
- Comptroller of the Currency, 12
- Computer errors, 36
- Conduct of customer business:
 - Bankers Trust (BT) case, 41, 77–79, 367
 - Enron case, 79–80
 - other cases, 80–81
- Constant-maturity Treasury (CMT), 295
- Contagion, 482
 - credit contagion, 108–109, 115, 126–129
 - market contagion, 109–111, 115, 129–131
- Contango, 303
- Contingency plans:
 - for disaster risk, 36
 - for funding liquidity risk, 43
 - for model risk and evaluation control, 213–214
- Contingent credit default swap (CCDS), 531
- Contingent immunization strategy, 136, 161
- Contingent premium options, 377–378
- Continuous review, 232–234
 - analysis of overrides, 233–234
 - back-testing, 233
 - daily P&L reconciliation, 232–233
- Contracts, risk of unenforceable, 37–40
- Control variate technique in modeling, 360
- Convergence position, 414
- Convexity:
 - convexity adjustments, 297
 - of credit instruments, 453–454, 457
 - defined, 330
 - price-vol-matrix and, 330, 334–335, 343
 - of single-payout options, 370
- Convexity risk, 305–307
- Cooley, Thomas, 121
- Copula methodology, 98–99, 180, 211, 482, 487, 489–490, 497–498
- Cordell, Larry, 83, 93
- Correlation between price and exercise, 422–424
- Correlation-dependent interest rate options, 425–439. *See also* Correlation-dependent options
 - Brace-Gatarek-Musiela (BGM) models, 425
 - described, 362–364
 - Heath-Jarrow-Morton (HJM) models, 425
 - intensity of use, 363
 - relationship between forwards treated as constant, 426–429
 - relationship between swaption and cap prices, 437–439
 - term structure models, 430–436
- Correlation-dependent options, 404–425
 - correlation between price and exercise, 422–424
 - described, 362
 - index options, 413–414
 - interest-rate options (*see* Correlation-dependent interest rate options)
 - linear combinations of asset prices, 405–409
 - nonlinear combinations of asset prices, 417–422

- Correlation-dependent (*Continued*)
 - options to exchange one asset for another, 415–417
 - risk management of options on linear combinations, 409–413
- Counterparty credit risk, 505–531
 - of CDS-bond basis risk, 456
 - exchange-traded derivatives, 128–129, 506–511
 - over-the-counter derivatives, 128–129, 508–509, 512–531
 - overview, 505–506
- Counterparty risk groups (CRGs), 526–531
- Counterparty Risk Management Policy Group, 74
- Countrywide, 84
- Cousin, Areski, 482
- Coval, Joshua, 500–501
- Cox-Ross-Rubinstein binomial tree, 426–429
- Coy, Peter, 87
- Crack spread, 260
- Credit concentration, 486
- Credit contagion, in financial crisis of 2007–2008, 108–109, 115, 126–129
- Credit default swaps (CDS), 447–451
 - asset swaps versus, 447, 448
 - binary, 449–450
 - CDS-bond basis risk, 454–456
 - counterparty credit exposure through, 525–526
 - in credit contagion of 2007–2008, 108–109
 - credit risk management and, 445–446
 - legal basis risk, 450
 - loss given default (LGD), 447–448
 - in market contagion of 2007–2008, 109–111
 - Monte Carlo simulation and, 178
 - origins, 447
 - total return swaps, 276–278, 450–451
- Credit exposure mitigation techniques, 291
- Credit Grades, 475–476
- Credit instruments, 447–451
 - asset swaps, 447
 - bonds, 447, 454–456
 - collateralized debt obligations (CDOs) (*see* Collateralized debt obligations [CDOs])
 - convexity of, 453–454
 - credit default swaps (*see* Credit default swaps [CDS])
- Credit-linked note (CLN), 523–524
- CreditMetrics, 481–484
- Creditors:
 - moral hazard and, 14
 - outside monitors for, 11–12
- CreditPricer spreadsheet, 453, 551
- Credit rating agencies:
 - criticism of, 12
 - estimating probability of default, 459–464
 - in financial crisis of 2007–2008, 89–92, 111, 113, 117–118
 - information asymmetry and, 11, 12
 - investment bank reliance on, 99–100
 - relationship with investment banks, 90
 - use of forecasts, 90
- Credit risk management, 279, 445–503
 - counterparty (*see* Counterparty credit risk)
 - credit instruments, 447–451
 - large money moves and, 195–196
 - legal risk versus, 39
 - loan-equivalent approach, 513–515
 - models of short-term credit exposure, 451–456
 - multiname (*see* Multiname credit derivatives)
 - portfolio (*see* Portfolio credit risk)
 - risk reporting for market credit exposures, 456–457
 - single-name (*see* Single-name credit risk)
- Credit spread curve, 456–457
- Credit value adjustment (CVA), 521
- Creswell, Julie, 101

- Crosbie, Peter, 476–477, 478
Cross-currency swaps, 525
CrossHedge spreadsheet, 420, 441–442, 527, 550–551
Crouhy, Michel, 67
Crush spread, 260
Csiszar, Imre, 200, 486
Culp, Christopher, 76, 206
- Daily Telegraph*, 132
Daiwa Bank, 66
Danish mortgage structure, 116
Dash, Eric, 101
DataMetricsRatesData spreadsheet, 431
Davidson, Andrew, 85, 87, 90, 116, 424
Davies, Rob, 62–64
Ddeltavol (Taleb), 343–344
De Angelis, Anthony (“Salad Oil King”), 52
Default basket, 495
Default risk:
 comparison of rates of loss given default, 466
 correlation with market values, 525
 default percentages by year, 480
 estimating amount owed at default, 468–471
 estimating default correlations, 479–482
 estimating loss given default, 465–468
 estimating probability of default, 458–465
 five-year default rates, 463–464
 leverage in measuring, 73
 rating agency evaluations, 459–464
 statistical modeling, 464–465
Delbaen, Freddy, 188
Dembo, Ron, 396
Demeterfli, Kresimir, 368–369
Derivative Strategies, 61
Derman, Emanuel, 184, 209, 228, 241, 243, 245–247, 342, 361, 385–391, 403–404, 409, 436
DermanErgenerKani20 spreadsheet, 550
DermanErgenerKaniDoubleBarrier spreadsheet, 403–404, 441, 550
Derman-Ergener-Kani hedge:
 broader applications, 403–404
 comparison with other static hedging models, 387–390
 key points, 390
 unwinding, 391
DermanErgenerKaniPartialBarrier spreadsheet, 403–404, 550
DermanErgenerKani spreadsheet, 550
Derman-Kani dynamic hedge, 383
de Servigny, Arnaud, 459, 461–462, 465–467, 476, 478, 484
Deutsche Bank, 62, 104, 106
DeWit, Jan, 455
Diebold, Francis, 190
Digital options. *See* Binary options
Direct borrowing and lending, 264–267, 270, 280
Direct negotiation, winner’s curse and, 22–24
Disaster risk, 36
Distance to default, 476–477, 488–489
Divergence position, 414
Diversifiable/idiosyncratic risk, 141–142, 197
Dixit, Avinash, 4
Documentation:
 of contribution of risk positions, 203–204
 legal risk and, 37–40
 of model risk and evaluation control, 218–219
 of model verification, 220
Dollar gamma, 331
Dorobantu, Diana, 482
Double barrier options, 382
Dowd, Kevin, 170, 173, 175, 176, 178, 180, 181, 183, 185, 186, 188–192, 199, 203–204
Down and in (knock-in), 382
Down and out (knock-out), 382, 383
Drexel Burnham Lambert, 237
Drift, 382–383, 426
Drysedale Government Securities. *See* Chase Manhattan Bank/Drysedale Securities case

- Duc, Francois, 161
Dudewicz, Edward, 175
Duffie, Darrell, 128, 454, 455, 482, 483, 510, 517
Dunbar, Nicholas, 80
Dwyer, Paula, 66
Dynamic hedging strategies:
 dynamic delta hedging, 315–316, 320, 344–345, 375–376
 hedge slippage and, 17–18
 impact of drift and mean reversion, 328
 models for barrier options, 385–387
 Monte Carlo simulation versus dynamic delta hedging, 324
 nature of, 313–318
 path dependence of, 318–323
 performance of, 314
 simulation of dynamic hedging, 321–329
 for vanilla options, 313–329
- Eber, Jean-Marc, 188
Economic scenario stress tests, 193–197
Economist magazine, 79, 102–103, 140, 268–269, 448–450
Eichenwald, Kurt, 80
Eincomb, Stephen, 331
Eisman, Steve, 104
Elkind, Peter, 80, 523
Ellis, Katrina, 391, 404
Embrechts, Paul, 190–191
Enron, 9, 12–13, 40, 79–80, 523
Enterprise risk, 30, 44
Equity spot risk, 258–259
Equity tranches, 89, 92–93, 94–95, 494–495
Ergener, Deniz, 385–391, 403–404
ERisk, 66
European options. *See also* Vanilla option risk management
 American options versus, 426
 conventions for, 311–312
 intensity of use, 363
European swaptions, 430
- EVT spreadsheet, 548
Exchange-traded derivatives, 128–129, 506–511
 closeouts, 506, 507, 508, 509, 510, 511
 loss mutualization, 506
 margining, 506, 507, 508, 510, 511
 netting, 506, 507, 511
 novation, 506
Exotic option risk management, 359–443
 correlation-dependent interest rate options, 362–364, 425–439
 correlation-dependent options, 362, 404–425
 exotic options, defined, 311, 359
 intensity of use of option structures in various markets, 363
 path-dependent options, 362, 381–404
 single-payout options, 362, 364–378
 time-dependent options, 362, 378–381
 valuation reserves and, 152
Extrapolation approach:
 based on time period, 352–355
 extreme value theory (EVT) in, 190–191
Extreme value theory (EVT), 190–191, 198, 548
- Fabozzi, Frank, 176, 233
Factor-push stress tests, 199–200
Fair value:
 defined, 159
 risk measurement for position taking, 159–161
Falloon, William, 450
Fannie Mae, 84
Fargher, Neil, 479
Fay, Stephen, 57
Federal Deposit Insurance Corporation (FDIC), 319
Federal Reserve Bank of New York, 128
Federal Reserve Bank of Philadelphia, 93

- Federal Reserve Board (FRB), 123, 209, 213–223, 232, 234, 252
- Federal Reserve System, 12, 30
- Fernandes, Chris, 416
- FICO scores, 87
- Finance function, 8, 10
- Financial commodities, 254
- Financial Crisis Inquiry Commission, 85, 95
- Financial Crisis Inquiry Report (FCIR), 84, 85, 91, 92, 94–95, 98, 102, 106–108
- Financial crisis of 2007–2008, 1, 83–132
- actuarial versus financial risk management and, 3
 - broader lessons, 132
 - CDO creators in, 88–89, 111, 116–117, 496–497
 - credit contagion in, 108–109, 115, 126–129
 - credit rating agencies in, 89–92, 111, 113, 117–118
 - crisis in CDOs of subprime mortgages, 85–108
 - equity tranches in, 89, 92–93, 94–95
 - FCIR report on, 84, 85, 91, 92, 94–95, 98, 102, 106–108
 - insurers in, 96–97, 106–108, 114–115, 126
 - investment banks in, 93–106, 112–114, 118–126
 - investors in, 92–93, 111, 118
 - lessons for regulators, 115–131
 - lessons for risk managers, 111–115
- Li's Gaussian copula formula and, 98–99, 211, 482, 489–490, 497–498
- market contagion in, 109–111, 115, 129–131
- overview, 83–85
- spread of the crisis, 108–111
- subprime mortgage originators in, 86–88, 111, 116
- “too big to fail” mentality, 105–106, 114, 121, 124
- Financial disasters, 49–81
- conduct of customer business, 77–81
 - large money moves, 68–77, 195–196, 201–202
 - misleading reporting, 49–67
- Financial risk management:
- actuarial risk management versus, 2–4
 - broader applications of, 2
 - credit risk (*see* Credit risk management)
 - default risk (*see* Default risk)
 - essential components, 136–142
 - financial versus actuarial risk, 3–4
 - forward risk (*see* Forward risk management)
 - instruments that lack liquidity, 144–147, 150–151
 - options risk (*see* Vanilla option risk management; Exotic option risk management)
 - quantification in, 2
 - through risk aggregation, 4
 - risk control, 161–167
 - through risk decomposition, 4
 - risk measurement in, 133–161
 - spot risk (*see* Spot risk management)
- Financial Stability Board, 115, 120
- Financial Stability Forum, 114, 115, 120–121, 124, 126, 129, 130
- Finger, Christopher, 178, 200, 461, 465, 466, 481, 484
- Fitch, 89
- Floors/floorlets, 347, 430
- Flows:
- indexed, 295–299
 - in pricing illiquid flows by interpolation, 284–291
 - representing promised deliveries, 282, 293–295
 - stack-and-roll hedge and, 291–293
- Focardi, Sergio, 176, 233
- Fons, Jerome, 91

- Foreign exchange spot risk, 257–258
- Forward contracts, 253, 272, 273, 280.
See also Forward risk management models in which relationship between forwards is treated as constant, 426–429
- Forward prices, borrowing costs and, 303–304
- Forward rate agreements (FRA), 274–275, 280, 430, 432, 437–439
- Forward risk, 279
- Forward risk management, 263–310
asset-backed securities, 278–279
direct borrowing and lending, 264–267, 270, 280
factors impacting borrowing costs, 265–269, 299–304
firm-level risk management, 307
forward contracts, 253, 272, 273, 280
forward prices for different time periods, 259–260
forward rate agreements (FRAs), 274–275, 280
forward transactions, defined, 253
futures contracts, 272–274, 280
instruments, 269, 270–282
interest rate swaps, 275–276, 281
Kidder Peabody case, 21, 31, 51, 53–55, 57, 64–66, 234–235
models (*see* Forward risk models)
overlap between interest rate risk and credit risk, 266–269
overview, 263–269
repurchase agreements, 69, 271–272, 280, 294–295
risk comparisons, 280–281
risk management reporting system, 269, 304–308
spot versus forward positions, 263–264
total return swaps, 276–278
- Forward risk models, 265, 267–269, 282–299
flows representing promised deliveries, 282, 293–295
indexed flows, 295–299
pricing illiquid flows by interpolation, 284–291
pricing long-dated illiquid flows by stack and roll, 291–293
- Forward-start caplets, 432
- Forward-start options, 378–379
hedge at rollover, 380
- ForwardStartOption spreadsheet, 381, 550
- ForwardStart spreadsheet, 379
- FRA (forward rate agreements), 274–275, 280, 430, 432, 437–439
- Frailty analysis, 482
- Fraud risk, 31–35
deception about earnings, 31
deception about positions, 31–32
reducing, 32–35
- Freddie Mac, 84
- Friedman, Billings, Ramsey, 87
- Front office:
components of, 8
defined, 8
fraud risk and, 32–35
healthy skepticism about, 95–96
hedge slippage and, 17–18
information asymmetry and, 8–10
legal risk and, 38–39
modeling choices of, 238
nondeliberate incorrect information and, 35
risks that are difficult to identify, 231
“too big to fail” mentality and, 105–106, 114
- Funding cost, of CDS-bond basis risk, 456
- Funding liquidity risk, 42–44
asset liquidity risk versus, 42
components of, 42–43
defined, 30
- Futures contracts, 272–274, 280
- Futures exchanges:
credit in, 70
over-the-counter markets versus, 70

- Galai, Dan, 67
- Gamma:
 defined, 315, 330–331, 343–344
 dollar, 331
 hedging costs, 327
 price-vol-matrix and, 330–331, 334, 343–344
- Gap market risk, 523
- Garfield, Andrew, 80
- Gates, Bill, 138
- Gatheral, Jim, 369, 377, 386, 387, 435
- Gaussian copula formula, 98–99, 211, 482, 489–490, 497–498
- General Electric (GE), 54, 407–408
- Generalized autoregressive conditional heteroscedasticity (GARCH), 176
- Gibson Greetings, 41, 77–78, 367
- Giescke, Henning, 63
- Giesecke, Kay, 490
- Gilbert, W. S., 14, 303
- Gillen, David, 80
- Glasserman, Paul, 409, 490
- Global Industry Classification Standard (GICS), 259
- Global Legal Group, 116–117
- Golden parachutes, 106
- Goldman Sachs, 22, 80, 106, 203–206
- Gone on special, 294–295
- Government:
 conflict of interest, 12
 information asymmetry and, 11
 lessons from financial crisis of 2007–2008, 115–131
 outside monitors for, 12
- Grace period, 508
- Granger, Nicholas, 331
- Granite Capital, 66–67
- Granville-Barker, Harley, 18–19
- Greece, 80, 448
- Greenlaw, David, 110
- Greenspan, Alan, 106
- Gregory, Jon, 450, 490, 505–507, 509, 512, 515–519, 521–523, 525, 526, 528, 531
- Gros Lambert, Bertrand, 486
- Gross position regulation, 63–64, 65
- Group of Thirty (G-30), 115–119, 125, 129–131
 recommendations on trading risk, 136, 137, 147, 156–157, 159, 169
- Group of Twenty (G-20), 115, 127–128
- Grunkemeyer, Barbara, 87
- Gumerlock, Robert, 4
- Gupta, Ajay, 327
- Gupta, Vishal, 391, 404
- Gupton, Greg, 461, 465–467, 481, 484
- Guys and Dolls* (Runyon), 44
- Hamanaka, Yasuo, 66
- Hamilton, David, 467
- Hammond, John, 6
- Hansell, Saul, 55
- Hanweck, Gerald, 347
- Harris, Larry, 27
- Harvard Business School case studies, 75
- Hasanhodzic, Jasmina, 243
- Hatzius, Jan, 110
- Heath, David, 188
- Heath-Jarrow-Morton (HJM) model, 425
- Heat maps, 331
- Hedge funds, need for broader regulatory oversight, 131
- Hedge slippage, 17–18, 232, 243
- Hedging. *See* Dynamic hedging strategies; Static hedging strategies
- Helwege, Jean, 449, 450
- Henderson, Schuyler, 450
- Heston, Steven, 346, 386
- Heston model, 346
- High-yield debt, 307
- Himmelstein, Linda, 78
- Historical data:
 simulation of P&L distribution, 170–171, 173–174, 180–183
 stress tests relying on, 192, 197–201
- Holland, Kelley, 78
- Holton, Glyn, 177

- Huang, Yilin, 83, 93
- Huertas, Thomas, 131
- Hull, John, 4, 83–86, 89, 115, 142, 175, 176, 225, 229, 265, 274, 283, 288, 289, 295–297, 311–313, 348–350, 360–361, 370, 378, 379, 383, 392, 408, 409, 413, 415, 416, 418, 425, 426, 434, 435, 453, 455, 465, 471, 488, 490, 493–495, 509
- Hull-White model, 434–435
- IBM, 407–408
- Idiosyncratic risk. *See* Diversifiable/ idiosyncratic risk
- Iguchi, Toshihida, 66
- Illegal actions, risk of, 40
- Illiquid instruments. *See also* Collateralized debt obligations (CDOs)
- asset liquidity risk and, 142–147
 - choice of liquid proxy, 144–146, 243–245, 247–249
 - choice of model validation approach, 241–243
 - design of Monte Carlo simulation, 245–247
 - implications for marking to market, 247–249
 - implications for risk reporting, 249–250
 - model validation and, 241–250
 - risk management, 144–147, 150–151
- Illiquid positions, pitfalls in deriving valuations, 150–151
- Importance sampling, 185
- Incentive asymmetry, information asymmetry and, 9, 11
- Independent auditors, criticism of, 12–13, 79–80
- Indexed flows, 295–299
- described, 295–297
 - translation into fixed flows, 298–299
- Index options, 413–414
- IndyMac, 125
- Neichen, Alexander, 243
- Informationally disadvantaged, 21
- Information asymmetry:
- adverse selection and, 20–21
 - for creditors, 10
 - government regulation and, 11
 - incentive asymmetry and, 9, 11
 - moral hazard in, 7–16, 148–149
 - nature of, 8–9
 - outside monitors and, 10–16
 - potential solutions, 9–10
 - traders and, 148–149
- Initial margin, for exchange-traded derivatives, 508
- Insurers:
- AIG, 86, 93, 97, 102, 107, 108, 114
 - in financial crisis of 2007–2008, 96–97, 106–108, 114–115, 126
- Interest rate swaps, 80, 145, 275–276, 281
- Internal arbitrage, 167
- International Monetary Fund (IMF), 92, 93, 115
- International Swaps and Derivatives Association (ISDA), 38–39, 450, 467–468, 515–516, 530
- Interpolation approach:
- in building a volatility surface, 346–352
 - in model validation, 227, 230
 - pricing illiquid flows by interpolation, 284–291
 - seasonality of borrowing costs, 302
 - between strikes, 347–352
- Intraday margin calls, 509
- Investment analysts:
- conflict of interest, 12–13
 - information asymmetry and, 11–12
- Investment banks:
- capital requirements reform recommendation, 122–124
 - CDO creators in financial crisis of 2007–2008, 88–89, 111, 116–117
 - compensation reform recommendations, 120–122

- conflict of interest, 12–13
- failure to account for illiquidity of super-senior tranches, 101–102, 113
- faulty CDO models, 98–99, 112
- in financial crisis of 2007–2008, 93–106, 112–114, 118–126
- inadequate analysis of statistical hedging, 103–105, 113
- inadequate derivative protection, 96–97, 112
- inadequate stress tests, 102–103, 113
- losses in financial crisis of 2007–2008, 92, 93
- off-balance-sheet vehicles, 97–98, 112
- overreliance on VaR measures, 100–101, 113
- personnel risk and, 36–37
- recommendations for, 118–126
- reliance on external ratings, 99–100, 113
- risk management procedures reform recommendation, 119–120
- size and allowable activities reform recommendations, 124–126
- “too big to fail” mentality and, 11, 105–106, 114, 121, 124
- Investors, in financial crisis of 2007–2008, 92–93, 111, 118
- Irish central bank, 300
- iTraxx index, 493, 494

- Jackel, Peter, 175
- Jackwerth, Jens, 348
- Jacobs, Michael, 467, 469
- Jain, Gautam, 409
- Jameson, Rob, 29, 66
- Jett, Joseph, 51, 53–55, 57, 234–235
- Jewson, Stephen, 144
- Jorion, Philippe, 176, 191, 205, 206
- JPMorgan, 80, 105, 231
- JPMorgan Chase, 40, 79–80, 84, 106, 205–206, 249
- JumpProcessCredit spreadsheet, 475–476, 551

- Jump process models, 475–476
- Junk bonds, 307
- Jurek, Jakob, 500–501

- Kalotay, Egon, 467, 479
- Kane, Alex, 141
- Kani, Iraj, 385–391, 403–404
- Karagozoglu, Ahmet, 467
- Kashyap, Anil, 110
- Kealhofer, Stephen, 481, 484
- Keeney, Ralph, 6
- Kerviel, Jérôme, 61–66, 67
- Khuong-Huu, Philippe, 434
- Kidder Peabody case, 21, 31, 51, 53–55, 57, 64, 234–235
 - detection of unauthorized positions, 55
 - development of unauthorized positions, 54
 - failure to detect unauthorized positions, 54–55
 - further reading, 55
 - incident, 53–55
 - lessons to be learned, 55
 - result, 54
- Kim, Jongwoo, 178, 200
- King, Mervyn, 126
- Kirshner, Susan, 302
- KMV approach, 467, 476–484
- Knock-in (down and in/up and in), 382
- Knock-out (down and out/up and out), 382, 383
- Kolm, Petter, 176, 233
- Kooi, Mari, 66
- Kotowitz, Y., 7–8, 11
- Koutoulas, James, 511
- Krenn, Gerald, 200
- Kurer, Peter, 99–100

- Ladder options, 402–403
- Large complex financial institutions (LCFIs), 120
- Large homogenous portfolio (LHP), 487–490, 495

- Large money moves, 68–77. *See also*
 Long-Term Capital Management (LTCM) case
 Metallgesellschaft (MG) case, 75–77,
 135, 201–202, 273
 stress tests and, 195–196, 205
- Laurent, Jean-Paul, 490
- Law of one price, 239–241
- Lee, Roger, 386
- Lee, Yoolim, 80
- Leeson, Nick, 51, 55–57, 66
- Legal-basis risk, 38–39, 450
- Legal risk, 37–40
 bankruptcy and, 39–40
 defined, 30
 error in legal interpretation, 50
 mitigating, 37–40
 risk of illegal actions, 40
 of unenforceable contracts, 37–40
- Lehman Brothers, 455
- Leibowitz, Martin, 161
- Lender of last resort facilities, 129
- Leonhardt, David, 103
- Leverage:
 in capital structure analysis, 477
 as measure of default risk, 73
- Lewis, Michael, 102, 104
- Li, Ada, 128, 510
- Li, David, 98–99, 211, 488
- Li, Jingyi, 490
- LIBOR (London Interbank Offered Rate),
 69, 267–269, 275, 297–298, 447
- Limited partnerships, 80
- Linear combinations of asset prices,
 405–413
 approximation of option values,
 407–409
 derivative characteristics, 405
 derivative payoffs as linear functions
 of, 406–407
 risk management of options with,
 409–413
 rules for dynamic hedging, 407
- Lippmann, Greg, 104
- Liquid instruments, model risk and,
 237–241
- Liquidity risk. *See also* Asset liquidity
 risk; Funding liquidity risk
 costs of liquidation, 139
 defined, 3
 time required for liquidation,
 135–136
- Liquidity squeeze, 469
- Liquid proxy:
 control variate technique compared
 with, 360
 for derivatives with actuarial risk, 520
 for illiquid instruments, 144–146,
 243–245, 247–249
 reasons to use, 144–146
- Li's Gaussian copula formula, 98–99,
 211, 482, 489–490, 497–498
- Litterman, Robert, 203–204, 307
- Lo, Andrew, 84, 85, 131, 243
- Loan-equivalent approach, 513–515
- Loan-to-value ratios, 87
- Local volatility models, 385
- Log contracts, 367–369
- London Interbank Offered Rate
 (LIBOR), 69, 267–269, 275,
 297–298, 447
- Long, meanings of, 304
- Long-Term Capital Management
 (LTCM) case:
 bailout, 71–72
 large money moves, 68–77, 109, 122,
 160, 193, 195, 201–202, 206, 211,
 250, 318
 lessons learned, 73–75
 management style, 68–69, 73–74
 suggestions for improved practices,
 74–75
 types of positions, 69–70
 Union Bank of Switzerland (UBS)
 and, 59, 61, 318
- Lookback options, 402
- Loss given default (LGD), 447–448
 estimating, 464–468, 484–485, 487,
 491
- Lowenstein, Roger, 61, 72, 75, 81, 90,
 96
- Lubke, Theo, 128, 510

- Ludwig, Eugene, 51, 59
Lynch, Gary, 55
- Madan, Dilip, 364, 386
Madoff, Bernie, 17
Malcolm, Fraser, 37, 40
Marcus, Alan, 141
Margin Call (film), 1
Margin calls, 128, 509
Margining:
 exchange-traded derivatives, 506, 507, 508, 510, 511
 over-the-counter derivatives, 513, 514, 516–517, 518–520, 522–523, 531
- Mark, Robert, 67
Market contagion:
 in financial crisis of 2007–2008, 109–111, 115, 129–131
 need for broader regulatory oversight, 131
 need for more orderly bankruptcy proceedings, 131
 need to reduce procyclicality, 129–131
- Marketers, in front office, 8
Market makers/market making:
 gambling analogy, 26
 hedging in spot markets, 254–255
 impact of customer order flow in spot markets, 255–257
 liquidity risk/basis risk trade-off, 255–256
 market making, defined, 5
 models to perform risk decomposition, 25–26
 position taking versus, 24–27
 winner's curse and, 22
- Market risk, legal risk versus, 39
Market using. *See* Position taking
Marking to market, 95–96, 110, 130–131, 143, 147–152
 analysis of revenue and, 156
 caveats concerning, 150–151
 dollar versus Japanese yen, 147, 153
 in establishing exit prices, 148–149, 151–152
 for exchange-traded derivatives, 508
 by expert panels, 149–150
 exposure to market price shifts, 157–159
 frequency of, 147
 with illiquid positions, 146, 147–151
 liquid proxy for illiquid instruments, 144–146, 243–245, 247–248
 purpose of, 146
- Markowitz, Harry, 141
Martinuzzi, Elisa, 80
Matytsin, Andrew, 350, 386
Maurer, Samuel, 449, 450
Mayer, Martin, 55, 57, 67
MBIA, Inc., 107–108
McAdie, Robert, 455
McDonald, Robert, 260
McKay, Peter, 66
McLean, Bethany, 80, 90, 97, 106, 125, 523
McNeil, Alexander, 190
Mean reversion, 327–328
Mello, Antonio, 76–77
Merck, 407–408
Merger arbitrage, 26–27
Merrill Lynch, 67, 79, 94, 95, 106, 107–108
Merton model, 474–476, 551
MertonModel spreadsheet, 474, 477, 551
Metallgesellschaft (MG) case, 75–77, 135, 201–202, 273
Metropolitan Life, 4
Meucci, Attiolio, 176
Mezzanine tranches, 94, 102, 104–105, 494
MF Global, 511
Middle office, 10
 defined, 8
 fraud risk and, 32–35
 model verification and, 223
Mihm, Stephen, 126
Milken, Michael, 237
Miller, Merton, 76
Miller, William “520 Percent,” 18–19

- Misleading reporting, 49–67
- Allied Irish Bank (AIB) case, 31, 51, 57–59, 64, 65
 - Barings Bank case, 31, 51, 55–57, 64, 66
 - Chase Manhattan Bank/Drysdale Securities case, 45, 51, 52–53
 - deception about earnings, 31
 - deception about positions, 31–32, 49–51
 - Kidder Peabody case, 21, 31, 51, 53–55, 57, 64–66, 234–235
 - other cases, 66–67
 - risk of nondeliberate, 35–36
 - Société Générale case, 31, 61–66, 67
 - Union Bank of Switzerland (UBS) case, 59–61, 67
- Mixture of Normals spreadsheet, 5, 6, 175, 178, 207, 548
- Model risk, 209–252. *See also* Model risk evaluation and control
- defined, 209
 - illiquid instruments, 241–250
 - importance of, 210–212
 - liquid investments, 237–241
 - as operations risk, 35–36
 - trading models, 250–252
 - valuation of illiquid positions, 150
- Model risk evaluation and control, 212–237. *See also* Model risk board of directors role in, 219
- business unit accountability for, 215–219
 - capturing difficult-to-identify risks, 231
 - components of review, 214
 - continuous review, 232–234
 - documentation of, 218–219
 - model as term, 213–214
 - model validation, 212, 226–231
 - model verification, 212, 219–226
 - periodic review, 234–237
 - proprietary information and, 217–218, 250–252
 - roles and responsibilities, 214–219
 - scope, 213–214
 - senior management role in, 219
 - vender versus in-house models, 213–214
- Model validation, 226–231
- capturing difficult-to-identify risks, 231
 - choice of approach for illiquid instruments, 241–243
 - cost of hedging approach, 227–228
 - defined, 212, 226–227
 - illiquid instruments and, 241–250
 - interpolation approach, 227, 230
 - liquid instruments and, 237–241
 - matching to model purpose, 229–231
 - no-arbitrage principle and, 239–241
 - by outside reviewers, 238, 245–246, 250–252
 - prevailing market model approach, 228–229
 - of specific trading strategies, 230–231
- Model verification, 219–226
- of approximations, 223–226
 - components of, 220–221
 - of deal representation, 222–223
 - defined, 212
 - degree of complexity of models, 221–222
 - independent implementation, 220
 - model error and, 221–222
 - nature of, 219–220
 - rules, 220–221
 - suggested controls for computational approximation, 224–225
 - systems implementation, 220
 - testing on cases with known solutions, 220–221
- Money market mutual funds, in credit contagion of 2007–2008, 109
- Monopoly rents, 167
- Monte Carlo simulation:
- advantages of, 176–180, 245–246
 - computational alternatives to full simulation, 486–490

- of counterparty credit exposure, 517–518, 524–525
- disadvantage of, 199
- dynamic hedging of vanilla options, 321–329
- equal probability weights for all simulation runs, 189
- for illiquid instruments, 242–243, 245–247
- missing/nonsynchronous data in, 176–177, 181–182
- model verification using, 225–226
- of options hedging, 357–358
- of P&L distribution, 170–171, 175–183
- of portfolio credit risk, 482–486
 - with stress tests, 198–199, 200–201
 - stress tests versus, 192–193
- Monte Carlo stress tests, 200–201
- Moody's Investors Service ratings, 11, 89, 91, 108, 268, 459–464, 466, 480
- Moody's KMV, 467, 476–484
- Moosa, Imad, 47
- Moral hazard, 7–16
 - in analysis of insurance risks, 13–14
 - conflict between insiders and outsiders, 8–16
 - defining, 7–8
 - information asymmetry, 7–16, 148–149
 - in risk measurement, 135
 - taking large risk positions, 50–51
 - “too big to fail” mentality and, 11, 72, 105–106, 114, 121, 124
 - in value placed on earnings volatility, 15–16
- Morgan Grenfell Asset Management, 80
- Morgan Stanley, 4, 95
- Morgan Stanley Capital International (MSCI), 259
- Morini, Massimo, 209, 211, 212, 215, 220, 221, 226–227, 230–231, 235, 239, 471
- Morningstar, 258–259
- Mortgage brokers, in financial crisis of 2007–2008, 86–87
- Multiname credit derivatives, 493–501
 - CDO tranches and systematic risk, 500–501
 - modeling, 495–498
 - nature of, 493–495
 - risk modeling and reporting for, 498–500
- Myers, Stewart, 141
- Myktyka, Edward, 175
- Nagpal, Krishan, 461
- Narrow banks, 125
- NastyPath spreadsheet, 319, 345, 548
- National Association of Insurance Commissioners, 459
- National Westminster Bank, 67
- Netting:
 - exchange-traded derivatives, 506, 507, 511
 - over-the-counter derivatives, 513, 514, 515, 516, 517, 518, 520–521, 529, 531
- Neuberger, Anthony, 368
- New York Stock Exchange (NYSE), 277
- New York Times*, 87, 103
- No-arbitrage principle, 239–241
- Nobel Prize in economics, 20
- Nocera, Joe, 90, 97, 106, 125
- Nondiversifiable risk, 141–142, 197
- Nonlinear combinations of asset prices, 417–422
- Norris, Floyd, 107–108
- Northern Rock, 125
- Novation, exchange-traded derivatives, 506
- Numeraire, 305, 312
- O'Brien, Timothy L., 206
- Off-balance-sheet vehicles, 97–98
- Office of the Comptroller of the Currency, 87

- “Off-the-run” instruments, 238
- O’Kane, Dominic, 161, 455, 489, 490, 492–496, 500
- One-way markets, 151
- One-year tenor options, 432
- “On-the-run” instruments, 238
- Operational risk, 29–47
 - accounting risk, 30, 42
 - defined, 29
 - enterprise risk, 30, 44
 - funding liquidity risk, 30, 42–44
 - identification of risks, 44–45
 - legal risk, 30, 37–40
 - operational risk capital, 45–47
 - operations risk, 30, 31–37
 - reputational risk, 30, 41–42
- Operational risk capital, 45–47
 - bottom-up approach, 46
 - top-down approach, 46
- Operations risk, 31–37
 - defined, 30
 - disaster risk, 36
 - personnel risk, 36–37
 - risk of fraud, 31–35
 - risk of nondeliberate incorrect information, 35–36
- OptBarrier spreadsheet, 396, 440–441, 550
- Option-adjusted spread (OAS), 424
- OptionMC1000 spreadsheet, 358, 549
- OptionMCHedged1000 spreadsheet, 358, 549
- OptionMCHedged spreadsheet, 357–358, 549
- OptionMC spreadsheet, 357–358, 549
- OptionRoll spreadsheet, 352–355, 549
- Options risk management. *See also* Exotic option risk management; Vanilla option risk management
 - options conventions, 311–312, 426–427
 - options transactions, defined, 253
 - overview of options risk management, 313–318
- Options to exchange one asset for another, 415–417
- Option-theoretic approach, 471–479
 - jump process models, 475–476
 - KMV statistical analysis, 476–484
- Out-of-the-money calls, 319
- Overbeck, Ludger, 490
- Override analysis, 233–234
- Over-the-counter derivatives, 128–129, 508–509, 512–531
 - active management report, 526–531
 - closeout, 513, 515, 516, 517, 519, 520, 529
 - collateralization approach, 515–526
 - counterparty risk groups (CRGs), 526–531
 - loan-equivalent approach, 513–515
 - margining, 513, 514, 516–517, 518–520, 522–523, 531
 - netting, 513, 514, 515, 516, 517, 518, 520–521, 529, 531
 - overview, 512–513
 - wrong-way risk, 521–526
- Oyama, Tsuyoshi, 85
- Padovani, Otello, 404
- Paine Webber, 54
- Pandit, Vikrim, 96–97
- Parking, 31
- Parsons, John, 76–77
- Partial differential equations (PDEs), 383–384
- Partial-time barrier options, 382
- Path-dependent options, 17–18, 381–404
 - barrier options with rebates, 402
 - broader classes, 403–404
 - deriving the Carr hedge, 393
 - described, 362
 - dynamic hedging models for barriers, 385–387
 - in exotic option risk management, 381–404
 - intensity of use, 363
 - ladder options, 402–403
 - lookback options, 402
 - put-call symmetry, 391–392

- standard analytic models for barriers, 383–384
- static hedging models for barriers, 387–401
- in vanilla option risk management, 318–323
- Paulson, John, 104
- Pearson, Neal, 496
- Pension funds, as investors in financial crisis of 2007–2008, 92–93, 111, 118
- Performance attribution, 258–259
- Performance measurement, 205–206
- Periodic review, 234–237
 - changes in academic literature, 237
 - changes in market environment, 236–237
 - changes in market practices, 237
 - changes in population of transactions, 234–236
 - changes in technology, 237
- Perold, Andre, 16, 69, 75
- Personnel risk, 36–37
- Phantom profits, 167
- Physical commodities:
 - borrowing costs for, 303–304
 - defined, 254
 - financial commodities versus, 254
 - spot risk, 259–260
 - storage costs, 301, 302
 - transportation costs, 259–260
- Pindyck, Robert, 4
- Pin risk (Taleb), 377
- Pirrong, Craig, 128
- Plain-vanilla options. *See* Vanilla option risk management
- Ponzi, Charles, 18–19
- Ponzi schemes, 17–19, 156
 - broadened meaning, 17–18
 - hedge slippage and, 17–18, 232
 - Kidder Peabody case, 21, 31, 51, 53–55, 57, 64–66, 234–235
 - losses from unauthorized positions and, 51
 - original meaning, 17, 18–19
- Portfolio credit risk, 479–493, 515
 - computational alternatives to full simulation, 486–490
 - estimating default correlations, 479–482
 - Monte Carlo simulation of, 482–486
 - risk management and reporting for portfolio credit exposures, 490–492
- Portfolio insurance, 96–97, 315, 320
- Portfolio Risk Tracker, 484–485
- Portfolio theory, 141
- Position managers, in front office, 8
- Position taking:
 - defined, 25
 - gambling analogy, 26
 - instruments outside area of expertise, 164–165
 - market making versus, 24–27
 - models as forecasting tools, 25
 - risk measurement for, 159–161
- Power options, 366–367
- Predescu, Mirela, 455
- Price taking. *See* Position taking
- Price-vol matrix:
 - advantage of, 339
 - for being a short a call option, 331, 332
 - for a calendar spread, 336, 337
 - for a call spread, 331–334
 - interpolation results based on, 185
 - for a reduced risk portfolio, 336–341
 - in vanilla option risk management, 315–317, 323–324, 326, 329–344
- PriceVolMatrixCycle spreadsheet, 336, 548
- PriceVolMatrix spreadsheet, 329–344, 548
- Pricewaterhouse Coopers, 101, 122, 127, 129–131
- Prince, Chuck, 94
- Private equity funds, need for broader regulatory oversight, 131
- Procter & Gamble (P&G), 41, 77–78, 367

- Program trading, 251
- Proprietary trading, 125–126
- Prudential-Bache Securities, 80
- Pull to par, 429
- Pyramid schemes, 17. *See also* Ponzi schemes
- Quanto:
- nonlinear combinations of asset prices, 417–422
 - single-asset quanto options, 369–370
- Quanto worksheet, 370
- Rafael, Andrea, 201
- Raiffa, Howard, 6
- Rainbow contracts, 419
- Rajan, Raghuram, 122, 125, 126, 499
- Ramberg, John, 175
- Random matrix theory/shrinkage estimation, 176
- RateData spreadsheet, 308–309, 548
- Rates spreadsheet, 283–284, 288, 307, 310, 548
- Ratios worksheet, 138
- Rawnsley, Judith, 57
- Real options, 4
- Rebates, barrier options with, 402
- Rebonato, Riccardo, 6, 209, 228, 237, 241, 242, 327, 348, 434–435, 438
- Rebooking trades, 152–153
- Reduced risk portfolio, 336–341
- Rehedging, 196, 322, 327–329
- Reiner, Eric, 417
- Remargin period, 519
- Remolona, Eli, 84, 465
- Renault, Olivier, 459, 461–462, 465–467, 476, 478, 484
- Rennie, Andrew, 418
- Repurchase agreements (RPs), 69, 271–272, 280, 294–295
- Reputational risk, 41–42, 77–81
- accounting risk as form of, 42
 - Bankers Trust (BT) case, 41, 77–79, 367
 - defined, 30
 - large money moves and, 196
 - nature of, 41–42
- Researchers, in front office, 8
- Reserves. *See* Valuation reserves
- Resti, Andrea, 465, 467
- Revealing positions, problems of, 150–151
- Richardson, Matthew, 84, 118, 121
- Right-way risk, 522
- Risk-adjusted return on capital (RAROC), 44, 206
- Risk aggregation, 4
- Risk arbitrage. *See* Merger arbitrage
- Risk control, 161–167
- detailed limits on size of exposure, 162–165
 - incentive-based approach to, 161–163
 - internal hedging in, 166–167
 - risk decomposition and, 166
- Risk decomposition:
- defined, 4
 - models to perform, 25–26
 - reporting in, 203–204
 - risk control and, 166
- Risk Identification for Large Exposures (RIFLE), 231
- Risk* magazine, 99
- Risk management. *See* Financial risk management
- Risk Management Association, 485
- Risk managers, in front office, 8
- Risk measurement, 133–161
- analysis of revenue, 156–157
 - exposure to changes in market prices, 157–159
 - general principles, 133–144
 - instruments that lack liquidity, 144–147, 150–151
 - liquidation time and, 135–136
 - market valuation, 147–152
 - for position taking, 159–161
 - principles of risk management in, 136–142
 - rules for, 133–134

- stop-loss limit, 133–136
- valuation reserves and, 145, 146, 152–156
- RiskMetrics Group, 177
- Risk of fraud, 31–35
 - pressures, 32–35
- Risk of nondeliberate incorrect information, 35–36
- Risk reversals, 334–335, 377
- Roe, John, 511
- Roseman, Alan, 107
- Rosen, Dan, 485
- Ross, Stephen, 141
- Roubini, Nouriel, 126
- Royal Bank of Scotland, 67
- Rubinstein, Mark, 348
- Rullière, Didier, 482
- Runyon, Damon, 44
- Rusnak, John, 51, 57–59
- Russian debt default of 1998, 71, 206

- Salespeople, in front office, 8
- Salmon, Felix, 98, 209, 211
- Salomon Brothers, 68, 71
- Sanders, Anthony, 424
- Sarkar, Asani, 449, 450
- Saunders, Anthony, 120, 464, 481
- Scenario analysis, 211
- Schachter, Barry, 190, 191, 194
- Scheinkman, Jose, 307
- Scheuermann, Til, 190
- Schonbucher, Philipp, 99, 475, 483, 487–489
- Schorderet, Yann, 161
- Schuermann, Til, 85, 90
- Schutz, Dirk, 61
- Seasonality, of borrowing costs, 302
- Securities and Exchange Commission (SEC), 12, 153, 159
- Securities Industry Association, 210
- Seinfeld* (TV program), 1
- Sell side. *See* Market makers/market making
- Semi-American options, 426
- Semi-European options, 426
- Senior Supervisors Group report, 100

- Senior tranches, 494
- September 11, 2001 attacks, disaster risk and, 36
- Serrat, Angel, 268, 284, 298–299, 304–305, 308
- Shakespeare, William, 142–143
- Shareholders:
 - information asymmetry and, 11–12
 - outside monitors for, 11–12
- Shareholder value added (SVA), 44, 206
- Sharma, Pawan, 37, 40
- Sharpe, William, 141, 258–259
- Sharpe ratio, 5, 160
- Shaw, Julian, 175, 182
- Shiller, Robert, 103, 116, 132
- Shin, Hyun Song, 110
- Shirreff, David, 61, 72
- Shkolnik, Alexander, 490
- Short, meanings of, 304
- Shortfall/expected shortfall VaR, 187–189
- Shorting a call option, 331, 332
- Short squeeze, 300
- Short-term credit exposure, 451–456
 - CDS-bond basis risk, 454–456
 - convexity of credit instruments, 453–454
 - impact of bankruptcy law, 452–453
 - risk reporting for market credit exposures, 456–457
- Sidenius, Jakob, 490
- Sifakis, Carl, 18–19
- Simulation. *See also* Monte Carlo simulation
 - advantages of, 139–140
 - computational alternatives to full simulation, 486–490
 - historical data in, 5, 170–171, 173–174, 180–183
 - illiquid positions in, 146
 - Monte Carlo (*see* Monte Carlo simulation)
 - nature of, 138–139
 - of P&L distribution, 170–171, 173–187

- Simulation (*Continued*)
for risk measurement, 138–140
subjective judgment and, 4–6
- Single-asset quanto options, 369–370
- Single-name credit risk, 457–479
estimating amount owed at default, 468–471
estimating loss given default, 465–468
estimating probability of default, 458–465
option-theoretic approach, 471–479
- Single-payout options, 364–378
accrual swaps, 378
binary options, 371–377
contingent premium options, 377–378
convexity, 370
described, 362
intensity of use, 363
log contract swaps, 367–369
single-asset quanto options, 369–370
variance swaps, 367–369
- Singleton, Kenneth, 454, 455, 482, 483
- Sironi, Andrea, 465, 467
- Skew, 349–350
- Smiles, 348–349
- Smith, Adam, 11
- Smith, Roy, 120
- Smithson, Charles, 362, 363, 496
- Société Générale case, 31, 61–66, 67
detection of unauthorized positions, 64
development of unauthorized positions, 62
failure to detect unauthorized positions, 62–64
further reading, 66
incident, 61
lessons to be learned, 64–66
result, 62
- Sony Corporation, 417
- Sorkin, Andrew Ross, 109
- Soros, George, 116
- South Korea, 80
- SpecComm, 62–63, 63
- Speculation. *See* Position taking
- Spence, Michael, 20
- Split-fee options, 379–381
- Spot risk management, 253–261, 279
equity, 258–259
firm-level risk management, 257
foreign exchange, 257–258
overview, 253–257
physical commodities, 259–260
spot trades, defined, 253
- Spreadsheets:
AmericanOption spreadsheet, 427–428, 551
BasketHedge spreadsheet, 364, 368–370, 381, 412–413, 439, 549–550
BasketOption spreadsheet, 550
BinaryMC spreadsheet, 376, 440, 550
Bootstrap spreadsheet, 289, 548
calculating default rates from bond rates, 501–502
CapFit spreadsheet, 347, 549
CarrBarrierMC spreadsheet, 396, 440, 550
CarrBarrier spreadsheet, 394–396, 440, 550
CDO spreadsheet, 496, 551
comparing the jump process credit model to the Merton model, 502
CreditPricer spreadsheet, 453, 551
CrossHedge spreadsheet, 420, 441–442, 527–531, 550–551
DataMetricsRatesData spreadsheet, 431
DermanErgenerKani20 spreadsheet, 550
DermanErgenerKaniDoubleBarrier spreadsheet, 403–404, 441, 550
DermanErgenerKaniPartialBarrier spreadsheet, 403–404, 550
DermanErgenerKani spreadsheet, 550
EVT spreadsheet, 548
ForwardStartOption spreadsheet, 381, 550

- ForwardStart spreadsheet, 379
- generating fat tails in Monte Carlo simulations, 207
- interpolation, 308–309
- JumpProcessCredit spreadsheet, 475–476, 551
- maximizing diversification, 207
- measuring fat tails in historical data, 207
- MertonModel spreadsheet, 474, 477, 551
- MixtureofNormals spreadsheet, 5, 6, 175, 178, 207, 548
- Monte Carlo simulation of options hedging, 357–358
- NastyPath spreadsheet, 319, 345, 548
- OptBarrier spreadsheet, 396, 440–441, 550
- OptionMC1000 spreadsheet, 358, 549
- OptionMCHedged1000 spreadsheet, 358, 549
- OptionMCHedged spreadsheet, 357–358, 549
- OptionMC spreadsheet, 357–358, 549
- OptionRoll spreadsheet, 352–355, 549
- options portfolio risk measures, 356–357
- PriceVolMatrixCycle spreadsheet, 336, 548
- PriceVolMatrix spreadsheet, 329–344, 548
- RateData spreadsheet, 308–309, 548
- Rates spreadsheet, 283–284, 288, 307, 310, 548
- simulation of the impact of trading rules on expected return and risk, 260–261
- stack and roll, 309
- Swaptions spreadsheet, 437, 442, 551
- TermStructure spreadsheet, 432–433, 442, 551
- using Vasicek model for risk
 - measurement of CDO tranches, 502–503
 - value-at-risk computations, 206–207
- VaR spreadsheet, 138, 181, 189, 206–207 (example), 548
- VolCurve spreadsheet, 347, 549
- VolSurfaceStrike spreadsheet, 348, 352, 549
- WinnersCurse spreadsheet, 24, 548
- Squam Lake Group, 118, 119, 121, 127, 131
- Stack-and-roll hedge, 291–293
 - advantages of, 293
 - described, 291–293
- Stafford, Erik, 500–501
- Standard & Poor's (S&P) 500 stock index, 3–4, 141, 157–159, 259, 319, 371–377
- Standard & Poor's (S&P) ratings, 11, 89, 108
- Static hedging strategies:
 - for barrier options, 387–401
 - for exotic options, 361–362
 - flows representing promised deliveries, 293–295
 - indexed flows, 295–299
 - nature of, 313
 - pricing illiquid flows by interpolation, 284–291
 - quasistatic representations, 361–362
 - stack-and-roll hedge, 291–293
- Static overhedge, 375
- Statistical hedging, inadequate analysis in financial crisis of 2007–2008, 103–105
- Stay period, 519
- Stein, Roger, 464–467, 475, 479, 482, 483, 492
- Sticky delta, 342
- Sticky strike, 342
- Stiglitz, Joseph, 20, 126
- Stigum, Marcia, 51, 53
- Stochastic volatility models, 385
- Stop-loss limits, 133–136, 162–163
- Storage costs, 301, 302

- Stress tests, 1, 192–201
in assessing credit risk, 75
in capital requirements reform recommendations, 122–124
of counterparty credit exposure, 520
economic scenario stress tests, 192, 193–197
for exchange-traded derivatives, 509–510, 511
factor-push, 199–200
historical data stress tests, 192, 197–201
impact of large money moves and, 205
inadequate, 102–103, 113
large money moves and, 195–196, 205
Monte Carlo simulation versus, 192–193
Monte Carlo simulation with, 198–199, 200–201
overall measures of firm position risk, 201–205
overview, 192–193
performance measurement and, 205–206
for positions that achieve liquidity, 143
- Strickland, Chris, 386, 426
- Stroughair, John, 190
- Structured Finance Litigation blog, 92
- Structured investment vehicles (SIVs), 97–98
- Structurers, in front office, 8
- Subjective judgment. *See also* Stress tests
historical information versus, 5
in Li's Gaussian copula formula, 98–99
simulation and, 4–6, 140–141
- Subprime mortgage originators. *See also* Financial crisis of 2007–2008
in financial crisis of 2007–2008, 86–88, 111, 116
- Sullivan, Arthur, 14, 303
- Sumitomo Corporation of Japan, 66
- Suo, Wulin, 229, 455
- Super-senior tranches, 94–102, 104–105, 106–107, 113, 494
- Swaps:
accrual, 378
basis, 298–299
binary credit default, 449–450
credit default (*see* Credit default swaps [CDS])
cross-currency, 525
interest rate, 80, 145, 275–276, 281
log contract, 367–369
total return, 276–278, 450–451
variance, 367–369
volatility, 368–369
- Swaptions, 413–414
Bermudan, 432, 433–434
European, 430
relationships between cap prices and, 437–439
- Swaptions spreadsheet, 437, 442, 551
- Swensen, David, 176
- Swiss Bank Corporation (SBC), 60
- Synthetic tranches, 494
- Systematic/nondiversifiable risk, 141–142, 197
- Tadikamalla, Pandu, 175
- Taleb, Nassim, 14–15, 138, 300, 318–319, 334, 342–344, 347, 350, 355, 377
- Tanega, Joseph, 37, 40
- Technologists, in front office, 8
- Technology stock bubble (2001), 80–81
- Telecom, 263
- Term structure models, 430–436
- TermStructure spreadsheet, 432–433, 442, 551
- Tett, Gillian, 88, 91, 95, 97, 99–100, 102, 105
- Thaler, Richard, 21–22
- Theta:
defined, 343–344
price-vol-matrix and, 343–344

- Tickets in the drawer, 31
- Time-dependent options, 378–381
- cliquet options, 378–379
 - compound options, 379–381
 - described, 362
 - forward-start options, 378–379
 - intensity of use, 363
- “Too big to fail” mentality, 11, 72, 105–106, 114, 121, 124
- Total return swaps, 276–278, 450–451
- Totem Market Valuations service, 150
- Trade cancellation, 62–63, 65
- Trade compression, 530
- Traders:
- adverse selection and, 20–21
 - collusion and, 63, 65
 - conservatism versus independence and, 155–156
 - control personnel versus, 148–149
 - delta rehedging and, 196
 - detailed limits on size of exposure, 162–165
 - fraud risk and, 32–35
 - in front office, 8
 - G-30 recommendations on trading risk, 136, 137, 147, 156–157, 159, 169
 - incentive-based approaches in risk control, 161–163
 - information asymmetry and, 7–16, 148–149
 - monopoly rents and, 167
 - moral hazard and, 14–15, 148–149
 - positions in instruments outside area of expertise, 164–165
 - pressure to book immediate profits, 240–241
 - trading models and, 246–247, 250–252
 - valuation reserves and, 152–156
- Trading and Capital-Markets Activities Manual* (Federal Reserve System), 30
- Trading models, 250–252
- Transportation costs, in physical commodities spot risk, 259–260
- Treasury function, funding liquidity risk control and, 43–44
- Trinomial tree model, 425–426
- Tsiveriotis, Kostas, 416
- Tuckman, Bruce, 268, 284, 298–299, 304–305, 308
- Turner Review, 114–116, 120, 121, 124–125
- Twelfth Night* (Shakespeare), 142–143
- UniCredit Group, 63
- Union Bank of Switzerland (UBS):
- Amplified Mortgage Portfolio (AMPS), 103–104
 - analysis of financial crisis of 2007–2008, 95, 99–101, 103–104
 - VaR methodologies, 99–101, 104
- Union Bank of Switzerland (UBS) case, 59–61, 67, 335
- development of authorized positions, 60–61
 - further reading, 61
 - incident, 59
 - lessons learned, 61
 - result, 60
- Up and in (knock-in), 382
- Up and out (knock-out), 382, 383
- Utopia, Limited* (Gilbert & Sullivan), 14, 303
- Vacation policy, 63, 65
- Valuation reserves, 145, 146, 152–156
- aging reserve policy, 155
 - impact of exiting large positions, 154
 - model verification and, 226
 - objective standards for reserves, 153–156
 - to shield earnings from fluctuation, 248–249
- Value at risk (VaR) analysis, 1, 104, 136, 169–191
- back-testing, 191, 233
 - based on credit rating agencies, 99–100
 - based on historical variance/covariance, 170, 172–173

- Value at risk (VaR) analysis (*Continued*)
- in capital requirements reform recommendations, 122–124
 - counterparty credit exposure, 517–520
 - detail recorded on positions and market prices, 185–186
 - determining all market variables, 183–187
 - direct measurement of profit and loss, 170, 171–172
 - earnings volatility and, 205
 - for exchange-traded derivatives, 508, 509–510, 511
 - exotic derivative prices and, 186
 - extreme value theory (EVT) in, 190–191
 - in financial crisis of 2007–2008, 99–101, 172–173
 - for forward positions, 184
 - illiquid positions in, 146
 - importance sampling in, 185
 - liquidity considerations in, 146, 186–187
 - measures of profit and loss distribution, 187–191
 - nonstatistical measures versus, 169
 - for option positions, 184–186
 - overall measures of firm position risk, 201–205, 201–206
 - overreliance on, 100–101, 113
 - performance measurement and, 205–206
 - for positions that are born illiquid, 143
 - in risk control, 162
 - shortfall/expected shortfall, 188–189
 - simulations of P&L, 170–171, 173–187
 - for spot positions, 183–184
- Vanilla option risk management, 311–358
- building a volatility surface, 346–355
 - conventions, 311–312
 - delta hedging, 315–316, 320, 344–345
 - dynamic hedging strategies, 314, 318–329
 - overview, 313–318
 - risk reporting and limits, 329–344
 - tools in, 315–316
 - vanilla call spread, 372
 - vanilla options, defined, 311, 426
- van Nieuwerburgh, Stijn, 84
- VaR analysis. *See* Value at risk (VaR) analysis
- Varian, Hal, 406
- Variance gamma model, 386
- Variance swaps, 367–369
- VaR spreadsheet, 138, 181, 189, 206–207, 548
- Vasicek model, 98, 102, 487–489, 496, 501
- Vause, Nicholas, 530
- Vega:
 - defined, 315, 330
 - price-vol-matrix and, 330, 334–335
- Vigorish, 26
- Volatility surface, 346–355
 - extrapolating based on time period, 352–355
 - interpolating between strikes, 347–352
 - interpolation between time periods, 346–347
 - for pricing vanilla options, 346
- Volatility swaps, 368–369
- Volcker, Paul, 125, 132
- Volcker rule, 125
- VolCurve spreadsheet, 347, 549
- VolSurfaceStrike spreadsheet, 348, 352, 549
- Vosey Inheritance, The* (Granville-Barker), 18–19
- Wagner, Christoph, 490
- Wall Street Journal*, 21, 40, 116
- Walter, Ingo, 120, 121

- Wang, Jin, 175
Wang, Yuan, 449, 450
Washington Mutual, 84, 125
Wealth of Nations, The (Smith), 11
Weather derivative options, 144
Weinberger, Alfred, 161
Weiss, Gary, 55
Whaley, Elizabeth, 344
White, Alan, 409, 455, 490, 495
White, Lawrence, 84, 118
Williams, Jeffrey, 265
Williams, Meredith, 83, 93
Wilmott, Paul, 344
Wilson, Charles, 20
Wilson, Harry, 67
Wilson, Thomas, 201–202, 205, 484, 485
Winner's curse, 21–24
 application to trading, 22–24
 defined, 21
 mechanism leading to, 21–24
WinnersCurse spreadsheet, 24, 548
Winters, Bill, 526
Wired magazine, 98
Wolfe, Eric, 67
Wolfe, Lan-Ling, 424
World Bank, 115
Wrong-way risk, 521–526

Y2K crisis, 36, 192, 302
Yield curve, nonstatistical limits on
 yield curve shape, 307–308
Youngblood, Michael, 87

Zandi, Mark, 94, 130–131
Ziehmann, Christine, 144
Zou, Joseph, 409
Z-score model, 464–465