The McGraw-Hill Companies

isk mana

Jayanth Rama Varma

Derivatives and Risk Management

The McGraw·Hill Companies

Derivatives and Risk Management

About the Author

JAYANTH VARMA is a Professor in the Finance and Accounting Area at the Indian Institute of Management, Ahmedabad. He has carried out extensive research in the field of Indian financial markets and finance theory and published extensively.

Prof. Varma has substantial experience relating to the regulation of derivative markets in India. He has served as the Chairman of the Risk Management Group of the Forward Markets Commission which regulates commodity futures in India. He was also the Chairman of the Advisory Committee on Derivatives and of the Group on Secondary Market Risk Management set up by the Securities and Exchange Board of India (SEBI) which regulates equity and equity derivative markets. He has also been a member of the board of SEBI for four years.

Prof. Varma advises companies and financial institutions on issues related to risk management and finance and has also served on the board of several companies and banks.

Prof. Varma did his post-graduation in management from the Indian Institute of Management, Ahmedabad (IIMA), where he was awarded a Gold Medal for scholastic performance. Subsequently, he obtained his doctorate in management from IIMA. He is also a qualified cost accountant.

Derivatives and Risk Management

Jayanth Rama Varma

Indian Institute of Management Ahmedabad



Tata McGraw-Hill Publishing Company Limited

NEW DELHI

McGraw-Hill Offices

New Delhi New York St Louis San Francisco Auckland Bogotá Caracas Kuala Lumpur Lisbon London Madrid Mexico City Milan Montreal San Juan Santiago Singapore Sydney Tokyo Toronto



Published by the Tata McGraw-Hill Publishing Company Limited, 7 West Patel Nagar, New Delhi 110 008.

Copyright © 2008, by Tata McGraw-Hill Publishing Company Limited,

No part of this publication may be reproduced or distributed in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise or stored in a database or retrieval system without the prior written permission of the publishers. The program listings (if any) may be entered, stored and executed in a computer system, but they may not be reproduced for publication.

This edition can be exported from India only by the publishers, Tata McGraw-Hill Publishing Company Limited.

ISBN (13 digits): 978-0-07-060430-8 ISBN (10 digits): 0-07-060430-4

Managing Director: *Ajay Shukla* General Manager—Publishing: (B&E/HSSL and School): *V Biju Kumar* Editorial Manager—B&E: *Tapas K Maji* Junior Editorial Executive: *Hemant K Jha* Executive—Editorial Services: *Anubha Srivastava* Junior Manager—Production: *Medha Arora* General Manager Marketing—Higher Education & School: *Michael J. Cruz* Asst. Product Manager: *Vijay Sarathi Jagannathan* Production Controller: *Rajender P Ghansela* Asst. General Manager—Production: *B L Dogra*

Information contained in this work has been obtained by Tata McGraw-Hill, from sources believed to be reliable. However, neither Tata McGraw-Hill nor its authors guarantee the accuracy or completeness of any information published herein, and neither Tata McGraw-Hill nor its authors shall be responsible for any errors, omissions, or damages arising out of use of this information. This work is published with the understanding that Tata McGraw-Hill and its authors are supplying information but are not attempting to render engineering or other professional services. If such services are required, the assistance of an appropriate professional should be sought.

Typeset at Script Makers, 19, A1-B, DDA Market, Paschim Vihar, New Delhi 110 063, and printed at Lalit Offset, 219, FIE, Patparganj Industrial Area, Delhi 110092

Cover Design: K Anoop

Cover printed at: SDR Printers

RZXZCRCXDDRBX

The McGraw·Hill Companies

Preface

Why this book, and for whom?

I have been looking for a book which takes a holistic view of derivatives and risk management, takes a modern approach to the subject, discusses the traditional as well as emerging models, focuses on intuitive analysis but not at the expense of essential math, and gives a succinct appraisal of the derivatives environment in India and the world. Though there are many excellent books on derivatives and each has its own strengths, none was exactly what I wanted. After using several good books over the last two decades, I finally decided to write a book that approaches the subject my way. This book has, therefore, evolved out of my experience of teaching postgraduate and doctoral students of managerial finance as well as finance executives.

Derivatives and Risk Management is a much sought after course by the MBA students today. This is true not only in India but worldwide. Typically, the course is offered as a finance specialisation optional in the final semester or year of an MBA course.

As a future fund and investment manager, the MBA Finance student must have a good idea of the risks which his or her firm would be exposed to and how to use derivatives creatively and analytically to mitigate those risks. A variety of derivatives have flooded the Indian stock markets since the turn of the 21st century as a result of the ongoing globalisation phenomenon and the opening up of economies and financial markets. The US has been traditionally leading the scene with a very high share of the derivatives trading, but the gap between the US and the rest of the world has been closing steadily. Many emerging markets including India have thriving derivative markets today.

What sets this text apart?

Approach

Right mathematical rigour: A few years ago, I realized that many chapters in advanced texts on asset pricing that we use in doctoral courses are often mathematically simpler than corresponding chapters in the derivatives texts designed for MBA courses. The reason is that in doctoral courses, we start with risk neutral valuation as the key tenet of modern finance and relate all pricing models to this central idea. Historically, however, option pricing formulas were first discovered using the idea of dynamic hedging and replication in continuous time. Continuous time finance requires mathematical concepts and techniques (like diffusion processes and Ito's lemma) that are quite intimidating to most MBA students. Risk neutral valuation on the other hand involves nothing more complex than computing weighted averages and discounting cash flows to find their present value.

Risk neutral valuation: This book is based on the idea that by using risk neutral valuation as the central theme, it is possible to understand a great deal about option valuation without knowing too much about diffusion processes and Ito's lemma. The reader, however, has to accept the validity of the principle of risk neutral valuation. Books for doctoral students spend a lot of time and effort proving this principle from basic assumptions using a lot of continuous time finance. This book, since it is designed for MBA

vi | Preface

students, does not try to do that. It devotes a whole chapter to providing an intuitive explanation about why risk neutral valuation is correct, but avoids a rigorous proof of the principle. The result is an enormous reduction in the mathematics required for reading this book. In particular, knowledge of stochastic calculus is not required to enjoy this book.

Integration of derivatives with risk management—the case studies: This text gives equal importance to derivatives and to risk management. Risk management is a much broader field than derivatives and involves several different perspectives that are strategic in nature. Financial risk management in the modern business corporation requires consideration of a very rich set of issues that are best understood using case studies. This book therefore includes several cases about risk management and derivatives.

"Global-local" context: In the context of the increasing global openness of Indian financial markets, this book includes an adequate coverage of the world's leading derivative markets while also providing material specific to the Indian derivative markets. Though many examples are from Indian markets, the text does not hesitate to use examples from other countries especially for instruments which are less developed in India. Several of the case studies are also from outside India.

Text organisation

I have conceptualised the coverage of DRM in 22 chapters, each dwelling on a particular topical theme. These topics range from derivative instruments, derivatives markets and market dynamics to models, methodologies and strategies pertaining to derivatives as well as risk management. Although many concepts and topics in the subject are interrelated, I have tried to make each chapter self-sufficient. This lends a wholesome and well-rounded feel to each chapter.

The introductory chapter looks at the evolution, history and the taxonomy of the derivatives, as also at the dynamics of the derivatives markets and trading mechanisms. Chapters 2 to 5 deal primarily with forward and futures. Much of the discussion in these chapters dwells on hedging. A series of case studies are included after this chapter to get a deeper understanding of corporate hedging decisions.

The bulk of this book (Chapter 6 to Chapter 17) is devoted to options. One highlight of this crosssection of chapters is the methodology of risk neutral valuation. Chapter 8 discusses this methodology, which is used in the rest of the book to understand derivative valuation. This text uses the risk neutral approach consistently and this chapter is therefore critical for understanding the rest of the book. Another highlight is Chapter 10 which discusses the famous Black-Scholes model: the most famous formula in derivative valuation. Two case studies on convertibles included at the end of Chapter 17 serve to illustrate how the models described in this book can be used to analyse and evaluate complex real world financial instruments.

Chapters 18 and 19 discuss some of the basic fixed income or interest rate derivatives. This is a complex subject and specialized books have been written exclusively on interest rate derivatives. This book deals with only the elementary theory of these derivatives. One case study on Rupee-Dollar swap is included in Chapter 18 to reinforce the concept of currency swaps.

Chapter 20 presents an overview of derivative accounting which is important for most users of derivatives. The chapter includes a case that highlights some of the complexities involved.

Chapters 21 and 22 deal with risk management. Chapter 21 which is a continuation of Chapter 5 discusses corporate risk management. The chapter includes a comprehensive case that explores all aspects of risk management and the design of a complex hedging strategy. Chapter 22 is about risk management in financial institutions. This chapter is largely about value at risk (VaR) and stress testing. Two case

studies arising from uncertainties in demand and supply of a commodity are included in Chapter 21 for exposing the students to the peculiarities of corporate risk management.

Learning aids

The book contains a large number of illustrations in the chapters as well as exercises at the end of each chapter. The exercises are designed to provide a mechanism to the students to test their understanding of the concepts and frameworks discussed in the chapters.

The text includes case studies at the end of each chapter, based on real data and events from the financial sector, to help readers interpret theory in the context of practice in the world of finance.

The learning objectives at the beginning of the chapters clearly delineate the scope of learning in the chapter. The chapter summaries at the end of the chapters reinforce these learning objectives which the student would have learnt by studying the chapter.

A textbook cannot be complete without explaining the derivation of the key theoretical results. I have tried to segregate the relatively difficult mathematical derivations in separate optional sections without affecting the continuity and flow of the textual matter. Practitioners who are more interested in the concepts and their application rather than in the derivation would thus be able to skip the optional sections and enjoy the rest of the book.

JAYANTH RAMA VARMA

Acknowledgements

I am grateful to the Indian Institute of Management, Ahmedabad (IIMA) for allowing me to include a number of cases whose copyright vests with the Institute. I am also grateful to the co-authors of some of these cases—Prof. Samir Barua, Prof. V. Raghunathan and Prof. N. Venkiteswaran—for allowing these cases to be used.

I would also like to acknowledge the comments and suggestions that I received from Prof. Barua on several chapters of this book at an early stage when we planned to write this book together.

I have also benefited from the experience of using some of the ideas and material in this book while teaching courses related to financial markets to several batches of students at IIMA and also to participants in several executive education programmes.

I gratefully acknowledge the contribution of reviewers during the developmental stage of the book. I thank all of them for their constructive criticism. The book has benefited from their comments and suggestions. The reviewers are Prof Pitabas Mohanty, XLRI, Jamshedpur, Dr G David Raju, Andhra Loyola PG College, Vijaywada, Mr J N Mukhopadhyay, Globsyn Business School, Kolkata, Dr G Jayabal, Alagappa Institute of Management, Karaikudi and Mr Sandeep Verma, Sachdeva Institute of Technology, Mathura.

I have learnt a lot from innumerable books and papers on derivatives over the course of several years, and this book doubtless draws on ideas picked up from these sources. It is futile to even attempt to list all of them. Instead, at the end of each chapter, I have given a very selective list of suggestions for further reading. This is intended not as an exhaustive bibliography, but as an introduction to the vast literature that awaits the interested reader.

Needless to say, I am solely responsible for any errors in this book.

The McGraw·Hill Companies

Contents

Prefe Ackn	ace lowled	gements	v ix
1	INTI	RODUCTION TO DERIVATIVES	11
1.	1.1 1.2 1.3 1.4 1.5 1.6	What are Derivatives and What Do they Do? 1.1Development and Growth of Derivative Markets 1.2Types of Derivatives 1.4Uses of Derivatives 1.7Derivative Disasters 1.8Overview of the Book 1.11Chapter Summary 1.13Suggestions for Further Reading 1.13Problems and Questions 1.14	1.1
2.	FOR 2.1 2.2 2.3 2.4 2.5 2.6 2.7 2.8 2.9 2.10 2.11 2.12 2.13	WARD AND FUTURES MARKETS Forward Contracts 2.1 Differences between Forward and Spot Market 2.1 Futures Contracts 2.2 Futures Market Contract Design 2.4 Physical Settlement, Delivery Options and Cash Settlement 2.6 Futures Markets 2.8 Global Futures Market Size 2.9 Commodity Futures 2.10 Individual Equity Futures 2.11 Stock Index Futures 2.11 Currency Futures 2.12 Futures on Government Bonds, Notes and Bills 2.12 Cash Settled Interest Rate Futures 2.13 Chapter Summary 2.14 Suggestions for Further Reading 2.14 Problems and Questions 2.15	2.1
3.	COS 3.1 3.2 3.3 3.4 3.5 3.6	T OF CARRY MODEL FOR FUTURES AND FORWARDS A Simple Example 3.1 Cash and Carry Arbitrage 3.3 Reverse Cash and Carry Arbitrage 3.5 The Cost of Carry Model 3.6 The Cost of Carry Model with Continuous Compounding 3.7 Consumption Assets and Convenience Yields 3.9	3.1

The McGraw·Hill Companies

xii | Contents

3.7	Value of a For	ward Contra	act 3.10	

- 3.8 Relation between Futures and Expected Spot Prices 3.10
- 3.9 Backwardation and Contango 3.11 Chapter Summary 3.12 Suggestions for Further Reading 3.12 Problems and Questions 3.12

4. RISK MANAGEMENT USING FUTURES AND FORWARDS

4.1

5.1

6.1

7.1

- 4.1 Perfect Hedges and Known Exposures 4.1
- 4.2 Basis Risk and Optimal Hedge Ratio 4.1
- 4.3 Measuring Risk Exposure 4.6
- 4.4 Use of Betas in Index Future Hedges 4.7
- 4.5 Use of Modified Duration in Interest Rate Hedges 4.8
- 4.6 Hedging Currency Risk of Equity Positions 4.10
- 4.7 Tailing the Hedge 4.10
- 4.8 Rolling Hedges 4.14 Chapter Summary 4.15 Suggestions for Further Reading 4.15 Problems and Questions 4.15

5. HOW AND WHY DO FIRMS HEDGE?

- 5.1 Costs and Benefits of Hedging 5.1
- 5.2 Hedging Instruments 5.2
- 5.3 Non-financial Hedges 5.5
- 5.4 Risk Management Structures and Policies 5.6 Chapter Summary 5.8 Suggestions for Further Reading 5.8 Cases 5.9–5.33

6. OPTIONS AND THEIR PAYOFFS

- 6.1 European and American Calls and Puts 6.1
- 6.2 Payoff and Profit Diagrams 6.2
- 6.3 Options as Insurance 6.6
- 6.4 Determinants of Option Price 6.7
- 6.5 Bounds on Option Price 6.8
- 6.6 Put Call Parity 6.11 Chapter Summary 6.12 Suggestions for Further Reading 6.13 Problems and Questions 6.13

7. OPTION MARKETS

- 7.1 Exchange Traded Options and OTC Options 7.1
- 7.2 Options Market Contract Design 7.1
- 7.3 Option Exercise, Settlement, and Assignment 7.2
- 7.4 Options Markets 7.4

Contents | xiii

	 7.5 Commodity Options 7.7 7.6 Stock Index Options 7.7 7.7 Individual Equity Options 7.8 7.8 Currency Options 7.8 7.9 Interest Rate and Bond Options 7.8 Chapter Summary 7.9 Suggestions for Further Reading 7.9 Problems and Questions 7.9 	
8.	RISK NEUTRAL VALUATION	8.1
	 8.1 Risk Adjusted Discount Rates and Certainty Equivalents 8.1 8.2 Risk-neutral Probabilities 8.2 8.3 Cost of Carry Model Revisited 8.3 8.4 Expected Utility Derivation of Risk-neutral Probabilities 8.4 8.5 No Arbitrage and Risk-neutral Probabilities—Single Period 8.6 8.6 No Arbitrage and Risk-neutral Probabilities—Multiple Periods 8.8 8.7 Risk-neutral Valuation in Continuous Time 8.13 Chapter Summary 8.13 Suggestions for Further Reading 8.13 Problems and Questions 8.14 	
9.	THE BINOMIAL OPTION PRICING MODEL	9.1
	 9.1 Single Step Binomial 9.1 9.2 Option Delta and Delta Hedging 9.3 9.3 Multi Period Binomial Trees 9.3 9.4 Option Deltas in the Multi Period Model 9.7 9.5 Matching Volatility and Risk-free Rate 9.8 9.6 More Examples of Binomial Trees 9.8 9.7 Binomial Model for American Options 9.10 9.8 Binomial Trees and the Log Normal Distribution 9.13 Chapter Summary 9.14 Suggestions for Further Reading 9.15 Problems and Questions 9.15 	
10.	THE BLACK-SCHOLES OPTION PRICING MODEL	10.1
	 10.1 The Model of Stock Price Behaviour 10.1 10.2 The Log Normal Distribution 10.2 10.3 The Risk Neutral Log Normal Distribution 10.2 10.4 The Black-Scholes Formula 10.4 10.5 Black-Scholes Model With Dividends 10.6 10.6 Options on Stock Indices 10.8 10.7 Options on Currencies 10.10 10.8 Options on Futures 10.11 10.9 Options on Commodities 10.13 10.10 American Options 10.13 	

xiv | Contents

	10.11	Appendix on Log Normal Distribution10.17Chapter Summary10.18Suggestions for Further Reading10.20Problems and Questions10.21	
11.	USES 11.1 11.2 11.3 11.4	OF OPTIONS: SIMPLE OPTION STRATEGIES Buying Options Instead of Trading Forwards or Spot 11.1 Range Forwards 11.6 Bull and Bear Spreads 11.7 Covered Option Writing 11.9 Chapter Summary 11.10 Suggestions for Further Reading 11.10 Problems and Questions 11.10	11.1
12.	THE 12.1 12.2 12.3 12.4 12.5 12.6 12.7	GREEKS OF THE BLACK-SCHOLES MODEL Delta 12.1 Gamma 12.5 Theta 12.8 The Black-Scholes Equation 12.11 Vega 12.12 RHO 12.14 Numerical Example of Option Greeks 12.14 Chapter Summary 12.16 Suggestions for Further Reading 12.17 Problems and Questions 12.17	12.1
13.	COM 13.1 13.2 13.3 13.4 13.5	PLEX OPTION STRATEGIES Straddles 13.1 Strangles 13.2 Butterfly Spreads 13.4 Greeks of Option Strategies 13.6 Calendar Spreads 13.9 Chapter Summary 13.11 Suggestions for Further Reading 13.12 Problems and Questions 13.12	13.1
14	VOLA 14.1 14.2 14.3 14.4	ATILITIES AND IMPLIED VOLATILITIES Historical Volatility 14.1 Exponentially Weighted Moving Averages 14.3 Garch Method 14.5 Implied Volatility 14.7 Chapter Summary 14.8 Suggestions for Further Reading 14.9 Problems and Questions 14.9	14.1

The McGraw·Hill Companies

VOLATILITY SMILES AND IMPLIED RISK NEUTRAL DISTRIBUTIONS 15.1 15. 15.1 Revisiting the Black-Scholes Formula 15.1 15.2 Volatility Smile and the Risk-neutral Distribution 15.3 15.3 Option Combinations and Volatility Smiles 15.5 15.4 Term Structure of Volatility 15.5 15.5 Volatility Surfaces 15.6 15.6 Implications of Stochastic Volatility 15.7 Implications of Volatility Surfaces for Hedging and Option Greeks 15.8 15.7 Chapter Summary 15.8 Suggestions for Further Reading 15.9 Additional Suggestions for Further Reading 15.9 Problems and Questions 15.9 **16. EXOTIC OPTIONS** 16.1 Digital or Binary Options 16.1 16.1 16.2 Barrier Options 16.5 16.3 Asian Options 16.9 16.4 Chooser Options 16.11 16.5 Compound Options 16.11 16.6 Other Exotic Options 16.12 Chapter Summary 16.12 Suggestions for Further Reading 16.12 Problems and Questions 16.13 **17. WARRANTS AND CONVERTIBLES** 17.1 17.1 Warrants 17.1 17.2 The Warrant Valuation Model 17.1 17.3 Employee Stock Options 17.4 17.4 Convertible Bonds as Straight Bonds Plus Warrants 17.5 17.5 Convertible Bonds: Risk Neutral Valuation of All Cash Flows 17.10 17.6 Convertible Bond Valuation with Interest Rate Uncertainty 17.11 Chapter Summary 17.12 Suggestions for Further Reading 17.12 Cases 17.14–17.31 18. INTEREST RATE AND CURRENCY SWAPS 18.1 18.1 The Swap Markets 18.1 18.2 Using Swaps to Hedge Interest Rate and Currency Risk 18.3 18.3 Interest Rate Swap As Exchange of Floating Rate Bond for Fixed Rate Bond 18.4 18.4 Valuing Floating Rate Bonds 18.6

Contents | xv

- 18.5 Valuing Fixed Rate Bonds using Yield to Maturity (YTM) 18.8
- 18.6 Zero Rates and Forward Rates 18.10
- 18.7 Interest Rate Swap as Bundle of Forward Contracts 18.16
- 18.8 Currency Swaps as Exchange of Bonds 18.17
- 18.9 Currency Swaps as Bundle of Forward Contracts 18.19

xvi | Contents

Chapter Summary 18.20 Suggestions for Further Reading 18.20 Case on Swaps 18.21–18.32

19. CAPS, FLOORS AND SWAPTIONS

- 19.1 Caps and Floors 19.1
- 19.2 The Libor Market Model (LMM) for Cap and Floor Valuation 19.2

19.1

20.1

21.1

22.1

- 19.3 Caplet and Cap Volatilities 19.5
- 19.4 Swaptions and Callable Bonds 19.7
- 19.5 The Swap Market Model for Valuing Swaptions 19.10
- 19.6 Swaption Volatilities 19.13
- 19.7 Risk Neutral Valuation of Interest Rate Derivatives 19.14
- 19.8 Reconciling Lognormality Assumption for Caps and Swaptions 19.15 Chapter Summary 19.16 Suggestions for Further Reading 19.17

20. DERIVATIVE ACCOUNTING

- 20.1 Introduction to Derivatives Accounting 20.1
- 20.2 Derivatives and Fair Value Accounting 20.3
- 20.3 Hedge Accounting 20.7
- 20.4 Requirements for Hedge Accounting 20.13 Chapter Summary 20.14 Suggestions for Further Reading 20.15 Cases 20.16–20.22

21. CORPORATE RISK MANAGEMENT

- 21.1 Risk Management and Shareholder Value 21.1
- 21.2 Lenders, Employees and Other Stakeholders 21.2
- 21.3 Planning and Control Reasons for Hedging 21.3
- 21.4 Financial Distress 21.4
- 21.5 Cash Flow Hedges and Value Hedges 21.5
- 21.6 Capital Structure and Hedging 21.6
- 21.7 Is the Risk Department a Profit Centre or a Cost Centre? 21.7 Chapter Summary 21.8 Suggestions for Further Reading 21.9 Cases 21.10–21.32

22. RISK MANAGEMENT IN FINANCIAL INSTITUTIONS

- 22.1 Value at Risk 22.1
- 22.2 Historical Simulation 22.3
- 22.3 Delta-Normal Approximation 22.7
- 22.4 Delta-Gamma Approximation 22.10
- 22.5 Monte Carlo Simulation 22.12
- 22.6 Modelling and Estimating Correlations 22.14
- 22.7 Back Testing 22.14

Contents | xvii

- 22.8 Stress Testing 22.15
- 22.9 Internal Control Systems 22.17
- 22.10 Regulatory Considerations 22.18
 - Chapter Summary 22.19 Suggestions for Further Reading 22.20

Index

I.1–I.7

The McGraw·Hill Companies

Chapter **One**

Introduction to Derivatives

This chapter introduces the concept of derivatives and to this book itself. Many issues discussed here are elaborated upon at length throughout the work. The chapter begins with a description of what derivatives are and what they do. It then covers the growth and evolution of derivative markets with special reference to the explosive growth of financial derivatives globally since the 1970s. It gives a brief overview of the different types of derivatives and explains how these are used by corporations, banks, investors and individuals to manage their risks. It explains the benefits of derivatives as well as the risks that they pose. Some well known derivative disasters are discussed to point out the risks posed by derivatives, and the role played by management and regulators in mitigating these risks. The chapter concludes with an overview of the book.

1.1 WHAT ARE DERIVATIVES AND WHAT DO THEY DO?

A derivative is an instrument whose value is 'derived' from another security or economic variable. The dependence of the derivative's value on other prices or variables makes it an excellent vehicle for transferring and managing risk.

Consider for example a wheat farmer who enters into an agreement with a baker to sell his wheat crop at a fixed price when it is harvested three months later. This contract is a type of derivative known as a forward contract. It is a derivative because its value depends on the price of wheat which is the primary or underlying asset. This derivative helps both the farmer and the baker to manage their risks. The farmer gets certainty about the price realization and can plan accordingly. The baker also gets certainty about the availability and price of an important raw material for his business.

Derivatives perform several useful economic functions:

- 1 As already highlighted, derivatives allow risk to be managed by hedging or risk transfer.
- 2 Derivatives reduce transaction costs because it is often easier to buy and sell in the forward market than in the cash market. Since delivery and payment take place at a later date, a person can sell in this market even though she does not currently hold a stock of the underlying asset. Moreover, she may be able to buy it back and extinguish her positon in the forward market before the delivery date. This means that she does not give delivery of the asset at all. Derivative markets are therefore much cheaper and more convenient for assets that are hard to store and transport.
- 3 Lower transaction costs often lead to higher liquidity in the derivative market. It is not uncommon for the trading volume in the derivative market to be several times larger than in the cash market.
- 4 By reducing transaction costs and increasing liquidity, derivative markets improve price discovery and lead to more accurate prices in the cash market. This leads to better economic decisions.
- 5 Derivatives increase the attractiveness of the underlying asset by increasing its liquidity and by allowing its risk to be hedged. This brings new classes of investors into the fold of the asset and broadens its appeal.
- 6 In many markets, it has been found that the existence of derivative markets leads to lower volatility. This is due to better price discovery as well as a broader investor base.

1.2 | Derivatives and Risk Management

Despite the above said, derivatives are often severely criticized in the popular press. A variety of ills are attributed to derivatives. Some of these criticisms, however, do not stand up to careful analysis:

- 1 It is sometimes argued that derivatives increase volatility. Since price discovery often happens first in the more liquid derivative market, it often appears that the volatility begins in the derivative market and spreads to the less liquid cash market. However, the absence of a derivative market would have shown the same volatility to be manifested directly in the cash market. Most academic studies conducted over a wide range of underlying assets have shown that derivatives, if anything, reduce volatility. A good example of the misplaced cynicism is the crash of 1987 in the US stock markets which was initially attributed to trading in the futures markets. Subsequently, however, dozens of academic studies about the 1987 crash have exonerated the futures markets.
- 2 Derivative markets have been accused of stealing liquidity from the cash market. This is based on the mistaken idea that there is a fixed pool of liquidity that has to be shared between the cash and futures markets. This is, however, not true. Derivatives, in fact, increase the total pool of liquidity. Studies show that derivatives actually increase liquidity in the cash market.
- 3 There is justifiable concern that derivatives can facilitate manipulation of the cash market. The liquidity and anonymity of the derivative markets does make it easier to accumulate large positions in a secretive manner. Exchanges and regulators have however evolved a variety of measures to prevent this. Regulations require, for example, that large positions should be disclosed to the exchange and to the regulators. There are often limits placed on the size of positions that can be taken by individual entities. Modern derivative markets have been quite successful in dealing with the threat of market manipulation.
- 4 There is some concern that derivatives allow companies and financial institutions to remove assets from the balance sheet and hide the risks associated with these assets. This was certainly true in the past when disclosure requirements related to derivatives were very weak. Accounting standards, however, have improved in the last couple of decades. Regulators have also become more conscious of this problem and have found ways to counter it.
- 5. Many perceptive observers have worried about the complexity of derivatives and doubted the ability of the management to control the risks of these positions adequately. Some well-publicized derivative disasters have lent greater urgency to this worry. Later in this chapter, we will discuss these disasters in order to draw appropriate lessons on the uses and risks of derivatives.

1.2 DEVELOPMENT AND GROWTH OF DERIVATIVE MARKETS

1.2.1 Commodity Derivatives

Forward contracts are probably as old as international trade. When merchants sent consignments to far off countries, they often entered into agreements for the sale of the goods to be delivered as and when the goods reached the destination. These can be regarded as the forerunners of modern forward contracts and on this basis, derivative markets can be traced back to probably as early as 2000 BC.

It has been claimed that the world's first organized futures exchange was the Dojima rice futures market officially set up in 1730 in Osaka, Japan. When the feudal landlords in Japan needed money for their annual tribute to Tokyo, they sold 'rice tickets' representing rice in the warehouse or in the fields. These tickets were bought and sold in the Dujima market and could be regarded as the forerunners of modern exchange traded in futures contracts.

Futures in commodities have thus been around in various forms for several centuries. However, the futures exchanges set up in Chicago and elsewhere in the United States in the nineteenth century are often regarded as the first modern futures markets. Even today, these exchanges are among the largest and most liquid derivative exchanges in the world.

1.2.2 Financial Derivatives

The futures markets were mainly confined to commodities even as late as the year 1960. Agricultural commodities like soya, eggs, and wheat accounted for the bulk of trading. The rest of the derivative market consisted largely of futures contracts in various metals.

The first financial derivatives were created only in the 1970s, but within a couple of decades, they grew rapidly to dominate the global derivatives market. By around 1990, financial derivatives accounted for three-fifths of derivative trading and financial contracts, e.g., the US Treasury bond futures and the Eurodollar futures. These were the most actively traded futures in the world. Despite rapid growth in absolute volumes, agricultural commodities constituted only a fifth of the market with energy and metals accounting for the remaining one-fifth. Since then financial futures have continued to grow very rapidly and in 2005 they accounted for over 90 per cent of the world's futures markets and *all* the 20 most actively traded futures contracts in the world.

When Merton H. Miller wrote that 'My nomination for the most significant financial innovation of the last 20 years is financial futures,' most financial experts agreed with his assessment.

1.2.3 Why have Derivatives Grown so Rapidly?

A number of factors have contributed to the explosive growth of derivatives, particularly financial derivatives since the 1970s.

- 1. The first factor has been the increase in macroeconomic instability and associated increase in risk and volatility since the 1970s.
 - a. After the collapse of the Bretton Woods system, major countries adopted free floating exchange rates that led to a sharp rise in currency risk for businesses and financial institutions. As a result, in order to manage this risk, there was an explosion in the demand for derivatives.
 - b. The two oil price shocks in the 1970s highlighted the risk faced by businesses from volatile commodity prices and increased the demand for instruments that could hedge these risks.
 - c. Rising uncertainty in interest rates also increased the demand for interest rate derivatives to manage these risks.
- 2. Businesses at the same time, became more global in their activities. Increased foreign trade meant greater exposure to exchange rate risks for exporters and importers. In the 1980s, major countries started opening up their economies to international capital flows. As institutional investors built up global portfolios in response to these new opportunities, they needed new instruments to manage associated risks.
- 3. While the above factors doubtless increased the demand for risk management products, developments on the supply side were equally important.
 - a. The early 1970s saw the emergence of the first option pricing formula by Black and Scholes. This was followed by a series of refined and increasingly complex models for pricing more sophisticated derivative instruments.

1.4 | Derivatives and Risk Management

- b. The spread of cheap computing power in the 1980s made these models accessible to a wide range of financial users.
- 4. Financial deregulation in the meantime, led to product innovation and lower transaction costs. These developments continue even today with newer products being designed to hedge risks that were previously very difficult or expensive to hedge. At the same time, financial institutions have become more sophisticated in their use of risk management tools. New classes of investors like hedge funds who use derivatives far more intensively have also increased liquidity in derivative markets. Arbitrageurs armed with highly refined pricing models and vast computing resources have improved price discovery and eliminated pricing inefficiencies in the market.

1.2.4 Globalization of Derivative Markets

The US has been the leader in the development of futures markets since the nineteenth century. It continued to hold the dominant share (80 per cent) of world futures trading volume during 1985. The top 10 futures contracts in the world were US contracts. Since then, derivative markets have developed at an astonishing pace in the rest of the world. By 1990, the US share of global futures trading was only 61 per cent, and four out of the top 10 futures contracts were from outside the US. In 2005, the US share of global futures trading was only 42 per cent, and six of the top 10 futures contracts were from outside the US. With the exception of one contract from Asia (the Korean Stock Index futures contract), all the top futures contracts were from either the US or Europe.

1.2.5 Derivatives in India

Stock index futures were introduced in India in the year 2000. This was followed by options on the stock index and on individual stocks. Futures were also introduced on individual stocks and these have become so popular that India has become the largest single stock futures market in the world.

Futures are also traded on a number of agricultural commodities as well as on energy, precious and non-precious metals. Options on commodities have not yet been introduced.

Interest rate swaps are actively traded in India but other interest rate-related derivatives have yet to be introduced. Currency forwards and options are traded actively in the Indian foreign exchange market.

1.3 TYPES OF DERIVATIVES

1.3.1 Classification by Nature of the Derivative Instrument

The forward contracts described earlier in this chapter are very closely related to the underlying asset. The simplicity of the value relationship between these derivatives and the underlying asset is reflected in the fact that at maturity, the payoff of the forward contract changes by an amount identical to the change in the value of the underlying itself. As shown in Figure 1.1, a rise in the price of the underlying causes an equal gain in the forward and a decline in price causes an equal loss. More generally, derivatives are regarded as being forward-like if they show a linear relationship (the payoff diagram showing the relationship between the payoff at maturity and the price of the underlying is a straight line). Linear relationships typically arise when both parties have fixed obligations under the contract and there is no optionality for either party.

Introduction to Derivatives | 1.5



Figure 1.1 The Payoff of a Long Forward at Maturity is a straight line at 45° indicating that a rise in the price of the underlying causes an equal gain in the forward and a decline in price causes an equal loss.

By contrast, an option is a contract in which the option holder has a right to buy or sell without any obligation to do so. For example, Figure 1.2 shows the payoff of a call option which grants the holder the right (without any obligation) to buy the underlying at a strike price of 1500. If at maturity, the underlying is worth less than 1500, the holder does not exercise the option and the payoff is zero. If the underlying is worth more than 1500 at maturity, the holder will exercise the option, and the payoff is the difference between the price of the underlying and the strike price of 1500. The payoff diagram is no longer a straight line. The non-linearity of the payoff is the defining characteristic of option-like derivative contracts.



Figure 1.2 Payoff at maturity of a long call is zero below the strike price. A long calls payoff thus has the upside of a forward, without its downside.

Based on this discussion, derivatives can be classified as shown in Table 1.1.

1.6 | Derivatives and Risk Management

Swaps

of the derivative instrument itself			
Derivatives without optionality are linearly related to the underlying	Derivatives with optionality have a non-linear relationship with the underlying	Some derivatives combine characteristics of different instruments	
Futures	Options	Hybrids	
Forwards	Options on futures	Embedded options	

Caps and floors on interest rates Options on swaps, caps and floors (swaption, caption, floortion)

Table 1.1 Classification of derivative contracts according to the nature of the derivative instrument itself

1.3.2 Classification by Nature of the Underlying Asset

Another important way of classifying derivatives is according to the underlying from which the derivative derives its value as shown in Table 1.2.

Table 1.2	Classification of derivatives according to the nature of the underlying
	from which the derivative derives its value

Stocks	Foreign Exchange	Interest Rate/ Bonds/Credit	Commodities	Natural phenomena
Single stock futures	Foreign exchange forward contracts	Bond futures and options	Futures and options on agricultural commodities like wheat, soyabean, and milk	Weather derivatives related to rainfall and temperature
Single stock options	Currency futures	Forward rate agree- ments and interest rate futures	Energy derivatives like crude oil, natural gas, and electricity	Derivatives related to natural calamities like earthquakes and hurricanes
Stock index futures	Currency options	Caps, floors, swaps, and swap- tions	Futures and options on precious metals like gold and silver	
Stock index options	Currency swaps	Credit default swaps and other credit derivatives	Futures and options on industrial metals like copper and aluminium	

1.3.3 Nature of the Market

The third way of classifying derivatives is according to the nature of the market in which it is traded. A large number of derivatives are traded in organized futures and options exchanges. These contracts tend to be standardized and designed for anonymous trading between parties who do not necessarily know each other. This requires the exchange design systems to mitigate the credit risk occurring if a party to a derivative contract fails to honour his obligations at maturity.

There are, however, a large number of derivative contracts that are traded on an 'over the counter' or OTC basis between two counterparties who negotiate the terms of the instrument between themselves.

This also means that the two parties assess each other's creditworthiness and negotiate appropriate arrangements for mitigating the credit risk.

Some important examples of exchange traded and OTC derivatives are presented in Table 1.3.

Table 1.3 Classification of derivatives according to the nature of the market in which the derivative is traded

Exchange traded derivatives	Over The Counter (OTC) derivatives
Futures contracts on stocks, currencies, and commodities	Forward contracts on stocks, currencies, and commodities
Exchange traded options on stocks, currencies, and commodities	OTC options on stocks, currencies, and commodities
Swap note futures and interest rate futures	Interest rate swaps, caps, floors, and forward rate agreements

1.4 USES OF DERIVATIVES

1.4.1 How Corporations use Derivatives

Companies use currency forwards and other derivatives to hedge their exports, imports, and other foreign exchange exposures. They use commodity derivatives to hedge raw material consumption and inventories as well as their output prices and inventories. For example, an electrical goods manufacturer might use copper futures to hedge the cost of copper which is a major raw material for it. A gold mining company might use gold futures to hedge the selling price of its output. Companies also use interest rate derivatives to hedge their borrowing costs.

1.4.2 How Mutual Funds and Investment Institutions use Derivatives

Investment institutions use currency forwards and other derivatives to hedge their international asset and liability portfolios. They use swaps and other interest rate derivatives to protect their portfolios from the effects of interest rate risk. They use commodity futures to invest in asset classes, in which they find it difficult to invest directly. Investment institutions also sell options to earn premium income and enhance the returns on the portfolio. Hedge funds and other aggressive investors use derivatives to speculate in various financial markets or to arbitrage between different markets.

1.4.3 How Financial Service Firms, Banks, and other Dealers use Derivatives

Banks and securities firms use derivatives to hedge their inventories of securities. For example, a stock broker might carry large inventories of shares as part of his trading activities. He might use stock index futures to eliminate the market risk of these inventories. Banks often act as dealers in derivative markets to earn dealer spreads by buying a derivative from one customer and selling the same to another at a higher price. They may also seek to make profits by carrying on arbitrage between different markets. Some firms may also speculate on different prices and earn trading profits by taking positions.

1.4.4 How Individuals use Derivatives

Many individuals do speculate on asset prices. A famous example was the late Nobel Prize winning economist, Milton Friedman, who wanted to short sell the British pound in 1969 when his economic

1.8 | Derivatives and Risk Management

analysis convinced him that a devaluation of that currency was imminent. He was unable to do so in the OTC markets available at the time and this experience made him a strong supporter of financial futures in Chicago.

Individuals who manage their own investment portfolios might also use derivatives for the same reasons as investment institutions. In addition, they may use derivatives to hedge their non-tradeable assets. For example, an individual who holds non-transferable employee stock options granted by his employer might use derivatives to hedge the risk of these options.

1.5 DERIVATIVE DISASTERS

This section discusses some of the well-known derivative disasters that have resulted in derivatives getting a bad name in the public imagination. The purpose of this discussion is to identify the lessons to be learned from these disasters so that the benefits of derivatives can be obtained without exposing the company to such risks.

1.5.1 Disasters caused by a Rogue Trader

In several cases, the disaster was caused by a single rogue trader flouting the policies of the company and assuming risk on a scale not known to the senior management. When a trader incurs a loss, he is often tempted to cover it up and hope that the loss will be recouped in the coming days. Traders who are more aggressive, adopt a doubling strategy in order to accomplish this.

Doubling is a classical device adopted by many gamblers. A gambler who loses 1,000 in the first round stakes 2,000 in the second round so that if he wins, he recoups the past loss and makes a profit of 1,000. If he loses again, he stakes 4,000 on the third round so that if he wins, he recoups the past losses (of 1,000 + 2,000 = 3,000) and makes a profit of 1,000. If he loses, he stakes 8,000 on the fourth round so that if he wins, he recoups the past losses of 1,000. If he wins, he recoups the past losses of 1,000. If he loses, he stakes 1,000. If he wins, he recoups the past losses of 1,000 and makes a profit of 1,000. If he loses, he stakes 1,000. If he loses, he stakes 16,000 on the next round and so on. The game is played in this way in the hope that he will not be afflicted with a continuous string of bad luck, and would win at some stage. Whenever he wins, he wipes out all past losses and is left with a profit.

The risk entailed in the doubling strategy is that the gambler may be wiped out completely before his luck favors him. A string of losses may exhaust his financial recources entirely and he can no longer continue with the doubling strategy. For this reason, prudent gamblers do not use the doubling strategy. They prefer to accept an initial loss of \$1,000 and continue the game at a modest level so that it does not entail the risk of bankruptcy. For the same reason, companies do not like the doubling strategy and do not allow their traders to adopt it.

Banks and financial institutions have strict internal rules on the size of positions that any of their traders can take. This rules out the use of the doubling strategy and thus protects the bank from disaster and bankruptcy. Smart traders, however, seek to circumvent these rules. They try to hide their trades or misrepresent the risks involved so that the senior management does not realize what is happening. By the time it finds out, it is too late.

Nick Leeson and Barings Bank A classic example of this phenomenon is the case of Nick Leeson and Barings Bank. Leeson told his superiors that he was carrying out arbitrage between Barings Bank and Nickei futures traded in Osaka and Singapore. He claimed that Barings' membership of both Osaka and Singapore exchanges allowed him to make risk-free profits by exploiting narrow price differences between the two exchanges for the same derivative contract. What his superiors did not realize

Introduction to Derivatives | 1.9

was that Nick Leeson had created an account numbered 88888 ostensibly as an 'error account' to net out minor trading mistakes. Leeson ensured that the Barings software was designed so that positions in the error account were not reported to his superiors at all.

Leeson took large, highly risky positions in the error account and most of his profits actually came from these positions. When the positions started making a loss, Leeson hid the losses in the error account and embarked on a doubling strategy. He managed to convince his superiors that his risk-free arbitrage business needed large amounts of cash to support it. The cash of course was actually needed to finance his losses. As he continued to lose money, the positions became so large that they could no longer be sustained. The episode ended in a loss of \$1.4 billion and caused Barings Bank to go bankrupt.

It is clear that the primary reason for this disaster was a total failure of internal control systems. Leeson's ability to open an error account, change the software to hide the error account, and draw immense amounts of cash from the head office without anyone seriously investigating his actions, all point to severe management failure. In a well-designed organizational structure, the functions performed by Leeson would have been divided between several individuals so that one or two persons could not hide everything.

What role did derivatives play in the disaster? First, the immense liquidity of the derivative market allowed large positions to be taken very easily and quickly. It is this liquidity that makes derivatives attractive to such rogue traders. Doubling strategies can be adopted only in very liquid markets and derivatives tend to fit the bill excellently. Second, derivatives provide enormous amounts of leverage. It is possible to take a large derivative position with relatively small initial investment. Large cash flows arise only when the position makes a loss. Finally, derivatives are complex and make it easier for smart traders to fool their superiors, especially when those superiors do not understand much about derivatives. All these characteristics make it important for banks that trade derivatives to design sound internal control systems to keep rogue traders at bay.

Hamanaka, and Sumitomo Hamanaka was a copper trader at Sumitomo who was nicknamed 'Mr 5 per cent' in recognition of the fact that his position in the world copper market was large, relative to the total size of the market. Over a period of 10 years, Hamanaka hid many of his trades and losses. He refused several promotions to retain control of his copper positions. It was only when Sumitomo ultimately reassigned his duties that it was revealed that Hamanaka had lost \$1.8 billion. Apparently, Hamanaka was trying to corner the global copper market to push up the price of copper. When the strategy failed and prices collapsed, his large long positions produced huge losses. Again, it was the leverage, liquidity, and convenience of the derivative markets that made them so attractive to Hamanaka.

1.5.2 Defective or Inappropriately Implemented Strategies

Another class of derivative disasters stem from defective or inappropriate derivative strategies that were implemented with full knowledge of the top management. The examples of Metallgesselschaft, and Proctor and Gamble illustrate this kind of situation:

Metallgesselschaft MG Refining and Marketing (MGRM) was the US subsidiary of a large German conglomerate Metallgesselschaft. MGRM was engaged in a variety of businesses related to oil exploration, refining, storage, and transport. In early 1993, when oil prices declined below \$20 per barrel, MGRM saw a business opportunity in selling long-term oil contracts to corporations who wanted

1.10 | Derivatives and Risk Management

to lock in the low prices. MGRM sold 10-year contracts totalling 160 million barrels of crude and hedged the risk by buying an equal amount of short term crude oil futures. The contracts were rolled over at maturity to maintain the hedge.

Towards the end of 1993, crude oil fell below \$15. At this point, MGRM was sitting on large profits on its 10-year contracts because its customers were committed to paying MGRM say \$20 per barrel over the next ten years while crude was worth only say \$15. At the same time, MGRM was sitting on large losses on the futures hedge because it was committed to buy crude at \$20 when it was worth only \$15. Ideally, the profits on the long-term contracts should have offset the losses on the futures hedge and insulated MGRM from movements in the price of crude oil. That was the function of the hedge.

There was a major complication, however, as losses in the short-term futures contracts had to be paid out in cash while the profits on the long-term contracts would accrue over the next 10 years. This meant a huge demand for liquidity, which caused the German parent firm to panic. The parent company management was worried that the entire business model was flawed and decided to terminate everything to avoid further losses in future.

There were several deficiencies in the hedging strategy chosen by MGRM. The risks involved in hedging long-term supply contracts with short-term futures had not been fully appreciated and addressed. Later in this book, we will see how MGRM did not decide the crucial issue of hedge quantity in an optimal way. Most importantly, MGRM did not pay adequate attention to the need for large reserves of liquidity to manage its business model. There was little doubt that this liquidity could be provided only by the parent. MGRM had not negotiated the requisite liquidity cover from the parent firm.

Proctor and Gamble In the early 1990s, Proctor and Gamble entered into some highly complex derivative transactions with Bankers Trust. The essence of the transaction at its inception was that at the low level of interest rates that prevailed in the market, the derivative gave Proctor and Gamble an even lower rate of interest on its borrowing. However, the derivative contained a complex and rather bizarre formula based on bond yields and bond prices at different maturities. The effect of this formula was that when interest rates rose dramatically in 1994, the derivative imposed huge losses on Proctor and Gamble. Proctor and Gamble then went to court claiming that Bankers Trust had misrepresented the nature of the derivative and had not explained its complexities properly to them. The case was ultimately settled out of court. It is evident from the bizarre formula that was employed, that the derivative was totally inappropriate to the hedging needs of Proctor and Gamble. The losses were a result of this inappropriate hedging strategy.

1.5.3 Consciously Chosen High-Risk High-Return Strategies

Some derivative disasters were the result of high-risk high-return strategies consciously adopted by the management. This is typically the case with hedge funds that promise high returns and take large risks in order to earn those returns. As long as the going is good, these strategies produce spectacular returns that attract new investors. But when market conditions change, the losses are equally spectacular. Investors often wonder how the same team that produced such excellent results in the past could go so wrong so awfully within a short span of time. One must understand that the initial run of excellent returns and the spectacular collapse are both results of the same high-risk high-return strategy pursued by the fund.

A good example of this is the hedge fund Long Term Capital Management (LTCM) that failed in 1998. With a management team that included veteran bond market traders and Nobel laureate finance

Introduction to Derivatives | 1.11

academics, LTCM pursued a strategy of bond market arbitrage with very high levels of leverage. The arbitrage opportunities that they pursued were very small in percentage terms, but because the fixed income markets are so large, it was possible to make huge profits by taking very large positions financed with borrowing. The leverage and liquidity advantages of derivatives were critical to LTCM's strategy. Tiny profits were amplified by the huge amount of leverage to produce stellar returns for the first few years. When things went wrong in 1998, the collapse was equally rapid as the leverage amplified the losses as effectively as it had amplified the profits in the past. LTCM had also ignored the possibility that the markets could become illiquid and make it difficult for LTCM to implement its chosen strategies. LTCM's size at the point of its collapse was so large that regulators worried about the systemic effect of its failure on the broader financial system. The US Federal Reserve Board orchestrated a partial rescue and orderly liquidation of LTCM.

1.5.4 Lessons Learnt from Derivative Disasters

The benefits of derivatives flow from their liquidity, convenience, and leverage. These same characteristics make derivatives attractive to rogue traders and aggressive investors seeking high returns through high-risk strategies. The complexity of derivatives makes it difficult for senior management and external stakeholders to fully understand the strategies being used. All these factors play a role the derivative disasters.

The principal lesson to be learned from the disasters is that management must understand these characteristics and ensure that they have well-designed internal control systems to ensure that derivatives are appropriately used. The Board of Directors also have a responsibility to lay down risk management policies and ensure that these are adhered to. Regulators also may have a role in ensuring appropriate accounting and disclosure of derivative positions, and in seeing to it that the Board and the top management take their roles seriously.

Later in this book, the chapters on risk management (chapters 21 and 22) will describe how companies and financial institutions can guard themselves against such disasters by using well-defined risk management policies and internal control systems. Derivative disasters can be done away with if these measures are firmly entrenched in a company.

1.6 OVERVIEW OF THE BOOK

The next four chapters of this book primarily deal with forward and futures. Chapter 2 discusses the forward and futures contracts, highlighting the differences between the two. The mechanics of futures markets and their settlement mechanisms are explored and an overview of the global futures markets is presented. Chapter 3 discusses the pricing of forward and futures contracts using the cost of carry models. Chapter 4 analyses the use of futures and forwards for hedging and risk management. The measurement of risk exposures and the design of optimal hedges is discussed. Chapter 5 deals with how and why firms hedge. The reasons for hedging, the choice of hedging instruments and the use of non-financial hedges are explored. A series of case studies are included after this chapter to get a deeper understanding of corporate hedging decisions.

The bulk of this book from Chapter 6 to Chapter 17 is devoted to options. Chapter 6 describes the major kinds of options and their payoffs, and explores the determinants of option price. Chapter 7 discusses the major global options markets on various underlyings. Chapter 8 is one of the most impor-

1.12 | Derivatives and Risk Management

tant chapters in the book. It presents the risk-neutral valuation methodology which is used in the rest of the book to understand derivative valuation. Historically, Black-Scholes and other derivative valuation models were first obtained using complex arbitrage arguments. Only much later was it realized that risk-neutral valuation provides a simpler and more intuitive route to understanding option valuation. This book consistently uses the risk-neutral approach and Chapter 8 is, therefore, critical to the understanding of the whole book. Chapter 9 uses risk-neutral valuation to understand the binomial option pricing model, which is an important model of option valuation. Chapter 10 discusses the famous Black-Scholes model which is the most famous formula in derivative valuation.

Chapter 11 and 13 present various trading and hedging strategies using options. Chapter 11 deals with simple strategies that are designed to take a directional view on the underlying and involve only one or two options. Chapter 13 discusses more complex strategies that involve more than two options, or are designed to take a view on the volatility of the asset rather than the direction of its price movements.

Chapter 12 deals with the Greeks of the Black-Scholes model. The Greeks represent the sensitivities of the option price to various pricing parameters and are critical for hedging the option, and understanding its risk profile.

Chapters 14 and 15 deal with volatility which is the only unobserved pricing parameter in the Black-Scholes model. Chapter 14 discusses the estimation of volatility using historical data as well as the implied volatility of other traded options. Chapter 15 deals with the use of the volatility smile to adjust the volatility input into the Black-Scholes model so that the formula can be used even when there are modest departures from the assumptions of the Black-Scholes model.

Chapter 16 describes exotic options like binary, barrier, and Asian options. The valuation of these options as well as their hedging are discussed. Chapter 17 discusses warrants and convertible bonds which are important examples of embedded options. The use of the binomial option pricing model to handle the complexity of some of these option structures is highlighted. A couple of cases illustrate the issues involved.

Chapter 18 and 19 discuss some of the basic fixed income or interest rate derivatives. This a complex subject and specialized books have been written exclusively on interest rate derivatives. This book deals with only the elementary theory of these derivatives. Chapter 18 discusses interest rate swaps and currency swaps while Chapter 19 covers the standard market models for caps, floors, and swaptions.

Chapter 20 presents an overview of derivative accounting which is important for most users of derivatives. Most of the time, when companies use derivatives to hedge risk, they are concerned about the accounting treatment of the derivatives. The availability of hedge accounting is often a critical consideration while using derivatives. The chapter includes a case that highlights some of the complexities involved.

The final two chapters, 21 and 22, deal with risk management. Chapter 21 which continues from where Chapter 5 ends, discusses corporate risk management. The discussion in this chapter links hedging decisions and capital structure decisions, and brings out the importance of cash flow hedges. The chapter includes a comprehensive case that explores all aspects of risk management, and the design of a complex hedging strategy. Chapter 22 is about risk management in financial institutions. This chapter is largely about Value at Risk (VaR) and stress testing.

Introduction to Derivatives | 1.13

Chapter Summary

A derivative is an instrument whose value is 'derived' from some other security or economic variable. The dependence of the derivative's value on other prices or variables makes it an excellent vehicle for transferring and managing risk. Derivatives bring greater liquidity to the market, improve price discovery, broaden the range of participants, and reduce volatility.

While some forms of derivatives are several thousands of years old, modern organized derivative exchanges date back to the nineteenth old century. Modern financial derivatives have developed largely since the 1970s and have grown to dominate the global derivative markets.

The relationship of the derivative with the underlying can be linear as in the case of forward contracts or nonlinear as in the case of options. The underlying can be a stock or a stock index, bonds or interest rates, exchange rates, commodities or natural phenomena like weather. The derivative can be traded in an exchange or Over the Counter (OTC) between two counterparties.

There have been several derivative disasters in recent years. Many of these are because of rogue traders who exploited the failure of internal control systems to supervise them adequately. Some disasters were due to deficiencies in the design of the hedging strategies while others were simply the result of aggressive investors pursuing high-risk high-return strategies.

Managements, boards, and regulators must understand the liquidity and leverage characteristics of derivatives and ensure that there are well-designed internal control systems to ensure that derivatives are used appropriately.

Suggestions for Further Reading

The seminal paper on the economic function of derivative securities is

- Arrow, K. J. (1964), "The Role of Securities in the Optimal Allocation of Risk-bearing", *The Review of Economic Studies*, 31(2), 91–96.
- The Metallgesselschaft disaster is discussed in:
- Culp, C. L. and M. H. Miller (1995), "Metallgesselschaft and the Economics of Synthetic Storage", *Journal of Applied Corporate Finance*, 7, 62–76.
- Culp, C. L. and M. H. Miller (1995) "Hedging in the theory of corporate finance", *Journal of Applied Corporate Finance*, 8, 121–127.
- Edwards, F. R. and M. S. Canter (1995) "Collapse of Metallgesselschaft: Unhedgeable risks, poor hedging strategy or just plain bad luck", *Journal of Applied Corporate Finance*, 8, 86–105.

The strategy followed by Nick Leeson at Barings is described in:

- Brown, Stephen J. and Steenbeek, Onno W. (2001) "Doubling: Nick Leeson's trading strategy", *Pacific-Basin Finance Journal*, 9, 83–99.
- The leveraged swap transaction between Bankers Trust and Proctor & Gamble is analysed in depth in:
- Srivastava, Sanjay (1998), "Value at Risk Analysis of a Leveraged Swap", Workshop on Risk Management, Isaac Newton Institute for Mathematical Sciences, Cambridge University, http://www.gloriamundi.org/picsresources/ sanjay.pdf

The Sumitomo debacle is explained in:

- "Case Studies: Analysing Sumitomo", International Financial Risk Institute, http://riskinstitute.ch/134800.htm The impact of derivatives on the volatility and liquidity of the underlying markets is analyzed in:
- Darat, Ali F.; Shafiqur Rahman, and Maosen Zhong (2002) "On the Role of Futures Trading in Spot Market Fluctuations: Perpetrator of Volatility or Victim of Regret?", *Journal of Financial Research*, 25(3), 431-444.
- Kumar, Raman; Atulya Sarin and Kuldeep Shastri (1998) "The Impact of Options Trading on the Market Quality of the Underlying Security: An Empirical Analysis", *Journal of Finance*, 53 (2), 717-732.
- Mayhew, S(2000) "The impact of derivatives on cash markets: What have we learned", Unpublished paper, University of Georgia, http://www.terry.uga.edu/finance/research/working papers/papers/impact.pdf

1.14 | Derivatives and Risk Management

- The evolution of financial derivatives is described in the following paper which is also the source for the quotation in the text from Merton Miller about financial futures being the most significant financial innovation:
- Miller, Merton H. (1986) "Financial Innovation: The Last Twenty Years and the Next", *Journal of Financial and Quantitative Analysis*, 21(4), 459-471.
- Data on futures industry volumes is published in the Futures Industry Magazine which is available at the web site of the Futures Industry Association:
- http://www.futuresindustry.org/

The Dojima rice futures market is described in:

Wakita, Shigeru (2001) "Efficiency of the Dojima rice futures market in Tokugawa-period Japan", *Journal of Banking and Finance*, 25(3), 535-554.

Problems and Questions

- 1. India is a major consumer of gold but imports most of its requirements. What economic function is performed by a gold futures market in India in which only Indian entities are allowed to participate?
- 2. How would your answer to the previous problem change if some Indian entities are allowed to participate in both the Indian gold futures market and the international gold futures market?
- 3. The US government at one time considered the possibility of creating a terrorism futures market in which payoffs depend on the occurrence of terrorist attacks. The idea was dropped in the face of strong political opposition. What economic function would such a market have performed?
- 4. Attempts are being made to develop derivative contracts tied to the occurrence of catastrophic natural calamities. What economic function would serve contracts serve? Should insurance companies support or oppose the creation of such markets?
- 5. Is an airline that buys new aircrafts while still losing money following a doubling strategy? If not, how does it differ from the doubling strategy adopted by Nick Leeson?
- 6. A bank that funds long term loans with short term deposits faces a liquidity risk. How does this differ from liquidity risk involved in Metallgesselschaft's hedging strategy?
- 7. After the Hamanaka/Sumitomo attempt to corner the copper market came to light, the London Metal Exchange came in for severe criticism for its alleged failure to police market. What do you think should the exchange have done when Hamanaka was doing his trade?

The McGraw-Hill Companies

Chapter **Two**

Forward and Futures Markets

This chapter introduces the simplest and most fundamental derivative contracts—forwards and futures. These are contracts in which the price has been determined but the settlement (delivery and payment) are to be made at some future date. This chapter explains the differences between futures and forwards and discusses the intricacies of designing a successful futures contract. It also provides information about the important futures markets worldwide in commodities, equities, bonds, interest rates, and currencies. The differences between cash settlement and physical settlement are also highlighted.

2.1 FORWARD CONTRACTS

A forward contract is among the oldest and simplest of derivative contracts. It is simply a purchase or sale transaction in which the price and other terms have been agreed upon, but the delivery and payment are postponed to a later date. Many ordinary commercial transactions are forward contracts in the strict sense, but the term is typically reserved for situations where the time period involved is beyond normal commercial terms.

For example, if a person buys an equity share in the stock market, typically both payment and delivery take place after one or two days. In the strict sense, they are one-day or two-day forward contracts. However, as the time period is very short and is the norm in this market, these are not conventionally described as forward contracts. Similarly, standard dealings between a manufacturer and its suppliers in which a shipment is made two or four weeks after the receipt of the order are not usually regarded as forward transactions. However, a stock market transaction in which delivery takes place after two or four weeks would normally be regarded as a forward transaction because the norm there is to settle within a couple of days.

Forward contracts have been in use for thousands of years all over the world because there are very natural reasons to postpone delivery and payment. For example, centuries ago, a farmer may have contracted to sell his crop to the local trader long before it was ready for harvesting. A travelling merchant in ancient times may have contracted to sell his wares in a foreign land before he actually shipped out the goods from his homeland.

In these examples, the forward contract is a transaction between two parties who know each other well and have learnt to trust each other through their business relationship. Modern forward contracts are different in the sense that at least one party in the deal is often a financial intermediary who thinks of it as a financial derivative.

2.2 DIFFERENCES BETWEEN FORWARD AND SPOT MARKET

The definition of a forward contract suggests two key differences between forward and spot markets:

1. In a spot contract, the payment takes place immediately or almost immediately. In a forward contract, payment is deferred till the expiry of the contract. This implies a gain for the buyer,

2.2 | Derivatives and Risk Management

equal to the interest for the period of the contract. This gain is reflected in the pricing of the forward contract, as we shall see in the next chapter.

- 2. In a spot contract, the delivery takes place immediately or almost immediately. In a forward contract, however, delivery is deferred till the expiry of the contract. This has several implications. First, if the seller already possesses the asset in question, he earns the dividends or other income that the asset produces, but has to expend the costs of storage. For financial assets, the storage costs may be negligible but for other commodities, the storage costs may be quite substantial. Second, it is possible for a seller to sell forward something that he does not possess. In fact, it is possible to sell something that does not even exist at the time of entering into the contract. All of this also impacts the forward price.
- 3. In a spot contract, since delivery and payment take place almost simultaneously, it is possible for a buyer and seller to carry out a transaction even though they do not know each other at all. In a forward contract, both parties have to worry about whether the other would perform the contract on the maturity. Forward contracts can therefore, be concluded only between parties who know each other or have created other mechanisms to ensure default-free performance of contractual obligations. The role of financial intermediaries in many modern financial forward markets is partly to intermediate the credit risk. Two companies that hesitate to deal with each other because of the credit risk may both be willing to deal with a bank which buys from one at a lower price and sells to the other at a higher price. The profit that it earns is partly a reward for bearing credit risk. Assessment of credit risk is anyway the principal function of a bank, and banks are therefore very well positioned to perform this role.

2.3 FUTURES CONTRACTS

As forward contracts entail credit risk, these are mainly suitable for large companies and institutions that are well known to each other, or to their banks. They are less suited to individuals and other small entities. As mentioned before, in an example in 1969, the Nobel prize winning economist, Milton Friedman anticipated a devaluation of the British pound and wanted to sell the pound forward. He found that banks that were not willing to transact with him, were willing to oblige companies that wanted to undertake similar transactions.

This means that a market can thrive only on forward contracts as long as it attracts interest only from a relatively small group of parties that can easily establish its creditworthiness. When a market acquires a much wider interest and becomes very liquid, there is need to create mechanisms so that it is accessible to anybody who wants to trade.

The most effective way to provide open and equal access to all is to organize the activity of trading in an exchange. Exchange traded contracts for delivery at a future date are known as futures contracts to distinguish them from forward contracts.

A futures contract has the following key features that differentiate it from forward contracts:

- Exchange Traded: As already pointed out, futures contracts are traded in an exchange. Thus, there is a great deal of transparency about past trades, prices, and volumes.
- **Standardized:** Futures contracts are standardized in many respects. First, delivery dates are standardized. A forward contract may specify any delivery date by mutual consent of the two parties, but a futures contract on the same asset, typically has a fixed set of delivery dates. There is often only one delivery date in any calendar month and contracts are often named after that
Forward and Futures Markets | 2.3

month. For example, market participants will talk about the September 2007 gold futures contract. Second, contracts are also standardized in terms of the type and quality to be delivered. For example, gold futures contracts typically require delivery of gold bars of 0.995 purity. By contrast, in the spot market, gold is traded in bars, biscuits (wafers), coins and jewellery of purity ranging from 0.995 down to 0.916 (22 carat gold) or even down to 0.75 (18 carat gold). Thus a gold forward contract may stipulate delivery of any form or purity of gold at any date, but the gold futures contract is standardized in terms of date, form, and purity of the gold.

- Mark to Market: Futures contracts deal with credit risk by using an intricate system of margins. Most importantly, futures exchanges impose a system of daily mark to market. Suppose an investor buys gold futures when the price is \$600, but at the next day, its price falls to \$590. The buyer would suffer a loss of \$10 and the seller would gain \$10. The seller would now obviously worry about whether the buyer will honor the contract at maturity. If the buyer were to default, the seller would lose \$10 because she would have to sell in the spot market at \$590, instead of earning \$600 from the futures contract. The system of daily mark to market is that the gains and losses are paid out daily, instead of waiting for maturity. In the above example, the buyer would pay \$10 to the seller (the buyer in fact, would pay \$10 to the exchange which would pay \$10 to the seller). The effect of this mark to market is that a default by the buyer is no longer of any consequence as the seller has already pocketed the gains from the contract. If the gold price rises to \$605 the next day, this would imply a loss of \$15 (605-590) to the seller and an equal amount of gain to the buyer. The mark to market payment that day would be a payment of \$15 from seller to buyer (routed through the exchange).
- **Initial Margins:** The daily mark to market reduces credit risk dramatically, but does not eliminate it completely. The remaining credit risk is that the buyer or seller may default in making the mark to market payment. The exchange deals with this risk by collected an 'initial margin' from both buyer and seller. The margin has to be collected from both because the exchange does not know beforehand whether the price will move up or down, and therefore, whether the mark to market payment will have to come from buyer or seller.
- Novation: Yet another feature of most futures exchanges is the concept of novation. This means that the exchange becomes a counter-party for all trades. Thus if A buys a futures contract from B, the exchange replaces this transaction with two transactions. In one transaction, the exchange buys from B and in another transaction, the exchange sells to A. The advantage of this process is that now neither A nor B needs to know anything about each other as legally speaking, they are not dealing with each other at all. Even if A defaults, the exchange has to fulfil its obligation to B and try and sell the asset in the market at whatever price it can procure. The credit risk is thus intermediated by the exchange, just as in many forward markets the credit risk is intermediated by banks. As the exchange relies on sophisticated margining systems to manage the credit risk, it does not have to charge a large fee to cover the risk.
- Anonymous Trading: One advantage of novation is that it allows anonymous trading. A person buying a futures contract in an exchange does not care who is selling the contract. It is therefore possible to design a trading mechanism that hides the identity of traders from every one except the exchange and its regulators. The exchange needs to know the identity of the traders to implement the margining system, while the regulators need this information to investigate market irregularities that may occur.

2.4 | Derivatives and Risk Management

- Easy Close Out: Another advantage of novation is that it makes it very easy to unwind a futures transaction. If an investor has bought a future contract, he can go back to the market whenever he wishes and sell the future. The party to whom he sells the future may be completely different from the party from whom the contract was bought in the first place. However, the novation means that the original purchase and subsequent sale become transactions supported by the exchange. The position is that the investor had first bought from the exchange and has now sold to the exchange. The two transactions cancel out and the futures contract is simply extinguished. On the other hand, to completely unwind a forward transaction, one has to go back to the original counter-party and extinguish the contract by negotiating a cash settlement. For example, suppose a company bought dollars, 90 days forward at Rs 47 per dollar and a month later, when the rupee weakened by 50 paise, the company found that the previously forecast import transaction that it was hedging, would not materialize. The company in this case, may go back to the bank and request that the forward contract be cancelled, and that the bank make a cash payment of about 50 paise representing the gains that the company has made so far on the forward contract. However, this is something that the company has to negotiate with the bank and if the bank offers to pay only 45 paise, the company may be forced to accept. The only alternative available to the company would be to sell dollars 60 days forward to another bank at the ruling forward price. This may lock in the full 50 paise gain at the cost of having two forward contracts open. Each of the forward contracts may tie up valuable credit lines because when a bank undertakes forward transactions with a company, it regards the transaction as a credit exposure.
- Delivery Rare: One important consequence of standardization and ease of unwinding is that most futures transactions are unwound prior to maturity. Typically 95-98 percent of futures trades are closed out before maturity. Only the remaining 2-5 percent of all transactions actually result in payment and delivery. Standardization of the contract means that the actual grade, quality, and delivery date that the buyer is interested in does not match the standardized grade, quality, and delivery date of the futures market. The buyer would not, therefore, take delivery in the futures market. Instead he buys the desired quality at the desired time in the spot market and uses the futures market only to lock in the price. If the price of the asset rises between inception and the actual transaction date, the buyer has to pay a higher price in the spot market. However, it would make a corresponding gain in the futures market. The result is that he ends up buying at an effective price close to the original futures price. This is discussed in detail in the next chapter.

By contrast, many forward contracts result in actual delivery and payment. This is partly because the forward contract usually represents the actual needs of both parties. For example, in a forward contract between a farmer and a food company, the farmer has a crop to sell and the food company needs the crop as its raw material. It makes sense for both parties to settle the forward contract by actual delivery and payment. Moreover, the terms of the forward contract are usually tailor-made to suit the needs of the two parties and there is little reason, therefore, for settling it any other way. Finally, the difficulty of unwinding forward contracts also acts as an incentive to both sides to settle the contract by actual delivery and payment.

2.4 FUTURES MARKET CONTRACT DESIGN

An exchange that seeks to introduce a futures market on any underlying has to make several choices in designing the contract. A futures contract is highly standardized and it is up to the exchange to achieve this standardization.

Forward and Futures Markets | 2.5

- Underlying: The most important decision that the exchange has to make is to choose the precise quality and grade of the asset on which the futures contract would trade. For some financial assets, this choice is quite straightforward and obvious. In a currency futures, for example, the underlying is the foreign currency and there is little room for any debate on this. In some cases, however, the choice is far from obvious. In the case of a future on a stock index, there is a range of indices to choose from. In India, for example, there is a choice between the BSE Sensex and the S&P CNX Nifty to mention just two popular indices.
- Contract Size: The contract size is the minimum unit in which trading can take place. Contracts that are primarily targeted at large institutions can have large contract sizes often representing a value of a million dollars or more. Contracts oriented towards retail users have much smaller contract sizes. Sometimes, on the same underlying, there is a large contract co-existing with mini contracts which are of much smaller size. For example, the Chicago Board of Trade's main wheat contract is for 5,000 bushels (approximately 135 tonnes), but it also has a wheat mini contract for 1,000 bushels (approximately 27 tonnes). By contrast, the wheat contract at the Indian commodity futures exchanges—National Commodity & Derivatives Exchange Limited (NCDEX) and the Multi Commodity Exchange (MCX)-is for 10 tonnes. In case of equity index contracts, the contract size is defined by the index multiplier. In the case of the S&P 500 contract at the Chicago Mercantile Exchange, the multiplier is \$250 so that every one point rise in the index produces a \$250 gain to the buyer of the contract. The Chicago Mercantile Exchange also trades a mini contract on the S&P 500 where the multiplier is only \$50. By contrast, at the National Stock Exchange in India, the minimum market lot for the Nifty index implies a multiplier of only Rs 100 (approximately \$2) in August 2006, but the level of the Nifty index was about 2 $\frac{1}{2}$ times the level of the S&P 500 index.
- **Delivery:** The precise mechanisms for delivery have to be specified. This is discussed in detail in the next section.
- **Contract Months:** The maturity date is usually standardized with only one contract for each maturity month. The number of such contract months that trade at the same time is a choice that has to be made. One simple example is to have, say twelve contract months open at any time. In this system, as each contract matures, a new contract that expires 12 months later is introduced. A more complex example is the Eurodollar contract at the Chicago Mercantile Exchange which has quarterly contracts expiring in March, June, September and December each year for the next 10 years, along with four serial monthly contracts for the next four months for which there are no quarterly contracts.
- Quotes: The form in which prices are quoted is also standardized. Consider for example, the Eurodollar futures at the Chicago Mercantile Exchange which is a contract on the interest rate on three-month dollar deposits in London. The price is not quoted as the implied interest rate but as 100 minus the implied interest rate percent. So if the market expects the interest rate to be 5 percent, the quoted price is neither 5 nor 0.05, but 100 5 = 95. The purpose is to make its price behave like the price of a bond. If interest rates rise, bond prices fall and vice-versa. In the Eurodollar futures too, if the interest rate rises to 5.5 percent, the quoted price will change to 100.00 5.50 = 94.50.
- **Position Limits:** To prevent one or two very large players from manipulating the markets, many futures exchanges place limits on the maximum position that can be taken by a single person. This varies from contract to contract based on the size of the overall market and the ease of manipulating it.

2.6 | Derivatives and Risk Management

- **Price Limits:** In some futures markets, there is a limit on the maximum price movement permitted in a single day. This is similar to the price limits that are imposed in some stock exchanges as well. Some markets, however, operate without any price limits
- **Price Increments:** This is best explained with an example. Consider the futures on British pounds at the Chicago Mercantile Exchange (CME) in the US. It is customary in the cash market to quote the exchange rate (dollars per pound) to four decimal places. The futures contract follows the same convention. Since the CME pound futures is for £62,500, the minimum price increment of 0.0001 in the quoted price, translates into \$6.25 per contract.

2.5 PHYSICAL SETTLEMENT, DELIVERY OPTIONS AND CASH SETTLEMENT

One of the most important issues in designing a futures contract is the choice of delivery specifications. The major alternatives are physical settlement and cash settlement.

Historically, futures contracts were generally physically settled: at maturity, the seller delivers the underlying and the buyer makes the payment. The high degree of standardization of futures contracts can create problems for delivery because the cash market positions of many hedgers may differ from the precise specifications of the futures contract in terms of quality, location and timing. Exchanges worry about people trying to manipulate the futures market by cornering the available supply of deliverable grades of the asset at deliverable locations during the delivery period. To guard against this risk, exchanges often give a number of options to the seller to make delivery easier. Needless to say, these options can make things difficult for the buyer.

- **Timing Option:** In some contracts particularly financial derivatives, there is no timing option as there is a fixed delivery date. For example, the currency derivatives in the Chicago Mercantile Exchange require delivery on the third Wednesday of the contract month. (As the currency 'spot' market settles after two days, the seller has to initiate the delivery two business days before the third Wednesday and this is also when the futures contract stops trading.) However, the Treasury bond future at the Chicago Board of Trade does have a significant timing option as it allows the seller to deliver the bond at any time of his choice during the expiration month. The seller has to declare his intention to deliver two business days before the actual delivery date and the price at which the buyer pays is then the settlement price of the futures contract on the intention day. Similar timing options exist in contracts like corn and live cattle in the United States.
- End of Month Option: While discussing the timing option above, we had assumed that the seller chose the time of delivery but received the price on the day that he announced his intent to deliver. But this is not necessarily the case. In many instances, the futures stops trading before the delivery month begins while delivery can take place any time during the delivery month. This means that the seller is allowed to announce intent after the futures contract has stopped trading and so the invoice price is the price at which the futures last traded. This option, known as the 'end of month' option has the potential to cause a larger loss to the buyer as the asset price may have fallen between the last trading day of the futures and the intention day.
- **Quality Option:** The Treasury bond option at the Chicago Board of Trade is a good example of a future with a quality option as well. Though the future is on the 30-year-bond, it allows the seller to deliver any bond which has a residual maturity of at least years and is also not callable for at least the next 15 years.

Forward and Futures Markets | 2.7

Whenever there is a quality option, the exchange lays down conversion factors (or equivalently, premiums or discounts) for delivering different qualities. Suppose for example, the conversion factor for the bond that the seller chooses to deliver is 0.90 (equivalently 10 percent discount). The buyer then has to pay 90 percent of the futures settlement price (on the day that the seller announced his intention to deliver) times the contract size.

- Location Option: The futures contract specification lays down the delivery location. Sometimes, alternative delivery locations are permitted, but discounts and premiums are applied as in the case of the quality option. For example, the wheat contract at the National Commodity & Derivatives Exchange Limited (NCDEX) in India requires delivery in Delhi, but allows delivery at other locations with discount/premium. For example, for the contract expiring in October 2006, the discount for delivering in Ahmedabad instead of Delhi was Rs 66 per 100 kg.
- Quantity Variation: In some contracts, a quantity variation is permitted. For example, the Chicago Board of Trade gold contract is for 100 troy ounces. The seller can, however, deliver either a single gold bar weighing 100 ounces or three bars weighing a kg (32.15 ounces) each aggregating 96.45 troy ounces. Officially, the exchange allows a 5 percent quantity variation on either side so that the seller can deliver anything from 95-105 troy ounces, but in practice, the feasible delivery options are 100 ounces and 96.45 ounces. Needless to say, the buyer pays for the quantity actually delivered so that if 3 kilo bars are delivered, the invoice is for 96.45 ounces at the futures settlement price in dollars per ounce.

Many of these delivery options are motivated by the desire to prevent the futures market from being manipulated by some one cornering the bulk of the deliverable stock of the underlying. If the possibility of manipulation is too high, shorts would be reluctant to participate in the market. On the other hand, if delivery options become too wide, the contract may lose its hedging effectiveness, as there is no uncertainty as to what would be received in delivery. No futures contract can thrive without large participation by hedgers. Successful contract design is thus a question of achieving a fine balance between the interest of shorts and longs and reconciling hedging effectiveness with the need to reduce market manipulability.

One solution to all these problems associated with physical delivery is to use cash settlement. Here the idea is that the seller does not deliver the underlying and the buyer does not make the payment, but they only settle the price difference. The seller then sells the quantities and qualities in the spot market and the buyer buys them in the spot market at the times and locations of their choice. The futures market serves only to fix the price and is not used to execute the actual transaction itself.

Cash settlement is widely used today in a variety of contexts:

- Stock index futures are cash-settled because delivering the underlying is practically impossible. To deliver the cash index, a person would have to deliver each of the component stocks of the index in the precise proportions in which they are represented in the index. This would always involve fractional shares, not to mention odd lots. Cash settlement is therefore, much easier.
- Weather derivatives and economic derivatives are cash-settled because they cannot even in principle, be physically delivered. The Chicago Mercantile Exchange for example, trades derivatives based on the temperature in specific cities in the US. These are cash-settled on the basis of the weather index. It is obvious that physical settlement is out of the question (at least until humanity develops the technology to control the weather itself).
- Many contracts which can be physically settled are now cash-settled to reduce the cost of physical settlement. Futures on individual stocks can be physically settled easily, but in many exchanges,

2.8 | Derivatives and Risk Management

these are cash-settled. Euronext (Liffe) in London, trades futures on stocks from around the world and uses cash settlement to eliminate the costs of settling them using the settlement systems of those home countries. In India, futures on single stocks are cash-settled, but there is an ongoing debate on this issue.

2.6 FUTURES MARKETS

2.6.1 Understanding Futures Markets Quotes

The end of day summary for a futures contract (the S&P 500 September 2006 contract) is shown in Table 2.1. The first few items do not need much explanation as they are the familiar open-high-low-close prices that are published for stocks and other securities.

The settlement price is more or less the same as the closing price except that it is often an average of the last few minutes of trading, rather than the Last Traded Price (LTP) itself¹. In futures markets, the settlement price is extremely important because this is used to compute the mark to market margins.

The 'price change' is the change in the settlement price. The change of -30 points or -0.30 implies that the settlement price on the previous day was 1,302.30. Every long (buyer) would face a mark to market outflow because of the price fall while every short (seller) would receive a mark to market inflow for the same reason. As explained earlier, the S&P 500 index has a multiplier of \$250 and therefore the mark to market for a long holding one contract on this day would be \$250 times 0.30 or \$75.

The next item (open interest) is something that is unique to derivative markets. Unlike other securities or assets, derivative contracts do not exist until someone trades them. Consider the first day of trading of a new derivative contract. No futures contracts exist until the first trade. Suppose the first trade is for 100 contracts. At this point, 100 futures contracts exist as some person (say X) is long 100 contracts and somebody else (say Y) is short 100 contracts. We say that the open interest is 100 contracts. Now if the original buyer (X) sells 30 contracts to a new buyer (Z), the open interest is still 100 contracts as X is long 70, Z is long 30 and Y is short 100. Suppose Z now buys another 20 contracts from W. The open interest is now 120 contracts because the aggregate long positions are 120 (X = 70, Z = 30, W = 20) and the aggregate short positions are also 120 (Y = 100, W = 20). The open interest is defined as the aggregate of all long positions or equivalently, as the aggregate of all short positions. The total of all long positions is always equal to the total of all short positions so that the net stock of futures positions (long minus short) is always zero. The open interest is thus a measure of how much hedging or speculative interest there is in that contract. In Table 2.1, the open interest is 598,080 contracts down 10,900 from the previous day.

In stock markets, people often track the 52-week high-low of a stock. In the futures market, it is customary to track the high-low during the life of the contract. This is because many contracts have a life of only a few months and therefore the 52-week high-low is meaningless. In this example, the contract has traded in a range of 1164 - 1342 during its life.

Another quantity which is often mentioned in this context is the notional principal value of the contract. Since the settlement price is 1302 and the index multiplier is \$250, the notional principal value of the contract is $1302 \times $250 = $325,500$. Thus the open interest of 598,080 contracts represents a

¹ A similar distinction is often made in some stock markets too. In the National Stock Exchange, for example, what is published as the 'closing price' is not the LTP but is an average of the last few minutes of trading.

notional principal value of 598,080 times \$325,500 or \$195 billion. Similarly, the traded quantity of 38,934 contracts represents a traded value of 38,934 times \$325,500 or \$13 billion. These dollar values give us a fairly good idea of how active this contract is.

Table 2.1S&P 500 September 2006 Futures at the Chicago Mercantile Exchange on22 August 2006. This table gives the price, volume, and open interest data for the contract

S&P 500 September 2006 Futures				
Open Range	1,299.70 - 1,300.00			
High	1,305.50			
Low	1,297.20			
Close Range	1301.70 - 1302.20			
Settlement	1,302.00			
Price Change (points)	-30			
Total Volume	38,934			
Open Interest	598,080			
Change in Open Interest	-10,900			
Contract High	1,343			
Contract Low	1,164			

2.7 GLOBAL FUTURES MARKET SIZE

The world's financial futures markets traded over a quadrillion (or thousand trillion) dollars of notional value in 2005. The total open interest at the year end was \$22 trillion which represents about 50 percent of world GDP. More information is provided in Table 2.2.

	Open Interest (December 2005)		Trading during 2005	
	Contracts (million)	Notional Value (\$ trillion)	Contracts (million)	Notional Value (\$ trillion)
Financial Futures	107	22	3,172	1,006
All Futures	119		3,982	

Table 2.2Size of the global futures market. Source Bank for International
Settlements, *BIS Quarterly Review*, June 2006

Futures contracts are traded on physical commodities (agricultural, metals, and energy) as well as on financial assets (interest rates, equity indices, and currencies). The relative share of different assets in the global futures market is shown in Figure 2.1 and Figure 2.2. The first figure shows that in terms of contracts outstanding (open interest), interest rate futures and equity index futures dominate the market. Financial futures (which include currency futures in addition to interest rate and equity index futures) account for over 90 percent of the total contracts outstanding. Commodity futures were the earliest derivatives in the world, but now they constitute only a small part of the global derivatives market.

In terms of the number of contracts traded (Figure 2.2), financial futures still dominate with an 80 percent share. The relative rankings are unchanged but interest rate futures are now much larger than equity index futures. However, in terms of contracts outstanding, the two are of comparable size.

The McGraw·Hill Companies

2.10 | Derivatives and Risk Management



Figure 2.1 The share of different classes of underlying assets in the global futures markets in terms of contracts outstanding as of December 2005. A total of 119 million contracts were outstanding on that date. *Source:* Bank for International Settlements, BIS Quarterly Review, June 2006



Figure 2.2 The share of different classes of underlying assets in the global futures markets in terms of contracts traded during 2005. A total of 3.98 trillion contracts were traded during that year. *Source:* Bank for International Settlements, BIS Quarterly Review, June 2006

2.8 COMMODITY FUTURES

Commodity futures are divided into three broad groups:

• Agricultural: The Chicago Board of Trade (CBOT) trades futures on food grains (wheat, rice, and corn), oilseeds (soya), dairy products (milk and eggs), and livestock products (live cattle).

Forward and Futures Markets | 2.11

The New York Board of Trade (NYBOT) trades contracts on food products like coffee, sugar, and cocoa. In Europe, Euronext.liffe trades contracts on wheat, corn, coffee, sugar, and cocoa. In India, the National Commodity & Derivatives Exchange Limited (NCDEX) and the Multi Commodity Exchange (MCX) trades contracts on products like wheat, rice, pulses, sugar, castor seed, soya, and other oilseeds.

- Energy: The New York Mercantile Exchange (Nymex) in the United States and ICE Futures (previously known as the International Petroleum Exchange) in London trade contracts on crude oil, natural, gas, and electricity. In India, crude oil futures trade at the NCDEX and the Multi Commodity Exchange (MCX).
- Metals: The London Metal Exchange trades contracts on most important non-ferrous metals like aluminium, copper, tin, nickel, lead, and zinc as well as some important plastics. Precious metals like gold and silver are traded at the CBOT and the Nymex.

2.9 INDIVIDUAL EQUITY FUTURES

Futures on single stocks are a relatively new phenomenon in modern derivatives exchanges, though some form of this were traded in the stock exchanges under the nomenclature of 'account period settlement'. In the US there was a dispute on regulatory jurisdiction on these instruments between the securities regulator (SEC) and the futures regulator (CFTC) and this prevented single stock futures from being launched till the beginning of this century.

In the twentieth century, single stock futures were traded only in a handful of second-tier markets around the world. The best known of these were the single stock futures introduced in 1994 by the Sydney Futures Exchange (SFE) on a small number of leading Australian stocks. Single stock futures entered the mainstream markets when Euronext.liffe introduced cash-settled single stock futures on leading stocks from around the world in 2001. Subsequently, it also introduced physically settled single stock futures in 2003. Single stock futures started trading in the US in late 2002 at OneChicago, a joint venture of the CME, the Chicago Board Options Exchange and the Chicago Board of Trade.

Single stock futures were launched in India in 2001 and the market has grown to become the largest such market in the world. Other large single stock futures markets in the world are those of South Africa and Spain. Typically, single stock futures contracts do not go out beyond six months and most of the trading is done in the near-month contract.

2.10 STOCK INDEX FUTURES

Futures on stock indices today, are second only to interest rate futures in terms of size. They could be introduced only after cash settlement became acceptable in the early 1980s. The Kansas City Board of Trade launched the first index futures in the world in 1982 on a relatively less popular Value Line Index and the CME followed soon thereafter, with a much more successful contract on the S&P 500. Today index futures are traded in most countries in the world. In India (as in many other countries introducing the futures markets for the first time), index futures were the first futures contract to be launched.

Cash settlement allows the stock index to be treated exactly as if it were the price of an actual share. If an investor holds 100 shares of Infosys and the share price increases by Rs 5, she makes a profit of Rs 500. If she instead holds a Nifty futures contract with an index multiplier of 100 and the index increases by 5 points, she makes a profit of Rs 500. So holding a Nifty futures is like holding 100 shares of an imaginary company, whose share price at any point of time is equal to the Nifty index value at that point

2.12 | Derivatives and Risk Management

of time. The magic of cash settlement allows investors to trade these imaginary shares in the form of index futures.

Index futures are traded on a variety of indices in the same country. In the United States, there are contracts on the S&P 500, the Nasdaq 100 and the Dow Jones Index. Futures also trade on some industry and sectoral indices. In India, the National Stock Exchange trades futures on an index of technology stocks, as well as an index of banking stocks, but these are much less popular than the main index contract.

Index futures contracts usually go out only a few months and most of the trading takes place in the near-month contract.

2.11 CURRENCY FUTURES

Currency futures (introduced at the Chicago Mercantile Exchange in 1972) were the first financial futures in the world, but today they are dwarfed by the interest rate and stock index futures in terms of both volumes and open interest. It is not that currency risk is unimportant or that people do not want to hedge against it. The reason for the relatively low trading in currency futures is that 35 years after they were launched, much of the activity in currency derivatives continues to be in the forward market. This market intermediated by the banks is so liquid and flexible that here, futures have not been able to make much of a headway. Companies which are heavy users of currency derivatives prefer the convenience of forward markets to the daily mark to market and other operational difficulties of using futures. In this book, we will mostly focus on forwards rather than futures when we look at currencies.

2.12 FUTURES ON GOVERNMENT BONDS, NOTES AND BILLS

There are a number of futures contracts on government bonds of various maturities:

- The Chicago Board of Trade offers contracts on the 30-year US Treasury Bond, the 10-year US Treasury Note, the 5-year US Treasury Note and the 2-year US Treasury Note.
- The Chicago Mercantile Exchange offers contracts on the 13-week USTreasury Bills.
- Eurex trades futures on German government bonds—Euro-Bund futures are on bonds with a remaining term of 8.5 to 10.5 years, Euro-Bobl futures are on bonds with a remaining term of 4.5 to 5.5 years, and Euro-Schatz futures are on bonds with a remaining term of 1.75 to 2.25 years. These are the most heavily traded bond futures in the world and serve as the principal fixed income derivatives for the Eurozone.

Futures on government bonds are among the most complex futures contracts because of the complexities of the delivery mechanism. It is not uncommon for the open interest in a bond contract to exceed the total issue size of any individual bond. This rules out futures contracts which allow only one specific bond to be delivered. Instead the futures allow a wide range of bonds to be delivered but apply 'conversion factors' to account for the differences between these bonds.

The futures contract can be thought of as a futures on a hypothetical bond with a 6 percent coupon. This means that if a bond with a 6 percent coupon is delivered, it will have a conversion factor of 1.0. But if a bond with a 5 percent coupon is delivered, it is clearly less valuable than the hypothetical bond and must, therefore, have a conversion factor less than 1.0. The methodology used is to discount the cash flows (coupons and redemption value) of the delivered bond at an assumed yield of 6 percent to

Forward and Futures Markets | 2.13

determine the conversion factor. Suppose we find that at a 6 percent yield, the delivered bond has a value of only 91 percent of par, the conversion factor is 0.91. This would mean that the buyer who receives this bond as delivery would have to pay 0.91 times the futures settlement price on the day of delivery.

But there is one more adjustment to be made. In the bond market, bond prices are always quoted as clean prices. The full or dirty price (which is the actual amount paid) is equal to the clean price plus the accrued interest from the last coupon date to the settlement date. The reason is that the buyer would get the full half yearly coupon at the next coupon date even if he held the bond for only one month. In a way, five months of this interest belongs to the seller who held it for that period. In bond markets, therefore, the buyer compensates the seller for this by paying him the accrued interest. This practice extends to the futures market as well. Therefore:

Invoice Amount = Contract Size × Futures Settlement Price × Conversion Factor + Accrued Interest Accrued Interest = (Coupon /2) × (Number of Days since Last Coupon Payment / Number of Days in a Coupon Period)

Conversion Factor = Present Value at 6 percent yield of the delivered bond per \$1 of par value

Since it is the short who decides which bond to deliver, he obviously chooses the cheapest bond. He will look at each deliverable bond and compute:

- (a) The price at which this can be bought in the cash market
- (b) The price at which it can be invoiced while delivering in the futures market
- (c) The excess of the invoice price (b) over the cash price (a). This is the profit that can be made by buying the bond in the cash market and delivering it in the futures market.

He will then identify the bond for which this profit (c) is the highest. (If the quantity (c) works out to a loss for all bonds, he will find the bond for which the loss is the least.) This is the bond that he will deliver. Needless to say, from the long's point of view, this is the worst possible bond to receive, as it entails the maximum loss from receiving delivery in the futures and selling in the cash market.

Many market participants keep estimating the cheapest to deliver bond at expiry. They then regard the bond future as a future on the cheapest to deliver the bond. Bond futures contracts usually go out four quarters.

2.13 CASH SETTLED INTEREST RATE FUTURES

While the bond futures discussed above are physically settled, there are a number of futures contracts which are cash-settled. One can think of the interest rate itself as being the underlying for these instruments. From this perspective, the contract cannot be physically settled because the interest rate is not an asset that can actually be delivered.

The Eurodollar futures contract was conceptualized in this way and it was launched only after the regulators agreed to allow cash settlement in 1981. This was in fact the first-cash settled futures contract. The underlying for this is dollar Libor, the rate at which the top tier of banks accept dollar deposits in London. The British, Bankers Association surveys 16 eligible banks at 11 am everyday to find out the rates that they are offering. The four highest and four lowest quotes are dropped and the remaining eight quotes in the middle range are averaged and published as the Libor for that day.

The Eurodollar futures contract is based on a 3-month Libor with a notional value of \$1 million. A one basis point (one-hundredth of one percent) change in the interest rate, therefore, affects the contract

The McGraw·Hill Companies

2.14 | Derivatives and Risk Management

value by

$$0.01 \times 1\% \times \frac{1}{4} \times \$1,000,000 = \frac{1}{10,000} \times \frac{1}{4} \times \$1,000,000 = \$25$$

The price is quoted as 100 minus the Libor rate in percent so that a 5.25 percent Libor rate is quoted as 100.00 - 5.25 = 94.75. This makes the futures price behave like a bond price (it falls when interest rates rise, and rises when interest rates fall). The person who is long the Eurodollar futures is thus like a person who is long bonds—he loses money when interest rates rise and gains money when interest rates fall. This makes the contract easier to understand for those who are accustomed to trading bond futures.

The Eurodollar futures contracts are listed for maturities going out 10 years at quarterly intervals. Four nearby monthly contracts are also listed. There is a fair amount of liquidity even in distant contracts.

Chapter Summary

Forwards and futures are contracts in which the price has been determined but the settlement (delivery and payment) are to be undertaken at some future date. Forward contracts are negotiated between two parties who have assessed each other's the credit risk. Futures on the other hand are exchange traded, standardized contracts in which the credit risk is handled by margins and daily mark to market. Successful futures contract design hinges on the correct choice of the underlying and the delivery mechanism, as well as other details like the contract size, contract months, position limits, and margins. Most physically settled futures contracts allow considerable flexibility to the short in terms of delivery options. Excessive flexibility can however decrease the hedging effectiveness of the contract and impede its success.

Open interest is defined as the aggregate of all long positions or equivalently as the aggregate of all short positions. The total of all long positions is always equal to the total of all short positions so that the net stock of futures positions (long minus short) is always zero.

The global futures markets trade a quadrillion dollars of notional value annually. The outstanding open interest is about \$22 trillion or about 50 percent of world GDP. Financial futures (particularly interest rates and equity index futures) dominate the futures market today.

Interest rate futures and stock index futures are cash-settled. Bond futures typically require physical settlement and the delivery options mean that these are essentially contracts on the cheapest to deliver bonds. Commodity futures were traditionally physically settled, but in recent years, several of them have moved to cash settlement. Single stock futures are cash-settled in some markets and physically settled in others. By their very nature, weather derivatives and macroeconomic derivatives are cash-settled.

Suggestions for Further Reading

The websites of the major derivative exchanges provide detailed information and contract specifications of each futures contract. For example:

http://www.cbot.com/ for the Chicago Board of Trade (CBOT) in the United States

http://www.eurexchange.com/ for Eurex in Europe

http://www.cme.com/ for the Chicago Mercantile Exchange (CME) in the United States

http://www.ncdex.com/ for National Commodity & Derivatives Exchange Limited (NCDEX) in India Limited http://www.nseindia.com/ for the National Stock Exchange in India.

Forward and Futures Markets | 2.15

Problems and Questiones

- 1. In 2007, the National Commodity Derivative Exchange (NCDEX) in India faced a problem in connection with its jeera (cumin seeds) contracts. One section of traders argued that due to unfavourable climatic conditions, the bulk of the jeera crip in India was substandard. NCDEX responded by changing the contract specification of new jeera futures (contracts launched in May 2007 or later) to lower quality jeera to be delivered at a discounted price (http://www.ncdex.com/Upload/Circular_all/pdf/2054.pdf). No changes were made in existing contracts. Would the alternatives have been different if cash settlement had been used?
- 2. Daily mark to market is the norm in most futures markets though many exchanges make intra day margin calls on an exceptional basis when there is a large price move. What are the advantages and disadvantages of moving to hourly or more frequent mark to market margins on a routine basis? Is it possible to do mark to market on a real time basis?
- 3. Would advances in information and communication technology allow futures to become more like forwads in terms of customization? Does the success of auction sites like eBay suggest that non standardized futures contracts are feasible?
- 4. Can forward contracts become more like futures if they adopt mark to market practices and move to centralized clearing, netting and collateral management?
- 5. What are the advantages and disadvantages of changing the specification of index futures contracts to require the shorts to compensate the longs for dividends paid on index stocks?
- 6. Would your answer to the above problem be the same if the contracts were single stock futures?

The McGraw·Hill Companies

Chapter **Three**

Cost of Carry Model for Futures and Forwards

A forward contract is perhaps the simplest derivative contract. In this chapter, we derive the cost of carry model for a detailed valuation of forward and futures contract. Unlike what one might think at first glance, the forward price is not equal to the expected spot price at maturity. The key idea in finding the correct value is the notion of arbitrage. The chapter describes the cash and carry and reverse cash and carry arbitrage strategies that enforce the fair valuation of forward and futures contracts. Starting with the simplest case of a non-income earning financial asset, the chapter builds up the cost of carry model for income earning assets as well as for commodities that are bought for non investment purposes as well. It concludes with the relationship between the forward price and the current spot price (contango or backwardation) and describes the determinants of this relationship.

3.1 A SIMPLE EXAMPLE

This chapter starts with a very simple example to illustrate some of the issues that arise in valuing futures and forward contracts. Consider a one year forward contract on an equity share that is not expected to pay any dividends during the next year. The first thought that would come to any one's mind is that the forward price should be equal to the expected share price at the end of one year. Discounting may not appear to come into the picture since the payment in the forward contract takes place at the end of the year and the share price is expected to reap returns at the same date.

But suppose we do try to look at its present value. At first sight, it may appear that since the proposed forward price is equal to the expected stock price at maturity, the present values of the forward price and the stock price would also be equal. A moment's reflection shows, however, that this is not the case. The forward price is fixed right now and is therefore certain. This ought to be discounted at the risk-free rate. The stock price at maturity is, however, unknown and we know only its expected value. This expected value must therefore be discounted at a risk-adjusted discount rate which reflects the required rate of return on that stock. This immediately creates a problem.

If we fix the proposed forward price equal to the expected stock price and then discount the two quantities at two different rates, we will see that the present values are not equal. Suppose for example, the expected stock price is 132, the risk-free interest rate is 5 percent and the required rate of return on the stock is 12 percent. The present value of the stock price (discounted at 12 percent) is 120 and the present value of the proposed forward price of 132 (discounted at 5 percent) is 125.71. The buyer of the forward contract, therefore, makes a payment with a present value of 125.71, but receives something whose present value is only 120. He therefore, loses 5.71 in present value terms.

This suggests that the buyer of the forward contract must agree to a forward price whose present value is only 120. In other words, the forward price must be 126 which means it has a present value of 120 when discounted at the risk-free rate (5 percent) while the expected value of the stock price (132) discounted at the risk adjusted discount rate of 12 percent is also 120.

This analysis suggests a rather complicated way of arriving at the forward price. It appears that we need to know the expected stock price in future as also the required rate of return on the stock in

3.2 | Derivatives and Risk Management

addition to the risk-free rate of interest. An enormous simplification arises when we consider the current stock price in the 'cash' or 'spot' market. This price, also called the 'spot price', is the price for which the stock can be bought now against immediate payment. Elementary finance theory tells us that the spot price of a stock that is not expected to pay dividends during the next year should be equal to the expected stock price at the end of the one year period discounted at the required rate of return on the stock.

In the earlier example, the spot price should be $\frac{132}{1.12} = 120$. The forward price is obtained as $120 \times 1.05 = 126$. Alternatively, the present value (at 5 percent) of the forward price is equal to 120. This gives us a very simple rule for finding the forward price i.e.

- The present value (at the risk-free rate) of the forward price equals the spot price, or
- Forward price equals the spot price plus interest at the risk-free rate.

This rule involves only two things that are easily observed in the market — the current stock price and the risk-free rate of interest. We do not at all need to know the expected stock price in future, nor do we need to know the required rate of return on the stock. It is true that both these quantities do impact the forward price, but we do not need to know them because both these quantities are impounded in the current (spot) stock price. It is therefore sufficient to know the spot price.

It is worthwhile thinking a bit about this remarkable phenomenon where the current stock price conveys all requisite information to find the forward price. What is it, that makes this happen? The key is, as the stock does not give any dividend during the next year, there is no benefit from holding the stock for one year, except the value of the stock at the end of the year. Put differently, the current stock price is determined solely by what it is likely to be worth at the end of one year.

The stock is what we shall call an investment asset. This must be contrasted with a consumption asset that provides value through its use during the year. For example, the present day value of a laptop is not determined by its expected value at the end of one year. (Given the rate of obsolescence of computing equipment, the laptop at the end of one year might be worth only a fraction of what it is worth now). A major part of the value of a laptop computer arises out of the use of the computer during the course of the year.

An asset may act partly as a consumption asset and partly an investment asset. A large part of an agricultural crop, for example, might be stored immediately after the harvest in the expectation of a price rise in subsequent months. This would make the crop an investment asset, whose price is largely determined by its expected price in future. (However, unlike the stock, this asset might have significant storage costs that need to be accounted for.) On the other hand, shortly before a harvest, the price of the same crop would probably reflect its immediate consumption value. The price of the crop at this point could be well above what it would be in the post-harvest month. Therefore, no one would now hold the crop as an investment. Stocks of the crop would be held only to support the consumption needs of the pre-harvest period. Therefore, the crop would be a consumption asset at that point of time.

A financial asset can also behave like a consumption asset in given rare situations. For example, when two groups are vying for management control of a company, its stock might have a highly inflated price that reflects its voting power. It can be expected that the stock's price would fall after the control issue is settled. At this point, rational investors should not hold the stock as an investment asset. They should rationally sell it to one of the contending players, unless they hold it as a means to support one of the contending players. Cost of Carry Model for Futures and Forwards | 3.3

All these examples indicate that the simple equivalence — the present value of the forward price equals the spot price — rests on several simplifying factors:

- the stock pays no dividends or other income during the life of the forward contract
- the stock has negligible storage costs
- the stock is not a consumption asset that provides benefits from possession or use

Later in this chapter, we shall relax a number of these assumptions and examine their impact on the pricing of forward contracts. But before we do that, we shall look at forward prices using the idea of arbitrage. Arbitrage is one of the key ideas of derivative pricing¹. and it is important to understand how arbitrage considerations allow us to determine forward and futures prices.

3.2 CASH AND CARRY ARBITRAGE

Let us return to the earlier example of a non-dividend paying stock that is trading at 120 when the risk-free rate of interest is 5 percent. Using the idea that the forward price is equal to the spot price plus interest, the one year forward price was determined to be $120 \times 1.05=126$.

Suppose, the forward price were higher than 126, say 128. The following arbitrage opportunity would then arise:

- 1. Borrow 120 at the interest rate of 5 percent.
- 2. Buy the stock now.
- 3. Sell the stock one year forward at the forward price of 128.

Since the stock purchase is financed by borrowing, there is no cash flow at the inception of this arbitrage strategy. At maturity, the forward contract is settled by delivering the stock which has been bought now. The forward price of 128 is received and out of this, the loan is repaid with interest which on the loan of 120 at 5 percent amounts to 6. The total repayment is, therefore, 126. There is a clean profit of 128 - 126 = 2. This is an arbitrage profit because:

- the arbitrage position does not require an initial investment
- there is no negative cash flow at any point of time
- there is a positive cash flow of 2 at maturity

In other words, the arbitrage opportunity produces a profit of 2 without any risk out of a zero investment. It thus produces something out of nothing and is also known as a free lunch. However, this is not possible in an efficient market where it is possible to make a gain only by either making an investment or taking risks.

This happens because many people see an opportunity and start implementing the arbitrage strategy of buying the stock now and selling it forward. In the process, the current price of the stock goes up (due to buying pressure) and the forward price falls (due to selling pressure). This process stops only when the arbitrage opportunity vanishes. This happens when the forward price is equal to the spot price plus interest.

Cash and carry arbitrage involves buying the stock immediately in the cash market and carrying it till maturity. It ensures that the forward price of an asset cannot exceed the spot price plus interest.

^{1.} The other key idea in derivative pricing is risk-neutral valuation and we shall see in chapter 8, how this idea can be used in the valuation of forward and future contracts.

3.4 | Derivatives and Risk Management

In the presence of storage costs, forward price = spot price + interest costs + storage costs. Suppose for example, that crude oil is trading at 70 a barrel when the interest rate is 6 percent and monthly storage costs amount to 0.60. Assuming that the storage cost is paid at the end of the month, cash and carry would incur the following costs for a one month period:

- 1. Spot price = \$70
- 2. *Plus* one month interest at 6 percent on \$70 = \$0.35
- 3. *Plus* one month storage cost =\$0.60

Thus the total costs of cash and carry would be 70.00 + 0.35 + 0.60 = 70.95. If one month crude oil forward trades above \$70.95, it would be profitable to implement cash and carry arbitrage.

If the asset pays a dividend or other income, it reduces the cost of cash and carry arbitrage. Consider for example, a bond that carries a coupon of 8 percent which is paid semi-annually. Consider a six month forward contract on the bond. Assume that the bond has just paid a coupon and therefore the next coupon would be received just before the forward contract matures. If the bond is now trading at 105.50 and the interest rate is 4 percent, then the cost of implementing cash and carry are as follows:

- 1. Spot price = 105.50
- 2. Plus six months interest at 4 percent = $\frac{1}{2} \times 105.50 \times \frac{4}{100} = 2.11$
- 3. Less semi-annual coupon = $\frac{1}{2} \times 8$ percent $\times 100 = 4.00$ or $\frac{1}{2} \times \frac{8}{100} \times 100 = 4.00$

Thus the total costs of cash and carry would be 105.50 + 2.11 - 4.00 = 103.61. The price for a six month forward contract on this bond should therefore be 103.61, which is less than the current price of the bond of 105.50.

The principal barrier to cash and carry arbitrage is the ability to borrow money to buy the stock. If borrowing is difficult or expensive, then cash and carry arbitrage is impeded. For instance, it may be difficult to borrow large amounts of money. Or it may be possible to borrow only at a cost higher than the risk-free rate of interest. In given market, many investors would be subject to either of these borrowing restrictions. However, it may be sufficient for a few large investors to perform the arbitrage for the prices to fall back in line. Some large institutions may be able to borrow at close to the risk-free rate by providing their investment portfolio as the collateral. Moreover, some investors may be able to perform cash and carry arbitrage without actually borrowing money. Investors who are lending out their surplus funds could instead use these cash surpluses to buy stock. The opportunity cost of doing so is the interest earned on the surplus cash. This is analogous to borrowing at the risk-free rate of interest.

In non-financial assets, another barrier to cash and carry arbitrage is the ability to store the asset. If an investor wants to undertake cash and carry arbitrage in crude oil, he will need access to tanks where the crude oil can be stored. Fire being a prominent hazard, only specialized facilities can be used for this purpose. Most of these storage facilities might be owned by the big oil companies and might be fully utilized for operational stocks. Thus storage facilities may not be available on the scale required to implement cash and carry arbitrage. One way to get around this problem is to think of the storage cost as being extremely high instead of storage being unavailable. Cost of Carry Model for Futures and Forwards | 3.5

3.3 REVERSE CASH AND CARRY ARBITRAGE

We have seen that cash and carry arbitrage prevents the forward price of a non-dividend paying stock from exceeding the spot price plus interest costs However, to assert that this is the forward price of the stock, we need to explain why the forward price cannot be less than this. This is the function of reverse cash and carry arbitrage. This arbitrage as the name suggests, is the reverse of cash and carry arbitrage, in that, the seller sells in the spot market and buys in the forward market.

It is easiest to think of this as being done by somebody who already owns the stock. The person sells the stock in the spot market, lends the proceeds at the risk-free rate and buys the stock back in the forward market. He gives up the stock only temporarily because at the maturity of the forward, he reacquires the stock.

Let us return to the original example of a non-dividend paying stock that is trading at 120 when the risk-free rate of interest is 5 percent. Using the idea that the forward price is equal to the spot price plus interest, the one year forward price was determined to be $120 \times 1.05 = 126$.

Suppose, the forward price were lower than 126, say 123. The following arbitrage opportunity would then arise:

- 1. Sell the stock at 120.
- 2. Lend the proceeds (120) at the interest rate of 5 percent.
- 3. Buy the stock one year forward at the forward price of 123.

At the end of one year, the seller gets the loan amount of 120 back along with an interest of 6 percent. Out of 126, he pays 123 under the forward contract to buy back the stock. He now has the same stock that he had earlier and has also made a risk-less profit of 3.

Many other investors will also see the same arbitrage opportunity and will sell the stock in the spot market, and buy it back in the forward market. As they do so, the pressure on selling will cause the spot price to fall while pressure on the buying will cause the forward price to rise. The arbitrage will stop when the forward price is equal to the spot price plus interest.

The impact of storage costs and dividends/coupons on reverse cash and carry arbitrage is the mirror image of their impact on cash and carry arbitrage. A person who owns the stock and implements reverse cash and carry arbitrage avoids storing the asset and thus saves the storage costs. However, he foregoes any dividends or coupons from the asset during the life of the forward contract.

Consider for example, the crude oil example discussed earlier, where crude oil trades at \$70 a barrel when the interest rate is 6 percent and monthly storage costs amount to \$0.60. Assuming that the storage cost is paid at the end of the month, reverse cash and carry would produce the following cash flows at the end of a one month period:

- 1. Sale proceeds of one barrel at spot price = \$70
- 2. *Plus* one month interest earned at 6 percent on \$70 = \$0.35
- 3. *Plus* one month storage cost saved = 0.60

Thus the total cash flow at the end of one month from reverse cash and carry would be 70.00 + 0.35 + 0.60 = 70.95. If one month crude oil forward trades below \$70.95, it would be profitable to implement reverse cash and carry arbitrage. For example, if the forward price is \$70.15, then there is an arbitrage profit of 70.95 - 70.15 = 0.80 from this transaction because it costs only \$70.15 to buy back the barrel of crude while the spot sale, together with interest and storage costs saved, amount to \$70.95.

Similarly, consider the earlier bond example where an 8 percent bond is trading at 105.50 immediately after paying a semi-annual coupon and the interest rate is 4 percent. Reverse cash and carry for six months produces the following cash flow at the end:

3.6 | Derivatives and Risk Management

- 1. Sale proceeds of one bond at spot price = 105.50
- 2. *Plus* six months interest earned at 4 percent = 2.11
- 3. Less semi-annual coupon foregone = $\frac{1}{2} \times 8\% \times 100 = 4.00$ or $\frac{1}{2} \times \frac{8}{100} \times 100 = 4.00$

Thus the total cash flow at the end of six months from implementing reverse cash and carry would be 105.50 + 2.11 - 4.00 = 103.61. The price for a six-month forward contract on this bond should therefore be at least 103.61.

Reverse cash and carry arbitrage can also be implemented by people who do not own the asset. Such arbitrageurs must borrow the asset from someone else before they can sell it. Just as cash and carry arbitrage required the ability to borrow money, reverse cash and carry arbitrage requires the ability to borrow the securities concerned. In many cases, large institutions are able to borrow and sell securities that they do not own. This process is known as 'short selling' or simply 'shorting'. When reverse cash and carry arbitrage is implemented through shorting, the process is as follows:

- 1. Borrow the stock and sell in spot.
- 2. Lend the proceeds at the risk-free interest rate.
- 3. Buy the stock one year forward at the forward price.
- 4. At maturity, reverse all of the above. Settle the forward contract using the loan repayment, take delivery of the stock and use it to repay the borrowed stock.

Reverse cash and carry is difficult to undertake for those securities, which owners are unwilling to sell or lend. Earlier in this chapter, we talked of consumption assets that are valued for their use and not just for their resale value. Such assets would clearly be difficult to borrow because the lender foregoes the use of the asset during the period of the loan. Therefore, consumption assets can also be thought of as assets which are difficult or very expensive to short.

Reverse cash and carry is also impeded when there are no inventories to sell or borrow. For example, an agricultural commodity being traded just before a new harvest would fit into this category. Since the price is expected to fall after the harvest, there would be very little inventories of the commodity. The only existing inventories would be the ones that meet anticipated consumption during the period prior to the harvest. Thus, reverse cash and carry will not be possible. This again, is an example of a consumption asset, rather than an investment asset.

Short-selling is sometimes impeded by regulatory obstacles. Many financial assets which are clearly investment assets, might be subject to this problem. Historically speaking, many governments prevented short-selling of government bonds. These restrictions have been lifted in many countries but they remain in some others including India. At times, short sale restrictions are also imposed in equity markets to prop up prices during periods of crisis. During the Asian crisis of 1997 for example, Hong Kong imposed stringent restrictions on shorting equities. Under such conditions, reverse cash and carry arbitrage becomes difficult.

3.4 THE COST OF CARRY MODEL

Cash and carry arbitrage prevents forward prices from being above the theoretical price, while reverse cash and carry arbitrage prevents prices from dipping below the theoretical price. Together, the two ensure that the forward price equals the theoretical price.

What is this theoretical price? We have seen several versions of this:

- spot price plus interest costs
- spot price plus interest costs plus storage costs
- spot price plus interest costs minus dividends or coupons

All of these can be summarized in one general expression

spot price plus net cost of carry

Here 'net cost of carry' refers to the net costs (apart from the spot price) incurred in undertaking a cash and carry arbitrage. Needless to say, this is also the net cash flow (apart from the spot price) that is obtained in a reverse cash and carry arbitrage:

• net cost of carry equals interest costs *plus* storage costs *less* income (dividend or coupons)

It is also important to take time value into account. While defining 'net cost of carry', we have implicitly assumed that interest costs and storage costs are incurred on the date of maturity, and that the income yield also arises on that date. In practice, these cash flows occur at various points through the life of the contract. To be more precise, therefore, we must present value everything.

$$PV (forward price) = spot price + PV (storage costs) - PV (income)$$
$$PV (F) = S + U - I$$
(3.1)

In this expression, F is the forward price, S is the spot price, U is the present value of storage costs and I is the present value of the income from the asset. Interest cost disappears because it is subsumed while computing the present value of the forward price. Interest cost is in fact, equal to the difference between the forward price and its present value.

This expression is very general, in that, it can take into account cash flows at the beginning, at the end, or at any time in between. It is, however, rather inconvenient. A much simpler expression is obtained if we assume that storage costs and asset yield are both evenly spread out over the life of the contract. This will require that we use continuous compounding.

3.5 THE COST OF CARRY MODEL WITH CONTINUOUS COMPOUNDING

What is the difference between 6 percent compounded annually and 6 percent compounded semiannually? Under semi-annual compounding, the interest for six months at 6 percent is applied at the end of the first half-year and the interest for the second half-year is computed on this enhanced principal. Therefore 100 of principal becomes:

$$100 + \frac{1}{2} \times 0.06 \times 100 = 100 (1 + 0.03) = 100 \times 1.03 = 103$$
 at the end of the first half year;

and the interest for the second half year is $\frac{1}{2} 0.06 \times 103 = 103 \times 0.03 = 3.09$. The total interest is thus 3.00 + 3.09 = 6.09 as against 6.00 under annual compounding. The interest is effectively 6.09 percent instead of 6.00 percent. The total amount to be repaid is then 106.09 which can also be got more simply as $100 \left(1 + \frac{0.06}{2}\right) \left(1 + \frac{0.06}{2}\right) = 100 \left(1 + \frac{0.06}{2}\right)^2 = 100 \times 1.03^2 = 106.09$. If the rate of interest in general is *r* compounded annually, the amount at the end of one year under semi-annual compounding is

The McGraw·Hill Companies

3.8 | Derivatives and Risk Management

obtained by multiplying the principal by $\left(1+\frac{r}{2}\right)^2$. If instead, compounding were done monthly, this process would be repeated 12 times in the year and the amount at the end of the year would be the principal times $\left(1+\frac{r}{12}\right)^{12}$. If *r* is 6 percent, we would get $1.005^{12} = 1.0617$ so that the interest rate is effectively 6.17 percent. Nothing prevents compounding from being done daily $\left(1+\frac{r}{365}\right)^{365}$ or even hourly $\left(1+\frac{r}{8760}\right)^{8760}$. If *r* is 6 percent, then daily compounding would give us 6.1831 percent and hourly compounding would give us 6.1836 percent.

Thus beyond a point, compounding more frequently makes very little difference. Mathematically, it is possible to imagine compounding being done every minute or every second, or being done continuously. It can be shown that if we compound continuously, the amount at the end of the year becomes e^r and the effective interest rate becomes $e^r - 1$. In these formulas, *e* is a mathematical constant² whose value is 2.71828. If r is 6 percent, then continuous compounding would give us an effective interest rate of $e^{0.06} - 1 = 1.061837 - 1 = 0.061837 = 6.1837$ percent.

It is also possible to go in the reverse direction. If the effective annual rate of interest is given as R, the equivalent continuously compounded rate of interest is given by $\ln (1 + R)$ where \ln denotes the natural logarithm. For example, if R is 6 percent, the equivalent continuously compounded rate of interest is given by $\ln (1.06) = 0.0583 = 5.83$ percent. This can be verified by noting that $e^{0.0583} = 1.06 - 1 = 0.06 = 6$ percent.

Using continuous compounding, it is possible to rewrite Eq (3.1) as follows:

$$PV(F) = S + U - I$$

$$e^{-rT} F = S + U - I$$

$$F = e^{rT} (S + U - I)$$
(3.2)

For a non-dividend paying financial asset, there is no income and no storage cost. Therefore, the expression simplifies as:

$$F = e^{rT}S \tag{3.3}$$

For a financial asset (no storage costs) with an income yield, the expression can be written as:

$$F = e^{rT} (S - I) \tag{3.4}$$

Now imagine that there is a continuous stream of income which occurs at a constant rate q as a fraction of the asset value. The income stream q offsets the interest cost r and the net cost of carry is at the rate of r-q. In this case, Eq (3.3) is modified to:

$$F = e^{(r-q)T}S \tag{3.5}$$

Even if the income stream is not actually continuous, we can find the equivalent continuous income yield by equating Eq (3.4) and Eq (3.5) to find the equivalent q.

^{2.} Mathematically, e is defined as $\lim_{n \to \infty} \left(1 + \frac{1}{n}\right)^n$. For the purposes of this book, it is sufficient to know its numerical value of 2.71828. In most business and statistical software e^r can be computed by using the expression exp(r).

Cost of Carry Model for Futures and Forwards | 3.9

We can rewrite Eq (3.4) as
$$e^{rT}S\left(\frac{S-I}{S}\right)$$
 and Eq (3.5) as $e^{rT}S e^{-qT}$. These two expressions become

equal if $\left(\frac{S-I}{S}\right) = e^{-qT}$ or $q = \frac{-\ln\left(\frac{S-I}{S}\right)}{T} = \frac{\ln\left(\frac{S}{S-I}\right)}{T}$. If we are willing to use this value of q, we

can always use Eq (3.5) instead of Eq (3.4).

We can also convert the storage cost U in Eq (3.5) into the form of a cost that is incurred continuously as a fraction u of the value of the asset³. We rewrite $e^{rT}(S+U-I)$ as $e^{rT}(S-I)\frac{S-I+U}{S-I}$. The first

part of this $e^{rT}(S-I)$ has already been rewritten as $e^{(r-q)}S$. The second half $\frac{S-I+U}{S-I}$ can be rewritten

as
$$e^{uT}$$
 by simply defining u as $\frac{\ln\left(\frac{S-I+U}{S-I}\right)}{T}$. We then get the more general version of Eq (3.5) as:
 $F = e^{(r-q+u)T} S$ (3.6)

3.6 CONSUMPTION ASSETS AND CONVENIENCE YIELDS

One way to look at consumption assets is that their use or possession yields a benefit over and above any explicit income stream. Consider for example, the safety stock of a raw material in a factory. Most factories would hold an inventory that exceeds the anticipated consumption of the raw material till the next replenishment of inventories. This safety stock is designed to protect against delay in delivery or unanticipated increase in consumption of any kind. The lead time-from placing an order till the receipt of fresh supplies—results in this safety stock, which clearly serves a useful function. The company clearly knows that this safety stock is expensive because of interest and storage costs. It still carries the stock because the expected benefit in terms of uninterrupted operation of the factory exceeds these storage and interest costs. This benefit is not an explicit and directly observable quantity like the coupon or dividend on a financial asset, but is nevertheless quite real and significant.

If we could observe and measure this benefit, we could use this (let us call it y) instead of the observable income yield q in Eq (3.6) and obtain the forward price formula for consumption assets as:

$$F = e^{(r - y + u)T}S \tag{3.7}$$

This quantity y is called the convenience yield as it reflects the convenience of having the commodity in stock, instead of placing an order and waiting for the supply to arrive.

The principal difference between Eq (3.6) and Eq (3.7) is that while q is easily observed and measured, y is neither directly observeable nor it is easily measurable. Instead, the convenience yield is defined using Eq (3.7). We can observe the forward price F and the spot price S, as well as the interest rate r. We can measure the storage cost u fairly well. We can then use Eq (3.7) to infer the value of y.

^{3.} Some storage costs are indeed a percentage of the value of the asset. Insurance costs for physical assets are usually linked to value. For financial assets, there is often a custody charge linked to value, though this cost is usually negligible.

3.10 | Derivatives and Risk Management

One might object to this method, in that, we are running around in circles as we are using the forward price to find the convenience yield and then using this convenience yield to find the forward price. In practice, we may use past data to find the convenience yield and hope that it will remain at the same level in future. This allows us to determine the forward price using Eq (3.7). In some commodities like crude oil, where convenience yields are large and volatile, people have built sophisticated models that attempt to predict the convenience yield based on variables like the level of inventories.

3.7 VALUE OF A FORWARD CONTRACT

There is a subtle difference between the forward price and the value of a forward contract. The forward price that has been determined so far is essentially the price that is equally acceptable to the buyer and the seller, and therefore, is the price at which the forward contract has zero value at its inception. If the forward contract had positive value at inception, nobody would sell forward, and if had a negative value, nobody would buy forward. At its inception, therefore, the forward contract has zero value.

But after some period of time, the forward contract would have a non-zero value. Suppose for example, an investor bought a three-month forward contract two months ago, at a forward price of 113. Let this be the delivery price. As two months elapse, the forward contract now becomes a one-month forward contract. The delivery price specified in the contract of course, remains fixed at 113. If the one-month forward price is now 115 clearly, our investor is sitting on a profit. If he were to enter into a forward contract, it would be at a price of 115, but his existing contract allows him to buy at 113. The existing contract, therefore, reduces his cash outgo at maturity by 2 from 115 to 113. The value of the forward contract is therefore, the present value of this 2.

In general therefore, the value f of a long forward contract with delivery price K is given by the equation f = PV(F - K) where F is the current forward price for the residual maturity of the contract. If the residual maturity is t and the continuously compounded interest rate is r, we can rewrite this as:

$$f = e^{-rt} \left(F - K \right) \tag{3.8}$$

Correspondingly, the value of a short forward contract is given by $-f = e^{-rt}(K - F)$. In the absence of any income, yield or storage costs, the forward price *F* is given by Eq (3.3) as $F = e^{rT}S$. We can then rewrite Eq (3.8) as

$$f = e^{-rT} (e^{rT}S - K) = S - e^{-rT} K$$
(3.9)

Similarly, if we use the more general Eq (3.2) where $F = e^{rT}(S + U - I)$, we get

$$f = S + U - I - e^{-rT} K$$
(3.10)

Alternatively, if we use Eq (3.5) where $F = e^{(r-q)T}S$, we get

$$f = e^{-qt}S - e^{-rt}K \tag{3.11}$$

3.8 RELATION BETWEEN FUTURES AND EXPECTED SPOT PRICES

We began this chapter by exploring the naïve idea that the forward or futures price is the expected spot price at maturity and saw that this idea is obviously mistaken in the case of a non- income earning financial asset. If we apply Eq (3.1) and note that storage costs and income are both zero, we find that the present value of the forward price equals the current spot price. In other words, the forward price is the current spot price plus interest at the risk-free rate. On the other hand, the expected spot price at

Cost of Carry Model for Futures and Forwards | 3.11

maturity equals the current spot price plus the expected return on the asset. The expected rate of return on the asset equals the risk-free rate plus a risk premium to compensate for the riskiness of the asset. Thus, the expected spot price at maturity exceeds the forward price because of a risk premium attached to holding the asset. Note that by contrast, the cash and carry arbitrage that determines the forward price is risk-free.

In general, Eq (3.1) contains two other terms which impact the forward price. But these also impact the expected spot price at maturity in a similar way. For example, if the asset earns an income, then the forward price is reduced, but the expected spot price at maturity is also reduced. (The expected spot price at maturity plus the income from the asset equals the current spot price plus the expected return on the asset.) Thus commonly speaking, it is true that the difference between the expected spot price at maturity and the forward price consists of a risk premium.

3.9 BACKWARDATION AND CONTANGO

In the simplest case of a non-income earning financial asset, the forward price is essentially the current spot price plus interest costs. A relationship where the forward price exceeds the spot price, is known as contango while the reverse relationship is known as backwardation. All non-income earning financial assets are always in contango. In terms of Eq (3.1), storage costs and income are both zero and the present value of the forward price equals the current spot price. As long as interest rates are positive, the forward price exceeds the current spot price. As negative interest rates are so rare as to be little more than an intellectual curiosity, we can say that these assets are always in contango.

Eq (3.1) tells us that the only way for an asset to be in backwardation is for it to have a large income yield. (Storage costs only push the asset further into contango.) Theoretically, it is possible for a financial asset to have an income yield that is large enough for it to overwhelm the borrowing cost and push it into backwardation. One example would be a forward contract on a high-yielding currency. Whenever the interest rate on a foreign currency exceeds the interest rate on domestic currency, the forward price of the foreign currency is less than the spot price. This can be identified as a case of backwardation. Another example would be a bond that has a coupon that is much larger than the current interest rate.

The most interesting examples of backwardation are those that are caused by a large convenience yield. Whenever a commodity is in backwardation, we know its convenience yield is very large. In fact it must be large enough to exceed the sum of the borrowing and the storage cost. When a commodity moves from contango to backwardation, it often signals market perception of a heightened probability of supply disruption, or demand shocks that cannot be met out of existing inventories. When such a commodity moves from backwardation to contango, it is a signal that inventories are now at comfortable levels to deal with the anticipated supply or demand shocks.

Another interesting case of backwardation arises when short sellers are caught on the wrong foot and are forced to deliver the asset that they have sold short. Their attempt to borrow or buy the asset in the spot market, causes the spot price to rise sharply. Since this steep rise is believed to be temporary, the forward price may not change much and the asset can move into backwardation.

Thus, the extent to which an asset is in backwardation or contango often tells us a lot about the nature of the market, about the magnitude of stocks being carried and attempts to manipulate the market. Backwardation and contango can be observed not only between spot and futures prices, but also between short maturity futures and long maturity futures.

3.12 | Derivatives and Risk Management

Chapter Summary

The theoretical forward price is the spot price plus net cost of carry where 'net cost of carry' equals interest costs *plus* storage costs *less* income (dividend or coupons). When we present value everything, we get the formula

$$PV (forward price) = spot price + PV (storagecosts) - PV (income)$$
$$PV (F) = S + U - I$$
(3.12)

In this expression, F is the forward price, S is the spot price, U is the present value of storage costs and I is the present value of the income from the asset. Rewriting this in terms of continuously compounded rates, we get the alternative formula:

$$F = e^{(r-q+u)T} S (3.13)$$

where r is the continuously compounded interest rate, q is the equivalent continuous dividend yield and u is the equivalent continuously incurred storage cost as a fraction of asset value.

The forward or futures price does not equal the expected spot price at maturity. The difference is due to the risk premium required to hold the asset.

If the forward price deviates from the theoretical price, there is an arbitrage opportunity. If the forward price is too high, cash and carry arbitrage can be used to exploit the opportunity by buying the asset spot and selling it forward. If the forward price is too low, reverse cash and carry arbitrage can be used to exploit the opportunity by selling (or shorting) the asset spot and buying it forward. These arbitrage possibilities are limited by constraints on borrowing or short selling.

The forward price must be distinguished from the value of a forward contract. The value f of a long forward contract with delivery price K is given by f = PV(F - K) where F is the current forward price for the residual maturity of the contract. In the absence of any income, yield or storage costs, this can be rewritten as $f = S - e^{-rT}K$ where t is the residual maturity of the contract and r is the continuously compounded interest rate.

A relationship where the forward price exceeds the spot price is known as contango, while the reverse relationship is known as backwardation. Non-income earning financial assets are always in contango. Assets that have a high income yield and commodities that have a high convenience yield can be in backwardation.

Suggestions for Further Reading

Much of the theory discussed in this chapter was developed in the early part of the twentieth century and the theory was more or less complete by the middle of the century. One of the classic papers of that era is:

Working, Holbrook (1949) "The Theory of Price of Storage", *The American Economic Review*, 39(6), 1254–1262.

The relationship between forward and futures prices is discussed in:

Cox, JC, Ingersoll, JE and Ross, SA (1981) "The Relation Between Forward Prices and Futures Prices", Journal of Financial Economics, 9(4), 321–346.

Problems and Questions

- 1. The shares of XYZ are trading at Rs 1,850 and no dividends are expected over the next month. What is the fair value of a one month future on this stock
 - (a) if the one month interest rate is 5% p.a on simple interest basis?
 - (b) if the one month interest rate is 5% p.a continuously compounded?
- 2. The Nifty index is trading at 4,100. Average dividend yield on the index is expected to be 0.6% p.a over the next quarter. The interest rate is 7.5% p.a continuously compounded. What is the fair price of the 3 month Nifty futures?

Cost of Carry Model for Futures and Forwards | 3.13

- 3. The one year interest rates in India and US are 6.5% and 4.5% p.a. semi-annually compounded. If the spot exchange rate is Rs 42.50/\$, what is the fair value of the one year forward exchange rate?
- 4. Storage costs of a particular commodity are Rs 500 per tonne. The spot price is Rs 85,000 per tonne and the interest rate is 5.25% p.a continuously compounded. If the commodity produces no income and no convenience yield, what is the fair value of the two month future?
- 5. In the above example, if the future is trading at Rs 84,000 what is the implied convenience yield?
- 6. The shares of ABC are trading at Rs 150. A dividend of Rs 35 has been announced and is payable after 35 days. The three month interest rate is 6% p.a. on simple interest basis. The three month future is trading at Rs 120. The three month interest rate is trading at Rs 120. If a person buys 1,000 shares of ABC and implements cash and carry arbitrage, what would her profits be at the end of three months? What would be her rate of return on the original investment?
- 7. The shares of Alpha Company are trading at Rs 750. No dividends are expected over the next three months. The three month interest rate is 5.45% continuoulsy compounded. The three month future is trading at Rs 755. A person implements a reverse cash and carry arbitrage by selling his holding of 5,000 shares and buying them back in the futures market. How much profit does he make at the end of three months?

Chapter Four

Objectives

Risk Management Using Futures and Forwards

This chapter looks at the use of futures and forward contracts in managing risks assuming that the company has decided to hedge the risks. The discussion of whether and why firms should hedge is the subject matter of the next chapter. The first step in hedging is the identification and measurement of risks, followed by the selection of the hedging instrument and determining the magnitude of the hedging required. The hedge then needs to be reviewed and adjusted over the course of its life. Many of these decisions have to be made with care when using futures or forward contracts that are not perfect hedges for the risks being managed.

4.1 PERFECT HEDGES AND KNOWN EXPOSURES

Consider a simple example of hedging a known exposure. Let's say, a company called XYZ has imported some machinery from the United States and the \$100,000 invoice is due for payment in 90 days. XYZ is worried that the US dollar may appreciate against the Indian rupee during this period and therefore goes to its bank for forward cover. The bank agrees to sell US dollars, 90 days forward, at a price of Rs 45.20 per \$.

This hedge completely eliminates the exchange risk for XYZ. On the due date it has to pay Rs 4.52 million to the bank, which has to sell \$100,000 at the aforesaid price. XYZ can then use \$100,000 to settle the invoice.

The transaction ensures complete certainty to XYZ about the rupee price of the machinery that it has imported. Whatever may happen to the exchange rate over the next 90 days, XYZ is sure that its total cost would be Rs 4.52 million—neither more nor less.

It is important to remember that in retrospect, XYZ may or may not be happy that it hedged the exchange risk. If on the due date, the exchange rate falls to Rs 45.10 per \$, XYZ would certainly wish that it had not bought the forward contract, as the machinery would have cost it merely Rs 4.51 million instead of Rs 4.52 million (that it now has to pay the seller under the forward contract). On the other hand, if the exchange rate at maturity rises to Rs 45.30 per dollar, XYZ would be very happy about its forward cover, as in the absence of that cover, the cost of the machinery would be Rs 4.53 million.

The forward hedge does not, therefore, ensure that XYZ gets a better deal. It only guarantees certainty about the amount that would be payable. Whether XYZ should or should not hedge to obtain this certainty is something that will be discussed at length in the next chapter. What this example shows is that once the decision has been made, then the forward contract provides a perfect hedge with no uncertainty at all, about the ultimate cost of the machinery.

4.2 BASIS RISK AND OPTIMAL HEDGE RATIO

The example in 4.1 above was very simple because the forward contract perfectly matched the risk exposure of the company. The liability was in US dollars and the forward contract was also for the same

4.2 | Derivatives and Risk Management

asset and at the same maturity date. In practice, such perfect hedges are not always available, especially when using futures contracts. As we have already seen, futures contracts are standardized contracts and therefore, they may not exactly match the risk that is being hedged.

Consider for example, a US company that plans to borrow \$100 million using a ten-year bond three months from now. It does not know what interest rate it would have to pay at that time and wishes to use a futures contract to protect itself from a rise in interest rates over the next three months. The Chicago Board of Trade offers futures on 10-Year US Treasury Notes, but there is a difference between the borrowing cost of the US government and the borrowing cost of a company. The interest rate that the company has to pay is equal to the risk-free interest rate plus a credit spread to compensate for the risk of default.

If the credit spread remains unchanged, then a 0.5 percent fall in the interest rate on the US Treasury will cause a 0.5 percent fall in the interest rate on the company's borrowing cost. That is why the company might consider using the Treasury futures to hedge its borrowing cost. However, the relationship is not perfect yet. For example, if there is a recession in the US economy, Treasury interest rates will tend to fall because of reduced economic activity and investment. However, the credit spread often tends to rise during a recession. This is because in a recession, corporate revenues and profits tend to fall, making the company less credit worthy. If the credit spread increases by 0.2 percent, then the company's borrowing cost falls by only 0.3 percent (0.5 percent fall in the risk-free rate minus the 0.2 percent rise in the credit spread).

How can the company take this into account while using Treasury futures to hedge its borrowing cost? The key idea is to adjust the magnitude of the hedge to reflect its imperfect nature. If the company were to hedge the borrowing of \$100 million fully, by selling \$100 million worth of Treasury futures, it would have been over-hedged. Suppose the modified duration^{1.} of the 10-year Treasury Note is 7 years, and assume for simplicity that the modified duration of the proposed debt issuance is also the same². When the Treasury yield falls by 0.50 percent, the price of the Treasury Note rises by 7×0.50 percent or 3.5 percent. On a \$100 million short hedge, there would be a loss of \$3.5 million. If the company's borrowing cost falls only by 0.3 percent (because of a 0.2 percent rise in the credit spread), then the present value of the debt servicing obligations of the company falls only by 7×0.3 % or 2.1 %. On the \$100 million borrowing that has been envisaged, the gain is only \$2.1 million. Thus the net hedged position suffers a loss of \$1.4 million. This of course, could work in reverse as well. If interest rates rise by 0.5 percent and the credit spread declines by 0.2 percent, then the short futures position would gain \$3.5 million and the borrowing cost would rise by only \$2.1 million for a net gain of \$1.4 million. But the whole idea of hedging is that the gains or losses on the hedge should offset the gains or losses on the position that is being hedged.

In the example that we have been discussing, 40 percent of any change in treasury interest rates is offset by an opposite change in the credit spread (40 % of 0.5 % = 0.2 %) and only 60 percent of the change is reflected as a change in the company's borrowing cost (60 % of 0.5 % = 0.3 %). If this pattern holds true in general, then there is a simple solution to the hedging problem, which is to hedge only 60

^{1.} The notion of modified duration is discussed in more detail later in this chapter. When interest rates rise (or fall) by 1 percent, the bond price falls (or rises) by d percent where d is the modified duration of the bond.

^{2.} If the company issues its bonds at par, the modified duration of these bonds will be slightly lower than that of the Treasury because of the slightly higher coupon. For the sake of simplicity, we will ignore this or assume that the bond is issued at a slight discount to achieve the same duration.

Risk Management Using Futures and Forwards | 4.3

percent of the risk exposure so that the company sells only \$60 million of Treasury Note futures to hedge the \$100 million planned borrowing. This way, when interest rates rise by 0.5 percent, the futures gains $7 \times 0.5 \% \times $60 \text{ million} = 3.5 \% \times $60 \text{ million} = 2.1 million . The company's borrowing cost rises by only 0.3 percent (because of a 0.2 percent fall in the credit spread) and the present value of the increase in the borrowing cost is $7 \times 0.3 \% \times $100 \text{ million} = 2.1 \% \times $100 \text{ million} = 2.1 million . The gains and losses are now equal and there is no over or under-hedging.

The problem thus boils down to determining the hedge ratio (60 percent in the above example). What is needed is an assessment of the relationship between changes in the value of the hedge instrument and changes in the value of the position being hedged. In most cases, the best way to make this assessment is to use statistical tools.

Figure 4.1 shows an illustrative scatter plot of change in corporate bond yields against treasury yields using hypothetical data. The horizontal axis shows weekly changes in the treasury yields in percent while the vertical axis shows changes in the corporate bond yields for the same week. We see that in almost all cases, both yields tend to rise together and fall together. There are a few exceptions where when the treasury yield rises slightly, the corporate bond yield falls a little or the other way around. It is true that any large rise or fall in treasury yields is accompanied by a change in corporate bond yields in the same direction. But even then, the relationship is not perfect. The oval figure shows five instances where the treasury yield rose by 0.09 percent. In these four instances, the corporate bond yield rose by amounts ranging from 0.03 percent to 0.07 percent for an average of 0.05 percent. While the average rise in corporate bond yields in this oval is about 50 percent of the rise in treasury yields, the individual instances range from about 30 percent to about 75 percent.



Figure 4.1 Illustrative scatter plot of change in corporate bond yields against treasury yields using hypothetical data. The horizontal axis shows weekly changes in the treasury yields in percent while the vertical axis shows changes in the corporate bond yields for the same week. We see that in almost all cases, both yields tend to rise and fall together. But the relationship is not perfect. The oval shows five instances where the treasury yield rose by 0.09 percent. In these four instances, the corporate bond yield rose by amounts ranging from 0.03 percent to 0.07 percent for an average of 0.05 percent.

4.4 | Derivatives and Risk Management

This is illustrative of what happens in real-world situations. There is a strong relationship between the hedge instrument and the risk being hedged, but the relationship is not perfect and can only be determined statistically. In Figure 4.2, the same hypothetical scatter plot of Figure 4.1 has been reproduced with a statistically best fitting straight line superimposed. The line is determined by linear regression and the regression equation is y = 0.6012 x-0.0017 where y is the corporate bond yield and x is the treasury yield.



Figure 4.2 Hypothetical scatter plot of Figure 4.1 with a statistically best fitting straight line superimposed. The line is determined by linear regression and the regression equation shows a slope of 0.6012. This means that on an average, a change in treasury yield causes a change in corporate bond yields which is in the same direction but only about 60 percent of the magnitude. The r-square is 0.8185 which means that nearly 82 percent of the change in corporate bond yields is explained by changes in treasury yields

The slope (or x coefficient) of the regression is 0.6012 which means that on an average, a change in treasury yield causes a change in corporate bond yields in the same direction but in only about 60 percent of the magnitude. The intercept or constant term is -0.0017 which is practically zero. This means that the regression line goes through the origin so that when the treasury yield is unchanged, the corporate bond yield is also unchanged on an average. The r-square or coefficient of determination of the regression is 0.8185 which means that nearly 82 percent of the change in corporate bond yields is explained by changes in treasury yields.

4.2.1 Hedge Ratio

What does this regression equation imply for the hedging decision? The first question is the magnitude of the hedge or the hedge ratio. The hedge ratio is given by the regression slope (or x coefficient) which

Risk Management Using Futures and Forwards | 4.5

in this example is 0.6012. The hedge ratio should be approximately 60 percent so that the \$100 million borrowing is hedged by a \$60 million short position in treasury futures.

4.2.2 Hedge Effectiveness

The second question is that of hedge effectiveness. The regression slope of approximately 0.60 means that on an average, the change in the corporate bond yield may be 60 percent of the change in treasury yields. Therefore the 60 percent hedge ratio ensures that, on an average, the change in the futures position will offset the change in the present value of borrowing cost. But this will be true only on the average. In an individual instance, the change in the corporate bond yield may be more or less than 60 percent of the change in treasury yields, and this will mean that the offset will not be exact. The hedged position would fine on average but could still make or lose money.

The hedge does reduce risk, but does not eliminate it. This is exactly what one would expect out of an imperfect hedge. The question is, what fraction of the risk is eliminated and what fraction remains? It is the r-square of the regression equation that provides the answer to this question. In Figure 4.2, the r-square is 82 percent implying that 82 percent of the risk has been eliminated, but that 18 percent of the risk remains. The hedge does reduce the risks dramatically but not completely. The residual risk that is not eliminated by the imperfect hedge is known as basis risk. This must be monitored carefully whenever an imperfect hedge is used.

4.2.3 Hedging and Variance Minimization

The statistical or regression approach can also be described in terms of minimizing the variance of the borrowing cost. A perfect hedge would fix the borrowing cost completely and thus reduce the variance to zero. Since the hedge instrument is not a perfect hedge, the variance cannot be brought down to zero, but can be reduced. The optimal hedge ratio minimizes the variance of the hedged position.

If y is the risk being hedged, x is the hedge instrument and h is the hedge ratio then the hedged position can be described mathematically as y-hx. The task is to determine the h that minimizes the variance of y-hx. Elementary statistics tells us that the variance of y-hx is given by

$$\operatorname{var}(y - hx) = \operatorname{var}(y) + \operatorname{var}(hx) - 2\operatorname{cov}(y, hx) = \sigma_y^2 + h^2 \sigma_x^2 - 2h\sigma_{x, y}$$
(4.1)

where σ_x and σ_y are the standard deviations of x and y, and $\sigma_{x,y}$ is the covariance between x and y. The optimal hedge ratio is determined by differentiating the right hand side of Eq (4.1) with respect to h and setting the derivative equal to zero. This gives us

$$2h\sigma_x^2 = 2\sigma_x, y \quad \text{OR} \quad h = \frac{\sigma_{x,y}}{\sigma_x^2} \tag{4.2}$$

This is exactly the same as the formula for the regression coefficient in linear regression. Thus, the variance minimizing hedge is given by the regression coefficient in a regression where the dependent or y variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the independent or x variable is the change in the values of the risk exposure and the

Using the formula for the covariance in terms of the correlation, ρ , and the standard deviations $(\sigma_{x,v} = \rho \sigma_x \sigma_v)$, we can rewrite Eq (4.2) as follows:

$$h = \frac{\sigma_{x,y}}{\sigma_x^2} = \frac{\rho \sigma_x \sigma_y}{\sigma_x^2} = \rho \frac{\sigma_y}{\sigma_x}$$
(4.3)

4.6 | Derivatives and Risk Management

Substituting this value of h into Eq (4.1) gives us the variance of the hedged position as

 $\operatorname{var}(y - hx) = \sigma_y^2 + h^2 \sigma_x^2 - 2h\sigma_{x,y} = \sigma_y^2 + \rho^2 \sigma_y^2 - 2\rho^2 \sigma_y^2 = \sigma_y^2 - \rho^2 \sigma_y^2 = (1 - \rho^2) \sigma_y^2$ (4.4) Since ρ^2 is the same as the r-square in simple linear regression, we find that the variance of the hedged position is less than the variance of the unhedged position by a factor of $1 - \rho^2$. The r-square (in percent) is thus the percentage reduction in risk achieved by hedging.

4.3 MEASURING RISK EXPOSURE

The discussion so far has assumed that the risk exposure is known. In practice, this is not always the case. Consider the following simple examples:

- 1. An exporter is worried about the exchange risk on export sales during the next quarter and is contemplating taking forward cover, but the quantum of export sales is not known with certainty. The management might have a target for export sales of Rs 1.25 billion, but the actual sales might be more or less than this.
- 2. An electrical goods manufacturer is worried about the price of copper and is considering hedging that risk in the futures market. The company has prepared a forecast of the consumption of copper during the next quarter. However, this forecast is subject to uncertainties about actual demand and production, the actual product mix and wastages, and operational efficiencies.

In practice, these uncertainties could be dealt with in either of two ways. The first possibility is that the company uses the forecast of the most likely magnitude of the risk exposure and hedges this exposure. The second approach might be to estimate the minimum level of risk exposure, hedge that and increase the hedge during the quarter, as the uncertainty gets resolved. In both of the above cases, the second approach may be justified on the ground that an over-hedging during a period of low demand (and therefore low revenues and profits) is a serious problem. On the other hand, an under-hedging during a period of high demand (and therefore high revenues and profits) is a much less serious problem.

A more serious problem of measuring exposure is when the nature of the exposure itself is not clear-cut. Consider the following examples:

- 1. A company is worried about the possibility of a large rise in transportation costs if the price of crude oil rises sharply. It knows that it has large exposure to crude oil prices, but finds it difficult to measure this risk exposure because the rise in transportation costs is not perfectly tied to crude oil prices. There is a lag between crude prices and the price of diesel. Not only that, the transport company may be forced to absorb a part of the rise in input costs in its profit margin, and may not pass it on wholly to its customers.
- 2. A company is worried about the exchange rate risk of imported raw material. However, the company has observed that the seller often moderates the rise in price of the raw material by reducing the dollar price when the dollar appreciates too much. Therefore, the extent of the risk exposure is difficult to measure.

These problems appear quite difficult, but there is a way of looking at them that allows one to use the methods developed in the previous section. If we do not think of the transportation cost problem as a crude oil price risk, but as using crude oil futures to hedge transportation cost risk, we know from the previous section that the optimal magnitude of the crude oil hedge can be determined by regressing

Risk Management Using Futures and Forwards | 4.7

transportation costs on crude oil prices and computing the regression slope. Therefore, we know the optimal crude oil hedge that is required.

This answer can also be interpreted a little differently as a measure of the company's crude oil risk exposure. In other words, we can interpret the regression coefficient as measuring crude oil risk and then interpret the hedge as a perfect hedge of the crude oil risk with crude oil futures. The interpretation that is chosen is largely a matter of taste and convenience, as the hedging decision is the same in either case. There are simply two different ways of interpreting this decision.

Similarly, in the second example, we could regard the risk, not as an exchange rate risk but as raw material price risk and regard currency forwards as an imperfect hedge of this risk. The optimal hedge is then given by the regression coefficient. Again, this hedging decision can then be reinterpreted as a perfect hedge of currency risk with currency forwards. In this case, the regression coefficient becomes a tool to measure the exchange risk exposure. These are just two different ways of interpreting the same hedging decision.

The exchange rate example can be extended to illustrate the advantage of the risk measurement interpretation. A company might import several different raw materials and components, and might also export some of its output. If it wants to know the net impact of all this, it could regress its profits on the exchange rate and use the regression coefficient as an estimate of its net exchange rate risk. In this situation, it is far more natural to treat the regression as a tool to measure the exchange rate risk exposure. In international financial management, it is common to refer to this kind of risk exposure as the 'operating' or 'economic' exposure to exchange rate risk. It is also common to use regression as a tool to measure this exposure.

It is moreover possible to have two levels of regression-based estimates in the same hedging decision. For example, under certain situations, a company might estimate its economic exposure to exchange rate risk over a long time horizon (say five years) using a regression or similar approach. It may then decide to hedge this risk using short maturity (say six month) currency forwards. There is now a basis risk because of the differing maturities of the risk and the hedge instrument. Again, an optimal hedge ratio may be determined using regression analysis.

4.4 USE OF BETAS IN INDEX FUTURE HEDGES

Though the principle of using the regression coefficient as the optimal or variance minimizing hedge ratio is a very general one, some special cases are of unique importance.

Consider a situation where an investor has bought a stock because he likes the company's technology and its management. He expects the company to do well but realizes, that if the broader market falls, he could end up losing money.

He can protect himself against this risk by using index futures. The analysis made earlier in this chapter tells us that the optimal hedge ratio is given by regressing the returns of the stock on the market returns and computing the regression coefficient. This regression coefficient is well known in portfolio theory and corporate finance, as the beta coefficient. Thus, the rule is that the magnitude of the index futures hedge is given by the beta of the stock. Since the beta can be greater than one or less than one, this is also an example where the rupee size of the hedge may be larger than the rupee size of the exposure being hedged. If for example, the stock that the investor has bought is a technology stock that has a beta of 1.5. Also suppose that he has bought 10,000 shares at a price of Rs 1,200 per share.

To compute the magnitude of the index futures hedge that is required, we proceed as follows. First of all, the investment that he has made is Rs 12 million (10,000 times Rs 1,200). The beta is 1.5, so the size

4.8 | Derivatives and Risk Management

of the hedge must be 1.5 times the position or $1.5 \times \text{Rs}$ 12 million = Rs 18 million. She will have to sell Nifty futures for an aggregate notional value of Rs 18 million. To determine the number of contracts required, we proceed as follows. For the Nifty futures, in 2006 the index multiplier was 100. Suppose the Nifty index value is 3,500. Then each Nifty futures contract has a notional value of Rs 0.35 million. To produce a hedge of Rs 18 million notional value, the number of contracts required is 18/0.35 = 51.43 contracts. As it is not possible to sell a fractional number of contracts, the hedge will consist of selling 51 Nifty contracts.

To see how this works, imagine that the Nifty index falls by 1 percent or 35 points. Since the index future has a multiplier of 100, each short Nifty contract produces a gain of Rs 3,500. The 51 short contracts produce a gain of Rs 178,500. Since the stock has a beta of 1.5, the 1 percent fall in the Nifty is expected to produce a 1.5 percent fall in the stock price. Since the stock price is Rs 1,200, it is expected to fall by Rs 18. The holding of 10,000 shares will lose out on Rs 180,000. The two offset each other almost completely. The slight difference of Rs 1,500 is because of the rounding of the optimal hedge from 51.43 contracts to 51 contracts.

4.5 USE OF MODIFIED DURATION IN INTEREST RATE HEDGES

We have already encountered the use of modified duration earlier in this chapter. The concepts of duration and modified duration are fundamental in valuing fixed income securities and managing interest rate risk. Duration of a bond is defined as the weighted average time when the present value of cash flows from the bond is received. It is computed by listing down all the cash flows (coupons and redemption values), computing the present value of each cash flow, multiplying the present value by the number of years after which it is received, summing up all these products, and dividing the sum by the price of the bond (or equivalently, the sum of the present values). Duration is used as a better measure of the term of a bond than its maturity.

A slight modification of duration is more useful in hedging-related decisions. When the duration is

divided by $1 + \frac{ytm}{2}$, where ytm is the semi-annually³ compounded yield to 2 maturity, we get what is

known as modified duration which is a measure of the sensitivity of the bond price to interest rates. More precisely, when interest rates rise (or fall) by 1 percent, the bond price falls (or rises) by d percent where d is the modified duration of the bond.

An example will make this clear. Consider a two year bond paying 8 percent coupons on a semiannual basis and suppose that the *ytm* of this bond in the market is 7.5 percent. The market price of this bond would then be 100.91 as computed below:

Time (years)	0.5	1.0	1.5	2.0
Coupon	4	4	4	4
Redemption				100
Total cash flow	4	4	4	104
Discount factor at 7.5 %	0.9639	0.9290	0.8954	0.8631
Present value	3.86	3.72	3.58	89.76
Bond price	3.86+3.72+3.58+89.76=100.91			

^{3.} In India, US and the UK, bond yields are by convention quoted on a semi-annually compounded basis. This convention leads to the term *ytm*/2 in the divisor in the formula for modified duration. If yields are annually compounded as in most of Europe, the divisor is simply 1+ *ytm*.
Risk Management Using Futures and Forwards | 4.9

The duration of this bond is 1.89 years as shown below. The duration is less than the maturity of two years because while the redemption is two years away, some coupon cash flows are received earlier and this pulls down the weighted average.

Time (years)	0.5	1.0	1.5	2.0
Present value	3.86	3.72	3.58	89.76
Time x Present Value	1.93	3.72	5.37	179.52
Sum of Time x PV	1.93+3.72+5.37+179.52=190.54			
Duration		190.54/100.91=1.8	39	

The modified duration is then seen to be $\frac{1.89}{1 + \frac{0.075}{2}} = 1.82$. This means that if the interest rate rises by

0.10 percent, the bond price would fall by approximately 1.82 times 0.10 percent or 0.182 percent. Since the price of the bond was 100.91, it should fall by 100.91×0.182 percent = 0.183. The new price should therefore be 100.73. When we recompute price the bond at a ytm of 7.6 percent, we find that the price does indeed fall to 100.73 as the following table shows:

Time (years)	0.5	1.0	1.5	2.0
Coupon	4	4	4	4
Redemption				100
Total cash flow	4	4	4	104
Discount factor at 7.6 %	0.9634	0.9281	0.8941	0.8614
Present value	3.85	3.71	3.58	89.59
Bond price	3.85 + 3.71 + 3.58 + 89.59 = 100.73			

Suppose an investor is holding \$25 million (face value) of these bonds and is trying to hedge the interest rate risk of this position using two-year Treasury Note Futures. Each futures contract on the two-year Treasury Note has a notional value of \$200,000.

As discussed in Chapter 2, Treasury Futures are often regarded as futures on the CTD bond. Suppose that the CTD has a price of 100.15, a modified duration of 1.75 years and a conversion factor of 0.9750. The conversion factor of 0.9750 means that when a seller delivers \$200,000 face value of the CTD against the futures contract, the buyer has to pay only 0.9750 times the futures' price.

This implies that the futures price at expiry must be $\frac{1}{0.9750} = 1.0256$ times the price of the CTD at expiry. In other words, when the CTD changes by 1 percent, the futures price changes by 1.0256 percent. Since the CTD has a modified duration of 1.75, its price changes by 1.75 percent when interest rates change by 1 percent. The futures price then changes by 1.0256 times of this so that every 1 percent change in the interest rate causes the futures price to change by 1.0256 × 1.75 % = 1.79 %.

On the other hand, when interest rates change by 1 percent, the investor's portfolio changes by 1.82 percent because its modified duration is 1.82. In other words, a 1.79 percent change in the futures price is associated with a 1.82 percent change in the value of the investor's portfolio. Put differently, every

1 percent change in the futures price causes a change of $1.00 \% \times \frac{1.82}{1.79} = 1.017$ percent in the value of the

portfolio. According to the discussion earlier in this chapter, the hedge ratio should be 1.017. The mag-

The McGraw·Hill Companies

4.10 | Derivatives and Risk Management

nitude of the hedge must be 1.017 times the value of the investor's portfolio. At a price of 100.91, the \$25 million (face value) of bonds has a market value of \$25.22 million. The hedge must therefore be \$25.22 million times 1.017 or \$25.66 million. Since each futures contract is for \$200,000 or \$0.2 million, the number of futures contracts required is 25.66/0.2 = 128.3. The actual futures position will be 128 contracts since a fraction of a contract cannot be traded.

4.6 HEDGING CURRENCY RISK OF EQUITY POSITIONS

Is it correct for a European investor to hedge a \$25 million portfolio of US equities against currency risk by selling \$25 million worth of dollars forward? The notion of the optimal hedge ratio provides an interesting answer to this question. To compute the hedge ratio, it is necessary to determine the average change in the euro value of the foreign equity portfolio when the dollar depreciates by 1 percent against the euro. If the dollar price of the stocks do not change while the dollar falls by 10 percent then the euro value of the portfolio does indeed fall by about 1 percent. But is it reasonable to assume that the dollar value of US stocks is unaffected by the movement in exchange rates? If the decline of the dollar reflects some problems in the US economy, then it is likely that the US stock market will also fall in dollar terms. Then the fall in the euro value of the portfolio would be greater than 1 percent.

On the other hand, suppose that the \$25 million portfolio consists largely of US companies with large exports to the rest of the world. A decline in the dollar makes US exports more competitive in the rest of the world. US exporters may then experience a rise in revenues and profit margins. Their stock price in dollars could surge in response to this rise in profitability. It is even conceivable that the dollar stock prices rise so much that the euro value of the portfolio rises despite the fall in the dollar.

In this case, it is difficult on theoretical grounds to make a good assessment of the likely impact of an exchange rate change on the value of the portfolio measured in home currency (euros). The statistical technique of regression then becomes very useful. It is possible to regress the euro returns of each stock on changes in the exchange rate to find its sensitivity to the exchange rate and therefore the hedge ratio.

For some companies that predominantly serve the domestic US market, this hedge ratio could turn out to be significantly above unity. For some large exporters, the hedge ratio may be well below unity and in exceptional cases, even be negative. Thus the optimal hedge ratio for the portfolio would depend on the precise composition of the portfolio.

4.7 TAILING THE HEDGE

Throughout this chapter, we have treated forward contracts and futures contracts as being equivalent for hedging purposes. We have also implicitly assumed in the previous chapter that the futures price is the same as the forward price. These assumptions need justification.

While forward and futures contracts differ in terms of a number of features like standardization, novation, clearing, and liquidity, the key difference from a valuation or hedging point of view is the daily mark to market. A change in the underlying causes an immediate gain or loss in the futures contract because of this mark to market. In the case of the forward contract, however, the gain or loss arises only at maturity when the transaction is settled. This is a non-trivial difference in financial terms because there is a difference in present value between the loss of a rupee in the present and the loss of a rupee at the end of six months.

To understand the implications of this difference, let us consider an example. Suppose a US company

Risk Management Using Futures and Forwards | 4.11

has to repay a $\in 10$ million liability at the end of three months. If the finance manager decides to hedge this using forwards, it is clear that she should simply buy a forward contract of $\in 10$ million. If she instead wants to hedge this risk in the currency futures market, would she be right in simply buying $\in 10$ million of futures? The answer is provided by the hedge ratio that we have discussed extensively in this chapter.

If the euro falls from 1.25 to 1.24, she would suffer an immediate mark to market loss of 0.01×10 million or 100,000 on the futures position. There is an offsetting gain in the euro liability of 100,000 but this gain arises only at the maturity of the forward contract. If the interest rate is 4 percent continuously compounded, then discounting the gain for three months gives a present value of only 100,000

 $\times e^{-0.04 \times \frac{3}{12}} = \$100,000 \times 0.99005 = \$99,005$. For the hedge to be correct, the loss on the futures contract should also be \$99,005 and not \$100,000. To achieve this, she must buy futures not for €10 million but for €9.90005 million. The mark to market loss on this long futures position when the euro moves from \$1.25 to \$1.24 would be \$99,0005 which is the same as the present value of the gain on the euro liability. The optimal hedge of €9.90005 million can be characterized in a different way as the present value of the euro liability:

 $10 \text{ million} \times e^{-0.04 \times \frac{3}{12}} = 10 \text{ million} \times 0.99005 = 9.9005 \text{ million}.$

There is a catch here. This adjustment of the hedge quantity is not something that can be decided at the time of inception and then left unchanged. With every passing day, the time for maturity comes down and the present value of the euro liability goes up. For example, consider the situation where the investor looks at the hedge again at the end of one month. At this point, the maturity of the euro liability is only 2 months away and assuming that the interest rate remains at 4 percent, the optimal hedge is \$10 million

 $\times e^{-0.04 \times \frac{2}{12}} = \$10 \text{ million} \times 0.99336 = \9.9336 million . She would then need to buy €33,100 of futures during the first month to increase the hedge from €9.900,500 to €9.933,600. After another month, with

maturity only a month away, the optimal hedge is \$10 million $\times e^{-0.04 \times \frac{1}{12}} = $10 million \times 0.99667 = $9.9667 million. She should, therefore, buy another €33,100 of futures during the second month. At expiry when the mark to market coincides with the maturity of the euro liability, the size of the hedge should be equal to €10 million, the value of the forward contract. So she should therefore buy another €33,300 of futures during the third month.$

Ideally of course, the hedge must be increased each day during the course of the month. In practice, however, this perfection is not possible because of trading costs and the minimum size of one futures contract. For example, the size of the Euro contract at the Chicago Mercantile Exchange (CME) is $\in 125,000$ and the size of the E-mini Euro contract is $\in 62,500$. The question of buying $\in 33,100$ or $\in 34,000$ of futures contracts does not therefore arise. It may be possible to adjust the size of hedge only once during the entire three month period. Moreover, in practice, small deviations from the optimal hedge are not significant and many traders may revise the hedge only at monthly or quarterly intervals.

This idea of a futures hedge which is smaller than the risk exposure at inception, but increases steadily to equal the exposure at maturity is known as tailing the hedge. In the case of short maturity exposures, the impact of tailing the hedge is not very important. In the above example of a three month liability, the adjustment amounts to only about one futures contract. The difference becomes much larger for long maturity liabilities. If in the above example, the euro liability had a maturity of ten years and the interest rate for this maturity was 4 percent continuously compounded, the optimal hedge would be:

4.12 | Derivatives and Risk Management

 $10 \text{ million} \times e^{-0.04 \times 10} = 10 \text{ million} \times 0.67032 = 6.7032 \text{ million}$, which is only about two-thirds of the euro liability. This position would perhaps be revised every quarter or so, that it steadily rises toward the $\epsilon 10$ million size of the euro liability. Tailing the hedge is absolutely critical in this case to maintain hedge effectiveness.

It should be emphasized that tailing the hedge is finally, simply an application of the theory of the optimal hedge ratio. When the futures moves by \$1, the change in the value of the euro liability (in present value terms) is only 0.99005 in the case of the three month maturity and 0.67032 in the case of the ten year maturity. The hedge ratio is therefore 0.99005 and 0.67032 respectively in these two cases. Therefore a $\in 10$ million risk exposure requires a futures hedge of only $\notin 9,900,500$ or $\notin 6,703,200$ respectively.

The analysis of tailing the hedge is also useful to establish that the futures price and the forward price are identical. Consider the earlier comparison of hedging a three month €10 million euro liability with forward contract and futures.

Let us first analyse the €10 million forward contract. Gain or loss happens on this, only on maturity and is equal to \$100,000 for every deviation of one cent between the forward price at inception and the spot price at maturity. The present value of a \$100,000 gain or loss discounted from the end of the third

month to inception is $100,000 \times e^{-0.04 \times \frac{3}{12}} = 100,000 \times 0.99005 = 99,005$. Therefore, the present value of the gains or losses on the $\notin 10$ million forward contract is 99,005 times the difference measured in cents between the forward prices at inception and the spot price at expiry, or 9,900,500 times the difference between these prices measured in dollars. Obviously, the gains or losses of the forward contract do not depend on the prices at intermediate dates between inception and maturity.

Something similar has to be first established for the tailed futures strategy. In this strategy, as already seen, the futures position is $\notin 9.9336$ million at the end of the first month, $\notin 9.9667$ million at the end of two months and $\notin 10$ million at expiry. Let us begin by assuming that the mark to market in the futures contract is monthly, instead of daily. This is simply to allow for the convenience of exposition, as the same logic extends to the actual daily mark to market scenario.

The present value of the mark to market gains or losses of a tailed hedge is determined solely by the cumulative change in the futures price from inception to expiry. It does not depend on monthly fluctuations at all. For example, if the futures price changes from \$1.25 at inception to \$1.22 at expiry, we can calculate the present value of all the mark to market losses of the tailed hedge without knowing the intermediate prices. For example, the futures price might have gone from \$1.25 to \$1.24 to \$1.23 and then to \$1.22 at the end of the first, second, and third months, or it might have risen to \$1.26 after one month and further to \$1.28 at the end of the second month, and dropped to \$1.23 in the last month. It might have dropped to \$1.22 in the first month itself and remained steady at this level till maturity.

There are numerous ways in which the futures price can go from \$1.25 to \$1.22 over a three month period. What is important is that the present value of mark to market losses of the tailed futures strategy does not depend on the path the futures price takes. It depends only on the total cumulative change in the price, from inception to expiry.

Let us see why this is so. At the end of the first month, the size of the tailed futures is $\notin 9.9336$ million. A \$0.01 (one cent) rise (or fall) in the futures price during this month causes a mark to market gain (or

loss) of \$99,330 at the end of the month. The present value of this gain or loss is $\$99,336 \times e^{-0.04 \times \frac{1}{12}} =$ \$99,336 \times 0.99667 = \$99,005 At the end of the second month, the size of the futures is €9.9667 million

Risk Management Using Futures and Forwards | 4.13

and a one cent rise (or fall) in the futures price during this month causes a mark to market gain (or loss) $\frac{2}{2}$

of \$99,667 at the end of the month. The present value of this is gain or loss is $$99,667 \times e^{-0.04 \times \frac{2}{12}} = $99,667 \times 0.99336 = $99,005$. Thus, it does not matter whether the one cent move happens in the first or second month; in either case, it causes a present value gain or loss of \$99,005. Similarly, at the end of the last month, the size of the futures hedge is $\notin 10$ million and a one cent rise (or fall) in the futures price during this month causes a mark to market gain (or loss) of \$100,000 at the end of the month. The present value of this gain or loss is once again \$99,005 because:

 $\$100,000 \times e^{-0.04 \times \frac{3}{12}} = \$100,000 \times 0.99005 = \$99,005$. Since a one cent move in the futures price causes a mark to market gains or losses of \$99,005 in terms of present value, regardless of when the move takes place, the aggregate present value of gains or losses is \$99,005 times the cumulative change in the futures price in cents. The cumulative change in the futures price is the same as the difference between the futures prices at inception and at expiry. Moreover, at expiry, the futures price is the same as the spot price. Therefore, the above result can also be restated as follows: the aggregate present value of gains or losses on the tailed futures is \$99,005 times the difference (measured in cents per euro) between the futures prices at inception and the spot price at expiry. Alternatively, it is \$9,900,500 times the difference measured in dollars per euro between the futures prices at inception and the spot price at expiry.

It is useful to look at this result slightly differently in terms of the terminal value at expiry (at the end of three months) instead of the present value at inception. Shifting from present value to terminal value, changes the multiplier of \$9,900,500 to:

 $9.9005 \text{ million} \times e^{0.04 \times \frac{3}{12}} = 9.9005 \text{ million} \times 1.01005 = 10 \text{ million}$. Thus if the tailed futures strategy were implemented and all mark to market gains and losses were reinvested at the risk free rate, the terminal value at maturity of all these gains or losses would be 10 million times the difference between the futures prices at inception and the spot price at expiry. 10 million times the futures prices in dollars per euro is simply the futures price of €10 million. Similarly \$10 million times the spot prices in dollars per euro at expiry is simply the spot price at expiry of €10 million.

Thus the cumulative reinvested gains and losses of the tailed futures would be the same as the difference between the futures price and the spot price at expiry of $\in 10$ million. Suppose the futures price at inception is \$1.25 per euro. Then the futures price of $\in 10$ million is \$12.5 million. The tailed futures would produce a gain or loss at maturity, equal to the difference between \$12.5 million and the spot price at expiry of $\notin 10$ million. Suppose an investor buys three month risk-free bonds with a redemption value of \$12.5 million and also implements the tailed futures strategy, at maturity, the value of his portfolio including reinvested gains and losses would be simply equal to the spot price at maturity of $\notin 10$ million. He can thus liquidate his portfolio and use the proceeds to buy $\notin 10$ million at the then prevailing spot price. To buy bonds with a redemption value of \$12.5 million, he would have to pay:

\$12.5 million $\times e^{-0.04 \times \frac{3}{12}} = 12.5 million $\times 0.99005 = 12.3756 million. Thus, by using the tailed futures, he can ensure that an investment of \$12.3756 million now, is exactly sufficient to buy $\notin 10$ million at the end of three months.

The net effect of this transaction is exactly the same as buying a forward contract for €10 million at a forward price of \$1.25. This will require a payment of \$12.5 million at the end of three months. For this

4.14 | Derivatives and Risk Management

money to be available at the end of three months, he would need to invest \$12.3756 million immediately. Since the end result of futures and forward contracts are identical, the futures and forward prices should be the same. If for example, the forward price were \$1.26 while the futures price were \$1.25, a person could sell euros forward while simultaneously implementing the long-tailed futures strategy, reinvesting all gains and losses at the risk-free rate. At maturity, he would end up buying euros at \$1.25 from the tailed futures and selling the same euros at \$1.26 under the forward contract, making a risk-free profit of \$0.01 per euro. In an efficient market, therefore, forward and futures prices should be the same.

During the course of these discussions, we assumed that the interest rate remains unchanged. This is not strictly necessary. What is required is that the changes in interest rates are uncorrelated with changes in the prices of the underlying. This means that if the underlying is itself an interest rate related instrument or correlated with interest rates, the equivalence of forward and futures prices breaks down. This is important for fixed income derivatives with very long maturities (for example, long maturity interest rate swaps). In most other cases, futures and forward prices can be assumed to be the same.

4.8 ROLLING HEDGES

Yet another interesting application of the idea of the optimal hedge ratio is the use of futures of a maturity different from that of the risk being hedged. For example, if a company has contracted to supply crude oil to its customers over the next five years at pre-determined prices and wants to hedge this risk using crude oil futures. Now crude oil futures are quite liquid at three month maturity, but the liquidity dries up rapidly at longer maturities and five year futures simply do not exist. Can the five year supply contract be hedged with three month futures?

Let us go back to the cost of carry model of futures pricing discussed in Chapter 3. Imagine for a moment, that there were no convenience yields or storage costs in crude oil. Then the three month futures would differ from the spot price by the interest cost for three months and the five year futures would differ from the spot price by the interest cost for five years. The difference in price between a three month and five year futures would be only the interest cost from the end of the third month to the end of the fifth year. If we assume that interest rates do not change, then the two prices would be perfectly correlated with each other because the price changes in both contracts can be perfectly predicted from the change in the spot price. The r-square of the regression of one price change on another would be unity, and the five year supply agreement could be perfectly hedged with even a three month contract. It must be ensured that tailing the hedge is correctly implemented as the risk exposure entails a long maturity.

Several strong assumptions have been made in arriving at this result. But in many financial futures, it is indeed possible to hedge relatively long maturity risks with short maturity rolling hedges. The hedge effectiveness is less than what it would be with matched maturity hedges, but the higher liquidity of short term futures compensates for this.

Crude oil, however, is very different in that, it is a commodity characterized by a large and highly volatile convenience yield. This means that sometimes the spot price and the three month futures move a lot while the five year futures remain largely unchanged. For example, suppose there is a short term supply disruption that leads to a temporary shortage of crude. Then the spot as well as the short maturity futures may rise sharply. Long maturity futures may be unaffected because the disruption is expected to be eliminated in a few months. This would reduce the correlation between the two prices.

Risk Management Using Futures and Forwards | 4.15

The three month futures would then become a much poorer hedge of the five year supply contract. This analysis helps us arrive at a clear analysis to understand the Metallgesselschaft disaster discussed in Chapter 1. Not only was the hedge ratio of unity used by Metallgesselschaft inappropriate, because of the need to tail the hedge and to account for the basis risk, the volatile convenience yield made the hedge itself poor.

Chapter Summary

The purpose of hedging is to reduce uncertainty about future costs, revenues or cash flows. Hedging does not necessarily increase expected profits, it only reduces its variance. In practice, most hedges are imperfect because the hedging instrument has a different underlying, a different maturity, or has different timing of cash flows. In these situations, the optimal or variance minimizing hedge ratio is the regression coefficient obtained by regressing the change in the value of the risk exposure on the change in the value of the hedge instrument. The effectiveness of the hedge is given by the r-square of the regression. The hedge ratio for hedging stock portfolios against market wide movements is given by the beta of the portfolio. For hedging a bond portfolio against interest rate risk, the hedge ratio depends on the duration of the portfolio and the duration of the CTD bond underlying the bond future as well as its conversion factor.

When hedging distant maturity risks using futures, the gains or losses of the risk arise at maturity, while the futures produce daily mark to market cash flows. To ensure the correct hedge in terms of present value, the futures position must be equal to the present value of the risk exposure. The hedge ratio will be well below unity at inception and will gradually rise to unity as maturity approaches. If the futures contract is for shorter maturities than the maturity of the risk exposure, the maturity mismatch will also require an adjustment in the hedge ratio. This basis risk will also reduce the hedge effectiveness.

Suggestions for Further Reading

The need for tailing the hedge while using futures contracts is discussed in:

Figlewski, S., Landskroner, Y. and Silber, W.L. (1991), "Tailing the Hedge: Why and How", *The Journal of Futures Markets* (1986–1998), 11(2), 201–212.

The notion of the optimal hedge ratio is elaborated in the following paper:

Stein, Jerome L. (1961) "The Simultaneous Determination of Spot and Futures Prices", *The American Economic Review*, 51(5), 1012-1025.

Figlewski, S and Kon, SJ (1982) "Portfolio Management With Stock Index Futures", *Financial Analysis Journal*, Jan-Feb 82, 52–59.

Problems and Questions

- 1. A fund manager wants to reduce the beta of her portfolio (of Rs 2.5 billion) from 1.1 to 0.95. If Nifty futures are trading at 3,900 and the multiple for the Nifty future is 100, how many Nifty futures should she buy or sell?
- 2. A bank has entered into a forward rate agreement with a customer wherely it has agreed to lend \$100 million to the customer in January 2011. How many Eurodollar futures must the bank buy or sell in January 2007 to hedge its interest rate risk?

(Recall that in Eurodollar futures, each basis point change in Libor is worth \$25). Do not forget to tail the hedge.

The McGraw·Hill Companies

4.16 | Derivatives and Risk Management

- 3. In the above problem, if the bank reviews its hedge in January 2008, what change will it need to make?
- 4. A company finds that there is a correlation of 0.80 between the aluminium price at the London Metal Exchange (LME) and the aluminium price in India. The volatility of the prices at LME and in India are roughly the same. If its expected consumption of aluminium is 5,000 tonnes, what is optimal hedge while using LME futures?
- A company obtained the following regression equation when it regressed the profits of its Indonesian subsidiary on the exchange rate.

$$Y = 12 + 25,000 X$$

where *Y* is the profits measured in millions of Indian rupees and *X* is the exchange rate (Indian rupees per Indonesian rupee).

- 6. A fund manager finds that the modified duration of her US bond portfolio is 5.3 years while the modified duration of the benchmark against which she is evaluated is 5.8 years. If her portfolio is \$20 million, how many US T-bond futures should she buy or sell to achieve the same duration as the benchmark? Assume that the cheapest to deliver bond has a price of 101.15, a modified duration of 8.75 years and a conversion factor of 0.9950.
- 7. Consider the Metallgesselschaft example discussed in the text. Assume that there is a squared correlation (*r*-square) of 0.80 between changes in the three months futures price of crude oil and changes in the five year forward price. If the three month futures are twice as volatile as the five year forward, then how many barrels of crude should Metallgesselschaft have shorted in the futures for every 1,000 barrels that it sold in five year forward contracts? Ignore tailing of the hedge.
- 8. In the above problem, what would be the impact of tailing the hedge if the interest rate is 5% continuously compounded.

Chapter **Five**

How and Why do Firms Hedge?

Objectives

This chapter discusses the costs and benefits of hedging and explains the reasons why companies choose to hedge against some risks. Once a decision to hedge has been made, there are a number of hedging instruments available to carry forward the process. This chapter highlights the advantages and limitations of financial hedges like derivatives and discusses some of the non-financial hedging instruments that may be appropriate under certain circumstances. Chapter 21 continues with relevant points discussed in this chapter, though at a more advanced level.

5.1 COSTS AND BENEFITS OF HEDGING

There are two types of costs associated with hedging. First there are the direct costs of hedging, for example, when buying and selling futures, the hedger incurs commissions and brokerage costs. Some derivatives for example, options (that we will discuss later in this book) also have an upfront cost or premium. In addition to these clearly quantifiable costs, there are also a number of costs incurred while managing the hedge. Risk exposures have to be measured, the hedge ratio has to be determined, and the hedges have to be periodically monitored. Futures contracts give rise to daily mark to market cash flows which have to be managed. The initial margin requirements for futures contracts also have a cost in terms of interest foregone on this margin. In case of forward contracts too, this cost is present, but is less explicit. For example, when a bank offers a forward contract to its customer, it regards the transaction as a 'non fund' exposure and implicitly reduces it from the aggregate funded credit that it is willing to extend to the client.

Another important type of cost that is involved in hedging is the opportunity cost that we discussed in the previous chapter. An example of a company known as XYZ was analysed in which XYZ bought dollars forward at Rs 45.20 per dollar to hedge the cost of imported machinery. This hedge does protect the company from a rise in the value of the dollar and brings certainty about the rupee payment that it has to make. But this certainty comes at a cost. If the dollar declines to Rs 45.10 per dollar, XYZ would rather that it had not taken the forward cover at all, as the forward contract forces it to pay Rs 45.20 for the dollars that are available in the market at Rs 45.10 each. This is an opportunity cost of hedging.

Why would XYZ then hedge? One way of looking at the situation is that Rs 45.20 is obviously a price that XYZ can afford to pay while still making a profit on its investment, or it would not have been importing the machinery at all. Thus, the inability to buy dollars cheaper than Rs 45.20 only means that it does not make even more profits than it anticipated. On the other hand, if the dollar rose sharply, the machinery could become so much more expensive that it would no longer remain a profitable proposition. Thus, the benefit of hedging is to protect the company against losses and the cost is what it foregoes—the opportunity to make more profits than it ever anticipated. Many companies prefer this certainty of reasonable profits to the uncertainty of profits being significantly higher or lower.

Hedging is thus, primarily a device to achieve greater certainty about cash flows (for example, costs, revenues, and investments). This greater certainty allows the company to focus on its business, which it knows well without having to worry all the time about things like exchange rates, that it does not know

5.2 | Derivatives and Risk Management

too much about. This is broadly the idea of core competence—a textile company for example, should focus on its strengths in designing and manufacturing textiles and leave exchange rate forecasting to banks and currency speculators. This is an important benefit of hedging – it allows management time to be focused on business, instead of speculation.

The certainty achieved by hedging has other benefits too. Planning and control, for example are much easier when risks are hedged. In the example of XYZ company, hedging makes it much easier to plan for cash requirements as there is complete certainty on the cash needed to pay for the imported machinery. In the absence of hedging, XYZ would have to plan for an extra cushion of cash to cover the possibility of larger cash outflows when the dollar strengthens. Of course, the extra idle cash has a cost in terms of foregone interest earnings.

Hedging also makes it easier to evaluate the performance of different divisions. For example, a division that imports a lot of raw material or exports a lot of its output, will have large swings in its profitability simply because of changes in exchange rates. It can become quite difficult to separate the effects of exchange rate fluctuations from changes in operating efficiency and performance. If the division hedges all its imports and exports, the problem goes away because the profits of the division are no longer affected by currency movements. Hedging is however not the only way to deal with this problem. A well-designed performance evaluation system can adjust for currency changes and compute the true performance of the division. But the point is that hedging simplifies things and makes it easy for managers to understand what is going on. It, therefore, allows them to focus on their business.

While hedging might be a luxury when it comes to small risks, it becomes absolutely essential when the risks are potentially large enough to threaten the very survival of the company. In the example of XYZ company, the effect of a small rise in the exchange rate may only be a reduction in profitability, or perhaps a small loss. If the transaction was much bigger and the potential adverse price changes more severe, such that the potential loss could make the company bankrupt, then the situation would be very different. Bankruptcy has huge costs. Employees and managers lose their jobs, customers and suppliers are badly affected, and lenders also lose a lot of money. All of these stakeholders would insist that large risks be hedged so that they do not incur these losses. Many lenders will for example, insist that a company insure its factories and other major assets. In some extreme cases, lenders even pay the insurance premium to protect themselves against loss.

Bankruptcy costs go beyond the large fees charged by lawyers and accountants as part of the court supervised liquidation. A distress sale of assets might fetch much less than what the assets are worth to a running business. Many intangible assets may prove worthless in bankruptcy, for example, a half finished research project might have no value at all to an outside buyer. Similarly, the company's brand and goodwill might also be eroded in the process of bankruptcy.

For all these reasons, most companies prefer to hedge against risks that are large enough to threaten their survival. For smaller risks, the choice is often driven by the managerial attitude towards risk taking. These issues will be discussed further in Chapter 21.

5.2 HEDGING INSTRUMENTS

5.2.1 Insurance

Insurance is perhaps the best known method of protecting against risks. Insurance simply transfers risk to a specialized intermediary in exchange for an insurance premium. This intermediary (insurance company) must then find some way of managing the risk that it has taken on. Traditionally, insurance

worked best with risks like fire or theft where most of the risk is eliminated by diversification. The insurance company insures a large number of companies in several different geographical regions and can then be reasonably confident that the aggregate claims will be less than the premium charged. There is also a re-insurance market where an insurance company can in turn, transfer risks to even more specialized intermediaries, or share risks with other companies.

In recent years, insurance companies have expanded their range of products to cover various risks that are more financial in nature. Insurance products thus compete with derivatives in several cases.

The law and practice relating to insurance is quite specialized and there are several conditions to be fulfilled before an insurance company pays its claims. For example, it is often necessary to demonstrate that loss has occurred. If a claimant fails to disclose all relevant information to the insurance company, it might invalidate an insurance contract. In many cases, insurance claims are intensely litigated and payment is delayed until litigation is complete. By contrast, in most derivative contracts, payment is automatic, immediate, and without any pre-conditions.

5.2.2 Forward Contracts

A forward contract is among the oldest and simplest hedging devices. It is simply a purchase or sale transaction in which the price and other terms have been agreed upon, but the delivery and payment are postponed to a later date.

Forward contracts were originally negotiated between two parties, both of whom were trying to hedge the risk of an anticipated future purchase or sale transaction. A flour mill, for example, may enter into a forward contract with a wheat farmer. The miller is worried that when he needs to buy wheat, the price may be too high. The farmer is worried that the price may be too low at the time of the harvest when he is ready to sell. The forward contract between the two entities hedges the risks of both parties.

These days, forward contracts are typically offered by banks who act purely as intermediaries. For example, in the foreign exchange market, an exporter does not usually enter into a forward contract with an importer. Both of them enter into opposite forward contracts with a bank that stands in the middle and earns a profit spread. The exporter and importer do not have to search each other out and negotiate terms.

Forward contracts can be highly customized to the needs of the hedger and therefore with very little basis risk. They are easy to use because of the lack of margins and mark to market cash flows. As a result, forward contracts have become quite popular as hedging instruments.

The principal disadvantage of forward contracts is that these customized contracts can be expensive to set up and unwind. For some highly customized forward contracts, there may only be a few banks willing to offer a quote. The company is then not sure whether it is getting a fair deal while entering into the forward contract. For most routine forward contracts, however, this is not a problem because a large number of competing quotes may be available.

There is a more serious problem when for some reason, the hedger wants to cancel the forward contract. For example, a change in business strategy may mean that the risk exposure no longer exists and therefore the hedger needs to terminate the hedge as well. The simplest way to do this is to negotiate a price with the bank or the other counter party to cancel the forward contract. The hedger may have little choice but to accept the price offered by the original counter party since it is dealing only with the bank and there are no competing quotes at this stage. The alternative is to do an offsetting forward contract with a new counter party.

5.4 | Derivatives and Risk Management

This would mean that the company has a forward contract with one bank to buy and a forward contract with another bank to sell. At maturity, the company simply buys from one and sells to the other. In some sense, this cancels the hedge, but at the cost of tying up credit lines with two banks. Though forward contracts do not have explicit margins, banks do reduce the credit limits that they would otherwise have extended to the company. This acts as an implicit margin. Cancelling the forward contract with the original bank releases the margin completely, while offsetting it with another bank, does not release the old credit limit and at the same time, helps tie up a new credit limit.

Another problem with forward contracts is that they are usually not available to individuals or small organizations. This is because the bank or other counter party has to make an evaluation of credit worthiness before offering a forward contract and therefore, may not want to deal with 'retail' customers. For most corporate users, this is not a problem as they usually have well-established banking relationships that give them access to these products. However, this does become an issue for individual investors.

Forward contracts in some cases, may be available only in large contract sizes. Crude oil forword contracts, for example, are usually entered into by oil companies with in their circle or with large customers, for typically large amounts of crude. A company could be quite large in its own business, but its need for crude oil may be really small, and it may find that the forward market is unsuitable for its hedging needs.

5.2.3 Futures

As discussed in Chapter 2, futures are similar to forward contracts but are standardized, exchange traded, and subject to initial and mark to market margins. The margining system is designed to make the market accessible to almost any one. This is a big advantage for retail investors and small organizations who may be effectively shut out of the forward market.

Because of the novation and standardization of these contracts, successful futures contracts are quite liquid and it is therefore easy for people to enter or exit these markets easily and at low cost. Novation means that a hedger who has bought a futures contract could cancel the position at any time by selling in the market to anyone. This is very different from the situation in the forward market where selling to a different counter party would create only two offsetting positions.

The major disadvantage of futures contracts is the basis risk induced by standardization. As discussed at length in Chapter 4, this requires the hedger to carefully monitor the hedge effectiveness and the hedge ratio.

Many companies who are infrequent uses of futures, find the daily mark to market feature of futures contracts quite inconvenient. However, active users of futures contracts do not find this to be a problem because once they have set up the systems and processes to handle these daily cash flows, the costs of managing the mark to market of additional futures contracts is quite negligible. In other words, there is a fixed cost of entering the futures markets, but the variable costs of entering into additional transactions is negligible.

Thus large companies and financial institutions with a large number of risk exposures and hedges are quite comfortable with futures contracts. Retail investors do not have access to forward markets and are therefore delighted to have access to futures. Companies in the middle may be large enough to have access to forward markets, but they are not large and complex enough to have already set up a futures trading system and collateral management system. These companies often exhibit a strong preference for forwards rather than futures.

However in certain conditions, even these mid-sized companies need to consider futures markets. A situation is where the company's risk exposure is indirect rather than direct. In this case, forward contracts are very inconvenient and cumbersome, while futures contracts may be more appropriate. Consider, for

example, a company which incurs large transportation costs that depend indirectly on crude oil prices. A forward contract on crude oil is inconvenient because the company does not directly use crude oil. Since futures contracts can be cash-settled, in this case, they are quite convenient.

Of course, it would be even simpler to have long-term contracts with truck operators to hedge transportation costs. This may be difficult if the company uses a large number of truck operators and moves products to a wide range of destinations. Under such situations, crude oil futures may be a very convenient tool to hedge the most important component of transportation cost risk.

5.2.4 Swaps

Swaps are essentially bundles of forward contracts that are used to hedge interest rate risk and exchange rate risk. These are discussed in detail in Chapter 18. They share many of the characteristics of forward and future contracts discussed above.

5.2.5 Options

Options are contracts that allow the hedger to transact at the pre-committed price or at the ruling market price, whichever is more advantageous. For example, a call option gives the holder the right to buy at a pre-specified 'strike' price but does not impose an obligation to buy at that price. If the market price is lower at maturity, the hedger is free to buy in the market and let the option lapse. If the market price is higher, then the hedger can use the option to buy at the strike price. This appears to give the hedger the best of both worlds. Unfortunately, there is no free lunch. The option costs money in the form of an upfront premium that has to be paid while entering into the option contract.

The use of options to hedge risks is discussed in Chapter 11.

5.3 NON-FINANCIAL HEDGES

While this book is concerned largely with derivatives and related financial instruments, it is important to recognize that non-financial hedges play an important role in risk management. In many real-life situations, derivatives and non-financial hedges are used in conjunction with each other to achieve the risk management objectives.

5.3.1 Contractual Risk Sharing or Risk Transfer

A very common method of managing risk is to share the risk with other parties or even to transfer it completely to them. For example, a company might have a long-term contract with a foreign supplier which specifies a rupee price for the goods that the company imports. This effectively and completely transfers the exchange rate risk to the foreign supplier. Alternatively, the agreement might involve risk-sharing under which a pre-specified fixed dollar price is converted into rupees at an exchange rate which is the average of a pre-specified fixed exchange rate, and the exchange rate prevailing on the date of supply. Under this arrangement, the exchange rate risk is shared equally between the company and its supplier.

Similarly, a company while engaging an engineering and construction contractor under a fixed price contract, might include escalation clauses under which various risks are shared between the company and the contractor.

Risk-sharing is possible with suppliers, customers, employees, and other stakeholders.

5.6 | Derivatives and Risk Management

5.3.2 Risk Reduction

Much of what is called 'risk management' is actually 'risk transfer'. Risk can be transferred to an insurance company, to a derivatives market, to a customer or supplier, and so on. But risks can also be reduced using a variety of technological solutions. For example, a boiler that can run on either oil or natural gas, reduces the fuel price risk. If for example, the price of natural gas rises abnormally due to supply disruptions, the boiler can be shifted to run on oil. Since the transportation infrastructure for oil and gas are quite different, this could give an enterprise, a substantial degree of protection. But again, there is no free lunch—dual fuel boilers cost more than boilers that run on only one fuel.

Currency risks can similarly be reduced through developing relationships with vendors in multiple countries. When the exchange rate of one supplier's currency appreciates, purchases from that supplier can be reduced by substituting other suppliers. This is a form of hedging and the hedging costs are the costs involved in maintaining multiple vendor relationships.

Another way to hedge currency risks is to locate manufacturing facilities in major markets instead of supplying them from home country facilities. Many Japanese automobile companies, for example, have set up factories in the US to partly reduce their exposure to the yen-dollar exchange rate.

5.3.3 Self Insurance and Post Loss Financing

The biggest impact of a risk event is often the cash flow uncertainty that it creates. This uncertainty can be managed by establishing a funding source to meet this cash outflow. In self insurance, the company creates a cash cushion financed out of notional insurance premiums. The idea is that, instead of paying insurance premiums to a insurance company, the company simply invests the premiums itself. Over a few years, this could build up to a cash cushion sufficient to absorb the cash outflow in the event of loss. Another alternative along similar lines is to negotiate a line of credit with a bank which would be drawn down in the event of loss, to meet the cash outflow.

5.3.4 Risk Avoidance

Many companies simply avoid completely getting into certain risky business. A company could avoid credit risk by making only cash sales. It could avoid the exchange rate risk of unstable currencies by invoicing exports to those countries only in dollars or other stable currencies. Most oil-producing countries, for example, invoice their oil exports to all countries in US dollars. (Incidentally, some of these countries do this even though they do not have very friendly relationships with the US.)

Risk avoidance has a cost in terms of the opportunity cost of business that is lost as a result of risk avoidance.

5.4 RISK MANAGEMENT STRUCTURES AND POLICIES

Risk management in a large company is a complex task requiring specialized knowledge of derivatives and finance. This often requires a specialized department with highly trained and skilled staff. This department needs to collect information thoroughly from the organization about the various risks that exist and decide on the methods to mitigate these risks. The organization of this task is discussed below in the section on risk management structures.

We must also recognize that risk management is a matter of strategic importance as it helps protect the company from bankruptcy or financial distress. As such, it requires the involvement of top management and even the board of directors of the company. This involvement typically takes the form of a risk

management policy, that lays down the goals of risk management, and the framework within which the risk department needs to operate. This is discussed in the section below on risk management policy.

5.4.1 Risk Management Structures

Large companies usually have a Risk Manager or a Chief Risk Officer heading the Risk Department. It is common for the Chief Risk Officer to report to the Chief Financial Officer of the company. Less commonly, the Chief Risk Officer may report to the Chief Executive Officer.

The Risk Department must collect information on all the risks arising in different divisions of the company, measure these risks, and aggregate them to determine the total risk exposure of the company. For example, one division of the company may be exporting \in 5.8 million of goods to Italy while another division may be importing \notin 9.3 million of machinery from Germany. On a net basis, therefore, the company needs \notin 9.3 – \notin 5.8 million or \notin 3.5 million. The Central Risk Department might then decide to buy \notin 3.5 million forward to protect against currency risk.

If each division had hedged the risks separately, the exporting division would have done \in 5.8 million of forward transactions and the importing division would have done \in 9.3 million of forward transactions. It would thus have paid commissions on a total of \in 15.1 million of forward transactions. The Central Risk Department on the other hand, pays commissions only on \in 3.5 million – a saving of over 75 percent.

For this to work, all risk exposures would have to be reported to the Central Risk department which must also design systems for accurate measurement of risks.

5.4.2 Risk Management Policy

The board of directors or the top management of the company typically establishes a formal risk management policy that defines the boundaries of the kinds of risks permitted and not permitted. For example, a company's policy may say that 50 to 100 percent of all export receivables will be hedged. This means that at all times, at least half of the receivables are hedged. Depending on the company's view on exchange rate risk and its volatility, this percentage could go up to 100 percent. Implicitly, the policy prohibits over-hedging (that is hedge ratios above 100 percent) even if the company has a very strong view on the expected depreciation of the foreign currency. Hedging is in other words, permitted but outright speculation is not allowed. Note however, that the ability to leave half of the receivables unhedged, permits a limited amount of speculation when the foreign currency is expected to appreciate.

The hedging policy also lays down the kind of instruments that are permitted. For example, it may permit currency forwards but prohibit all other derivatives. It may put restrictions on the banks or other counter parties with which the company can undertake derivative transactions.

When the company uses futures, the mark to market losses on the futures have to be paid in cash, while the offsetting gains on the hedged position are not realized in cash. This creates a need for liquidity. The hedging policy would place limits on the liquidity demands arising from the hedging strategy and outline the ways through which this liquidity demand would be fulfilled. Suppose for example, a company buys three month wheat futures to hedge the wheat that it needs to buy in the next quarter. If the wheat price happens to falls, there is a mark to market loss that has to be paid in cash immediately. This is offset by the anticipated reduction in the price of wheat that the company has to pay over the next quarter. This means that the mark to market loss has to be financed for the next three months. The hedging policy may prohibit hedging strategies that could produce excessive demands for liquidity. It may also lay down the sources of liquidity for the hedging strategies that are permitted.

5.8 | Derivatives and Risk Management

Chapter Summary

It is important to remember that hedging has costs and benefits, and therefore, there is no such thing as a free lunch. Hedging brings greater certainty in cash flows and operations; it facilitates planning and helps the firm focus on its core business. Hedging also gives greater certainty to customers, suppliers, employees, and other stakeholders who then might offer more favourable terms to the firm. While derivatives are a powerful and flexible hedging instrument, there are a number of other financial and non-financial hedges that may be appropriate under certain circumstances. Effective hedging usually requires a centralized that measures risks on a company-wide basis and decides on the strategies to hedge the net risks. The activities of the Risk Department needs to be controlled by a hedging policy laid down by the top management of the company.

Suggestions for Further Reading

Since the discussion in this chapter is continued in Chapter 21, suggestions for further reading are given at the end of that chapter.

CASES

Case 5-1 COIMBATORE YARNS' MARK RECEIVABLES (A)

Struggling to understand the subtleties of international finance, Mr. Shanmugham, Managing Director of Coimbatore Yarns, called his banker for advice. "The rupee has been unexpectedly firm in the last two months", replied the banker, "but our inflation rate is still higher than that of our trading partners and our balance of payments situation is still precarious. I expect the rupee to weaken in the next few weeks. Nevertheless, since you are a newcomer to international operations, I would advise you to take forward cover routinely until you develop enough experience and are able to take a view on the market".

Coimbatore Yarns was a small closely held company in the South Indian state of Tamil Nadu. In the last fifteen years of its existence, it had sold entirely in the domestic market and had never seriously explored the export market. This was partly because during most of this period, an overvalued rupee made Indian yarn uncompetitive in the world market. This position began changing in the early 90s when a balance of payments crisis forced the Indian government to embark on a programme of wide ranging economic reforms and liberalization. In July 1991, the rupee was devalued sharply; in March 1992, the Liberalized Exchange Rate Management System (LERMS) was introduced under which exporters had to surrender only 40% of their foreign exchange earnings at the official rate leaving the remaining 60% to be sold in the free market at a rate which was typically 20% higher than the official rate. In March 1993, all trade transactions were left entirely to the free market. A dollar of export earnings was now worth Rs 31 as against only Rs 20 two years back.

Coimbatore Yarns now started receiving offers from Bombay based brokers eager to buy yarn to fill export orders. Coimbatore Yarns saw no reason why it should sell its yarn to brokers and let them make large profits by exporting their yarn to Europe and elsewhere. It started exploring the market on its own. At the end of April 1993, it was rewarded with an export order from an Italian firm for nearly Rs 1 crore worth of yarn at a price about 15% higher than the domestic price. The yarn was to be shipped in the end of May on 60 days credit. Shanmugham was initially a little worried about the Italian lira's reputation as a weak currency, but his fears were allayed when the Italians readily agreed to accept invoicing in deutsche marks. Happy that his efforts had borne fruit, Shanmugham left the financial aspects of the transaction to be handled by his Finance Manager, Mr. Mahadevan, and resolved to pursue the export market more aggressively.

Mahadevan sought advice from his friends who worked in companies with substantial international transactions. They told him that the basic decision was whether to take forward cover or not. If Coimbatore Yarns waited for the payment to be received from the Italian importer, the amount of rupees that it would get in return for the DM 500,000 invoice value would depend on the exchange rate prevailing on that date. To get rid of this uncertainty, the company would enter into a forward contract to sell this DM 500,000 at the future date at a price specified now. This pre-specified forward rate was often quoted as an annualized percentage premium or discount relative to the spot rate prevailing now; for example, if the six month forward rate was 1.5% above the spot rate, this would be described as a 3% annualized forward premium. Mahadevan was told that the DM was currently trading at about Rs 19.80 and was at a forward premium

Reproduced with permission.

Written by Professors Samir K. Barua and Jayanth R. Varma.

Teaching material of the Indian Institute of Management, Ahmedabad, is prepared as a basis for class discussion. Cases are not designed to present illustrations of either correct or incorrect handling of administrative problems. Copyright © 1997 by the Indian Institute of Management, Ahmedabad.

5.10 | Derivatives and Risk Management

of about 4.5% (annualized). Mahadevan found to his consternation that there was no consensus among his friends about whether to take forward cover or not. It appeared to him though that importers were generally more eager to take forward cover than exporters.

Since this was the first such transaction at Coimbatore Yarn, Mahadevan decided to take the matter to Shanmugham for a final decision. That was when Shanmugham had the conversation with the banker reported at the beginning of this case. Shanmugham was a novice in international finance, but his well developed business acumen told him that the banker's last sentence of advice ran counter to the preceding two sentences of facts and analysis.

Case 5-2 COIMBATORE YARNS' MARK RECEIVABLES (B)

At the end of April 1993, Coimbatore Yarns received an export order from an Italian firm for nearly Rs 1 crore worth of yarn (see Coimbatore Yarns' Mark Receivables (A) for background information regarding this transaction). The yarn was to be shipped in the end of May on 60 days credit. Shanmugham was initially a little worried about the Italian lira's reputation as a weak currency, but his fears were allayed when the Italians readily agreed to accept invoicing in deutsche marks. Convinced that the German mark was a far stronger currency than the Indian rupee, Shanmugham decided to leave the DM 500,000 receivable uncovered. The actual movement of exchange rates (INR/DEM) was as follows:

Date	INR/DEM
End April	19.80
End May	19.70
End June	18.50
End July	18.00

Show the key balance sheet and income statement entries relating to the above transaction in the books of Coimbatore Yarn at each month-end.

In end-April, the DM had been at a forward premium of approximately 4.5% against the rupee. What would the key balance sheet and income statement entries have been if Coimbatore Yarns had taken forward cover on that date?

Case 5-3 WIPRO'S BILLION DOLLAR HEDGE

In September 2004, Wipro had an accumulated mark to market loss of over Rs 2.5 billion on a \$1 billion hedge that it instituted in March 2004. Back in March, when the Indian rupee was climbing rapidly, Wipro had changed its risk management policy and decided to hedge a major part of its anticipated dollar revenues for 2004-05. Since then, a change of government in India and major changes in the global economic environment had changed the picture dramatically. Over the six months from March 31 to September 30, the dollar had risen by 6%.

Written by Professors Samir K. Barua and Jayanth R. Varma.

Teaching material of the Indian Institute of Management, Ahmedabad, is prepared as a basis for class discussion. Cases are not designed to present illustrations of either correct or incorrect handling of administrative problems. Copyright © 1997 by the Indian Institute of Management, Ahmedabad. Reproduced with permission.

Background

Wipro is the second largest Information Technology company in India with \$1.3 billion (Rs 58.4 billion) of revenues in 2004. In terms of revenues, it is about 18% smaller than India's largest software company, Tata Consultancy Services and about 21% larger than the third largest company, Infosys Technologies Limited. Selected financial data about the company is presented in Exhibit 1.

Wipro Limited was incorporated in 1945 as Western India Vegetable Products Limited and was initially engaged in the manufacture of hydrogenated vegetable oil. Over the years, it diversified into the areas of Information Technology or IT services, IT products and Consumer Care and Lighting Products. For the fiscal year ended March 31, 2004, 92% of its operating income was generated from IT business segments.

Wipro's three principal business segments are:

- The Global IT Services and Products segment provides IT services to customers in the Americas, Europe and Japan. The range of services include IT consulting, custom application design, development, re-engineering and maintenance, systems integration, package implementation, technology infrastructure outsourcing, BPO services and research and development services in the areas of hardware and software design. This segment accounted for 75% of revenue and 85% of operating income for the year ended March 31, 2004.
- The India and AsiaPac IT Services and Products segment is a leader in the Indian IT market and focuses primarily on meeting the IT products and services requirements of companies in India, Asia-Pacific and the Middle East region. This business segment accounted for 16% of revenue and 7% of operating income for the year ended March 31, 2004.
- The Consumer Care and Lighting segment (the original core of the company) maintains a profitable presence in niche markets in the areas of soaps, toiletries, lighting products and hydrogenated cooking oils for the Indian market. This business segment accounted for 6% of revenue and 5% of operating income for the year ended March 31, 2004.

The Global IT Services and Products segment has been witnessing rapid growth as companies in North America and elsewhere have increasingly been using external professional services as an effective tool to meet their IT requirements. Outsourcing IT requirements enable companies to lower operating costs, realize productivity gains and convert a portion of their fixed costs into variable costs. By deploying high-speed communications equipment, companies can access skilled IT services from remote locations to meet their complex IT requirements in a cost-effective manner. India has emerged as the country of choice for offshore outsourcing of IT services. A McKinsey study conducted in 2002 for the Indian National Association of Software and Service Companies, or NASSCOM, estimates that the Indian IT services export, as well as the product and technology services, markets are expected to grow from approximately \$6 billion in 2002 to an estimated \$36 billion in 2008.

Prepared by Prof. Jayanth R. Varma on the basis of published material.

Teaching material of the Indian Institute of Management, Ahmedabad, is prepared as a basis for class discussion. Cases are not designed to present illustrations of either correct or incorrect handling of administrative problems. Copyright © 2005 by the Indian Institute of Management, Ahmedabad.

Reproduced with permission.

5.12 | Derivatives and Risk Management

Exchange Rate Risk

Approximately 75% of Wipro's revenues are earned in major currencies of the world while a significant portion of its costs is in Indian rupees. This mismatch is inherent in the Global Delivery Model that underlies the business model of Indian IT companies.

In this model, a small project team based at the client's location in North America or elsewhere defines the scope of the project, tracks project implementation, assists in installing the software and ensures its continued operation. However, the bulk of the development work on the project is performed offshore in India. The higher rate charged for performing work at client sites overseas does not compensate fully for the higher costs in these countries, and therefore, work done in India yields better profit margins. For this reason, Indian IT companies seek to move a project as early as possible to India. As of March 31, 2004, 73% of Wipro's professionals engaged in providing IT services were located in India.

An inevitable implication of this business model is a currency mismatch in which most revenues are in dollars and other foreign currencies while most of the costs are in Indian rupees. During most of the 1990s, Indian IT companies benefited hugely from the secular depreciation of the Indian rupee against world currencies. However, in 2002, the situation began to change.

Situation in March 2004

For over five years from mid 1997 to mid 2002, the Indian rupee depreciated steadily against the US dollar. However, after sinking to its lifetime low against the US dollar in May 2002, the Indian rupee turned around and began to appreciate. Having risen about 8% from its lifetime low, the rupee entered a period of range bound trading in the last quarter of 2003. The currency continued to face upward pressure from capital inflows and from a surplus on the current account, but the Reserve Bank of India intervened aggressively in the market to prevent any further appreciation of the currency. Exhibit 2 presents key macroeconomic indicators for India and Exhibit 3 presents data on the balance of payments.

Towards the end of the financial year 2003-04, the central bank faced a shortage of ammunition in its policy of sterilized intervention. The central bank's purchases of dollars from the market would in the normal course release rupees into the market. To prevent this expansion of the money supply, the central bank combined the purchase of dollars with a sale of government securities in what is known as sterilized intervention. The rupees released by buying dollars was sucked back by the sale of government securities. The difficulty was that after several months of intervention, the Reserve Bank of India had nearly depleted its stock of government securities. Unless this stock could be replenished, the Reserve Bank would have to either halt its intervention or allow the money supply to expand in response to unsterilized intervention.

In February 2004, the government agreed to replenish this arsenal of the Reserve Bank of India by issuing Market Stabilization Bonds over and above its normal borrowing programme. On March 25, 2004, the government and the Reserve Bank of India signed a formal Memorandum of Understanding for this purpose and announced a limit of Rs 600 billion (approximately \$ 13 billion) for the new bonds to be issued under this scheme during 2004–05. The auction calendar for the first quarter of 2004–05 was also announced at the same time.

In the meantime, during the last two weeks of March 2004, the markets were surprised as the rupee surged upwards by 4% while the central bank stayed away from the market. The markets speculated on the possible reasons for the failure of the central bank to dampen this volatility. Was this due to the

delay in launching the Market Stabilization Bonds or did it reflect a change in the exchange rate policy itself? Had the Reserve Bank of India concluded that economic fundamentals implied a stronger rupee and decided to let this correction take place? Or had it decided that the large capital inflows into the equity market should simply be allowed to drive up the rupee? Or had it now decided that rupee appreciation was the best way to counter inflation that was beginning to rise again after being dormant for a few years?

To the Indian exporter, all of these were frightening thoughts. While the Reserve Bank of India maintained a stony silence they were forced to reassess their hedging policies. At the same time, the forward premium which determines the cost of hedging had also become very volatile. Over most of February and March, the six month forward premium had been less than 0.5% so that dollars could be sold forward at a rate that assumed an annual depreciation of the rupee of less than 0.5% from the current exchange rate. During the last few days of March 2004 this premium spiked to almost 1%. Exhibit 4 provides a chart of the exchange rate and the forward premium.

The Billion Dollar Hedge

In response to this changed situation, Wipro made a dramatic change in its risk management policy. Its old policy called for hedging only foreign currency assets and liabilities, so that while existing dollar receivables were hedged, anticipated future receivables arising out of future revenues were left unhedged. In a dramatic departure from this old policy, Wipro now decided to hedge a significant portion of the anticipated U.S. dollar revenues for fiscal 2005.

The result of this major change in policy was a sharp rise in the foreign currency hedges. As of March 31, 2004, the total hedges amounted to almost \$1 billion in notional value consisting of the following:

- Forward contracts to sell US\$ 867 million.
- Net written put options to sell US\$ 113 million
- Total notional value US\$ 980 million

Wipro stated that

"As at March 31, 2004, a Rs.1 increase /decrease in the spot rate for exchange of Indian Rupee with U.S. dollar would result in approximately Rs. 900 million decrease/increase in the fair value of the company's forward contracts and net written options."

Accounting Policy

In its annual report, Wipro made the following statement:

"Further, during the year ended March 31, 2004, the Company has re-evaluated its risk management program and hedging strategies in respect of forecasted transactions. Effective March 2004, upon completion of the formal documentation and testing for effectiveness, the Company has designated the forward contracts in respect of forecasted transactions, which meet the hedging criteria, as cash flow hedges. Changes in the derivative fair values that are designated effective and qualify as cash flow hedges, under SFAS No. 133 Accounting for Derivative Instruments and Hedging Activities, are deferred and recorded as a component of accumulated other comprehensive income until the hedged transactions occur and are then recognized in the consolidated statements of income. The ineffective portion of a hedging derivative is immediately recognized in the consolidated statements of income.

5.14 | Derivatives and Risk Management

... Net written options are ineligible for hedge accounting under SFAS No. 133. Consequently, the [change] in fair value of the net written options is recognized in the consolidated statement of income."

The formal documentation required under paragraph 28(a) of SFAS No. 133 Accounting for Derivative Instruments and Hedging Activities is as follows:

"At inception of the hedge, there is formal documentation of the hedging relationship and the entity's risk management objective and strategy for undertaking the hedge, including identification of the hedging instrument, the hedged transaction, the nature of the risk being hedged, and how the hedging instrument's effectiveness in hedging the exposure to the hedged transaction's variability in cash flows attributable to the hedged risk will be assessed. ...Documentation shall include all relevant details, including the date on or period within which the forecasted transaction is expected to occur, the specific nature of asset or liability involved (if any), and the expected currency amount or quantity of the forecasted transaction. ... The hedged forecasted transaction shall be described with sufficient specificity so that when a transaction occurs, it is clear whether that transaction is or is not the hedged transaction."

In addition, paragraph 29(b) of SFAS No. 133 requires that "The occurrence of the forecasted transaction is probable."

Situation in September 2004

The first half of 2004-05 saw rupee depreciation of 6% instead of the much feared appreciation. In sharp contrast to the situation in 2003-04, the Reserve Bank was now intervening in the markets to support the rupee. Exhibit 5 presents monthly data on Reserve Bank interventions, reserves and capital inflows into India. Several factors combined to produce this outcome:

- 1. The prospect of rising interest rates in the US led to a reduction and even reversal of capital flows to emerging markets as a whole in May 2004.
- 2. This was coupled with the defeat of the incumbent government in the general elections in May. Uncertainty about the economic policies of the new government which depended on the support of left leaning parties led to a major fall in the stock market and contributed to the reduction of capital flows.
- 3. Rising oil prices in mid 2004 put downward pressure on the rupee because India imports a large fraction of its oil.
- 4. Inflation rose sharply in August partly in response to rising oil prices and this again impacted the rupee negatively.

Mark to Market Losses

In its financial statements for the quarter ended September 31, 2004, Wipro stated that its foreign exchange hedges had incurred a loss of Rs 2.8 billion (approximately \$60 million). These losses did not flow through into the reported profits of the company because of the hedge accounting discussed above. What happens under hedge accounting is that when the forecasted revenues actually materialize, they would be accounted for at the forward rate locked into in March 2004 even if the then prevailing exchange rate is more favourable than that forward rate. Thus if the rupee continues to be weak, the cost of hedging will reflect itself in lower revenues than if the company had been unhedged.

Required

1. In the circumstances of the case, would you say that Wipro's change of risk management policy amounted to speculating on the currency? Explain.

- 2. Was the size of the hedge inadequate, adequate or excessive?
- 3. Does the hedging policy need a rethink in September 2004?
- 4. Is it a good idea to sell put options to hedge the currency risk?
- 5. Do you think that the hedge accounting for these transactions is appropriate? Do you think hedge accounting reduces the flexibility of the company in dealing with exchange rate volatility?

Exhibit 1 Wipro Limited: Summary of Selected Consolidated Financial Data

	2002	2003	2004	2004
	Rs	Rs	Rs	\$
	million	million	million	million
Consolidated Statements of Income data				
Revenues:				
Global IT Services and Products				
Service s	21,457	30,118	43,343	999
Products	955	149	122	3
India and AsiaPac IT Services and Products				
Services	1,914	2,240	3,109	72
Products	5,037	5,801	6,305	145
Consumer Care and Lighting	2,939	2,942	3,567	82
Others	1,171	1,599	1,987	45
Total	33,473	42,849	58,433	1,346
Cost of revenues:				
Global IT Services and Products				
Services	11,419	17,635	27,853	642
Products	891	103	78	2
India and AsiaPac IT Services and Products				
Services	1,160	1,187	1,661	38
Products	4,268	5,100	5,643	130
Consumer Care and Lighting	1,999	2,008	2,355	54
Others	924	1,143	1,410	33
Total	20,661	27,176	39,000	899
Gross Profit	12,812	15,673	19,433	448
Operating expenses:				
Selling, general and administrative expenses	(4,359)	(6,193)	(8,450)	(195)
Amortization of intangible assets		(166)	(308)	(7)
Other operating income/(expenses)	(11)	172	226	5
Operating income	8,442	9,486	10,901	251
Gain/(loss) on sale of stock by affiliates, including	—	—	(206)	(5)
direct issue of stock by affiliate				
Other income/(expense), (net)	838	718	868	20
Equity in earnings of affiliates	147	(355)	96	2
come before taxes and minority interest	9,427	9,849	11,659	9 268

5.16 | Derivatives and Risk Management

Exhibit 1 (Contd.)

	2002	2003	2004	2004
	Rs	2003 Rs	Rs 2004	2004 \$
	million	million	million	million
Income taxes	(1,016)	(1,342)	(1,611)) (37)
Minority interest	—	(30)	(56)	(1)
Income from continuing operations	8,411	8,477	9,992	2 230
Earnings per share from continuing operations:				
Basic	36.39	36.66	43.20	1.00
Diluted	36.33	36.60	43.16	0.99
Cash dividend per equity share	0.50	1.00	1.00	0.02
Segment Information				
Revenue by Segment				
Global IT Services and Products	22,668	30,593	43,775	1,009
India and AsiaPac IT Services and Products	6,950	8,046	9,445	217
Consumer Care and Lighting	2,939	2,942	3,567	82
Others	916	1,268	1,646	38
Total	33,473	42,849	58,433	1,346
Operating Income by Segment				
Global IT Services and Products	7,609	8,281	9,300	214
India and AsiaPac IT Services and Products	578	539	761	18
Consumer Care and Lighting	404	422	546	12
Others	41	256	308	7
econciling items	(190)	(12)	(14)	—
Total	8,442	9,486	10,901	251
Consolidated Balance Sheet Data				
Cash and Cash equivalents	3,251	6,283	3,297	76
Investments in liquid and short-term mutual funds	4,126	7,813	18,479	426
Working Capital	18,495	21,473	30,649	706
Total assets	33,639	42,781	57,738	1,330
Total debt, including preferred stock	291	537	969	22
Total stockholders equity	27,457	35,431	46,364	1,068

Notes:

- The consolidated financial statements have been reported in Indian rupees, the national currency of India. Solely for convenience, the financial statements as of and for the year ended March 31, 2004, have been translated into US dollars at the exchange rate on March 31, 2004 of \$1 = Rs. 43.40.
- Total debt for the year ended March 31, 2000 includes preferred stock of Rs. 250 million.

Source: Wipro Limited Annual Report for 2003-04 (Form 20F) filed with the United States Securities and Exchange Commission.

Item	2003-04	2002-03	2001-02	2000-01	1999-00
Real GDP (% change)	8.2	4.0	5.8 P	4.4	6.1
Industrial Production (% change)	6.9	5.7	2.7	5.0	6.7
Agricultural Production (% change)	19.6	- 15.2	7.6	- 6.4	- 1.3
Gross Domestic Saving Rate (% of GDP)	n.a.	24.2	23.5 P	23.7	24.2
Gross Domestic Investment Rate (% of GDP)	n.a.	23.3	23.1 P	24.4	25.3
Central Government Finances (% of GDP)					
Total Revenue Receipts	9.5	9.4	8.8	9.2	9.4
Total Expenditure	17.1	16.8	15.9	15.6	15.4
Revenue Deficit	3.6	4.4	4.4	4.1	3.5
Fiscal Deficit	4.8	5.9	6.2	5.7	5.4
Domestic Debt	60.5	60.7	56.7	52.8	49.7
Monetary Aggregates (% change)					
Broad Money (M3)	16.6	12.7	14.1	16.8	14.6
Scheduled Commercial Banks (% change)					
Aggregate Deposits	17.5	13.4	14.6	18.4	13.9
Bank Credit	15.3	16.1	15.3	17.3	18.2
Wholesale Price Index (% change)					
Point-to-Point	4.6	6.5	1.6	4.9	6.5
Average	5.4	3.4	3.6	7.2	3.3
Consumer Price Index – Industrial Workers					
(% change)					
Point-to-Point	3.5	4.1	5.2	2.5	4.8
Average	3.9	4.0	4.3	3.8	3.4
BSE Sensitive Index (% change)	40.1	- 3.8	-22.0	- 8.4	41.4
Trade and Balance of Payments					
Exports in US \$ (% change)	20.4	20.3	- 1.6	21.0	10.8
Imports in US \$ (% change)	25.4	19.4	1.7	1.7	17.2
Current Account (% of GDP)	1.4	0.8	0.2	-0.8	- 1.0
Capital Account (% of GDP)	3.7	2.4	2.2	2.2	2.3
Foreign Exchange Reserves (US \$ billion)	112	76	55	43	39
External Debt (US \$ million)	113	105	99	101	98
Debt-GDP Ratio	17.6	20.2	21.1	22.6	22.1
Debt-Service Ratio	18.3	15.1	13.9	17.2	16.2
Exchange Rate (Rupee / US\$)					
High	43.45	47.51	46.56	43.61	42.44
Low	47.46	49.06	48.85	46.89	43.64

Exhibit 2 India: Select Macroeconomic and Financial Indicators

Source: Reserve Bank of India, Annual Report 2003-04, Appendix Table I.1

5.18 | Derivatives and Risk Management

Exhibit 3

India's Overall Balance of Payments

US \$ Million

	2003-04	2002-03	2001-02]
A. Current Account			
1. Exports, f.o.b.	62,952	52,512	44,915
2. Imports, c.i.f.	79,658	65,422	57,618
3. Trade Balance	- 16,706	- 12,910	- 12,703
4. Invisibles, Net	25,425	17,047	13,485
a) 'Non-Factor' Services	10,684	6,765	4,577
of which: Software Services Exports	12,200	9,600	7,556
b) Income	-4,703	- 4,935	- 3,601
c) Private Transfers	18,885	14,807	12,125
d) Official Transfers	559	410	384
5. Current Account Balance	8,719	4,137	782
B. Capital Account			
1. Foreign Investment, Net (a+b)	14,492	4,555	6,692
a. Direct Investment of which :	3,137	3,611	4,741
i. In India	4,675	4,660	6,131
Equity2,387	2,700	4,095	R e
invested earnings	1,800	1,498	1,646
Other Capital	488	462	390
ii. Abroad	- 1,538	- 1,049	- 1,390
Equity	- 282	-424	- 570
Re-invested earnings	- 1,176	- 519	- 699
Other Capital	-80	- 106	- 121
b. Portfolio Investment	11,355	944	1,951
2. External Assistance, Net	- 2,661	- 2,460	1,117
Disbursements	3,374	2,773	3,352
Amortisation	- 6,035	- 5,233	-2,235
3. Commercial Borrowings	- 1,853	- 2,344	- 1,576
Disbursements	6,253	2,833	2,696
Amortisation	- 8,106	- 5,177	-4,272
4. Short Term Credit, Net	1,560	979	- 891
5. Banking Capital	6,197	8,412	5,592
of which : NRI Deposits, Net	3,628	2,976	2,754
6. Rupee Debt Service	- 376	- 474	- 519
7. Other Capital, Net @	4,763	3,445	158
8. Total Capital Account	22,122	12,113	10,573
C. Errors & Omissions	580	730	402
D. Overall Balance [A(5)+B(8)+C]	31,421	16,980	11,757

Source: Reserve Bank of India, Annual Report 2003-04, Appendix Table VI.2



Exhibit 5 India: Select Monthly Data on FX Reserves, Intervention and FII Inflows

Month	Net purchase of dollars by RBI	Foreign currency reserves	Net FII Inflows
Aug 2003	2,352	82,624	492
Sep 2003	2,345	87,213	904
Oct 2003	1,593	88,674	1622
Nov 2003	3,449	92,148	883
Dec 2003	2,888	97,617	1549
Jan 2004	3,294	100,780	1147
Feb 2004	3,357	104,069	696
Mar 2004	3,382	107,448	1834
Apr 2004	7,427	113,011	846
May 2004	- 220	114,102	- 457
Jun 2004	- 413	114,151	- 477
Jul 2004	- 1,180	112,967	- 432

Source: Centre for Monitoring the Indian Economy

5.20 | Derivatives and Risk Management

Case 5-4 EUROTUNNEL GROUP: FOREIGN EXCHANGE EXPOSURE

The Eurotunnel Group operates a system of tunnels under the English Channel which provide the sole direct link between the motorway networks of the UK and Continental Europe and also permit passenger and through freight trains to travel directly between the UK and Continental Europe. Eurotunnel PLC (incorporated in the UK) and Eurotunnel SA. (incorporated in France) are the ultimate holding companies of the Eurotunnel Group, and their shares are listed in the form of Units (comprising one share in each company) on the London, Paris and Brussels stock exchanges. The French and UK governments have granted a 65 year Concession (expiring in July 2052) to the Eurotunnel Group companies to carry out the development, financing, construction and operation of the Channel tunnels; Eurotunnel is currently in the process of negotiating with the governments for extending the Concession till 2086. The Concession Agreement provides that the Eurotunnel Group shall abide by the principle of equal sharing between the UK and French companies of the total project cost as well as of all revenues and costs relating to the operation of the tunnels. The Eurotunnel group is managed by the Eurotunnel Joint Board which is essentially a common Board of Directors for Eurotunnel PLC and Eurotunnel SA. See Exhibit 6 for further details about the operations and structure of the group.

After experiencing financial difficulties in 1995, the Eurotunnel Group was forced to suspend payment of interest on most of its loans and begin negotiations with lenders for a restructuring of its indebtedness. In October 1996, the Steering Group appointed by the lenders reached an in principle agreement with Eurotunnel on the terms of restructuring. In May 1997, the detailed proposals were circulated to the shareholders and to the lenders for their approval. The restructuring agreement is expected to be completed during the first half of 1998. The date of completion is referred to below as the "Effective Date".

The nature of Eurotunnel's operations, its ownership structure, and the terms of its financial restructuring give rise to foreign exchange risks of various kinds as described below:

Translation Gains and Losses

Exhibits 1, 2 and 3 provide the balance sheet, profit and loss account and cash flow statement of the Eurotunnel Group. Translation gains and losses arise in the Eurotunnel Group's accounts due to the consolidation of the accounts of the EPLC Group (Eurotunnel PLC and its subsidiaries) and the ESA Group (Eurotunnel SA. and its subsidiaries). The notes to the accounts stated that the accounts of the ESA Group (which are in French Francs) have been converted into sterling as follows:

- share capital, share premium account, retained reserves brought forward, and concession fixed assets at historical rates;
- other assets and liabilities together with the profit and loss account, with the exception of depreciation at the rate ruling at the balance sheet date.

(At 31 December 1996, $\pounds 1 = FRF 8.90$; at 31 December 1995, $\pounds 1 = FRF 7.59$; at 31 December 1994, $\pounds 1 = FRF 8.35$. The financial projections prepared in connection with the debt restructuring assumed an exchange rate of $\pounds 1 = FRF 8.50$ during the period 1997-99)

Exchange differences arising out of the application of the above are included in the exchange adjustment reserve in the balance sheet. These exchange differences amounted to a loss of £404 million in 1995 and a profit of £713 million in 1996.

Prepared by Professors Jayanth R. Varma.

Case material of the Indian Institute of Management, Ahmedabad, is prepared as a basis for class discussion. Cases are not designed to present illustrations of either correct or incorrect handling of administrative problems. Copyright © 1998 by the Indian Institute of Management, Ahmedabad.

Foreign Currency Borrowings

Apart from the above, exchange differences also arise from retranslation at the year-end exchange rate of borrowings in foreign currencies and of inter company accounts. These differences which were charged to the profit and loss account amounted to profits of £89 million in 1995 and £91 million in 1996. Currently Eurotunnel has borrowings in sterling, French francs, US dollars, ECUs and Belgian francs. Under the restructuring proposals all the liabilities of the EPLC Group and the ESA Group are to be redenominated in sterling and French francs respectively, and the inter company accounts are to reduced substantially.

Operating Income and Costs

In its restructuring proposals, Eurotunnel stated that:

"The Eurotunnel Group generates revenue and incurs costs in both French francs and in sterling. Based on historic experience, Eurotunnel expects that a majority of revenues will continue to be generated in sterling as opposed to French francs, whilst a majority of operating expenditure is expected to be incurred in French francs."

The Operating Cash Flow Covenant

Under the restructuring proposals, Eurotunnel committed itself to achieving a minimum level of operating cash flows (before deducting leasing payments) on a half yearly basis till 2005. This minimum level was based on several sets of financial projections (See Exhibit 4) prepared by Eurotunnel:

- three year plan (1997–1999) showing two sets of projections: a base case and a downside case;
- extended projections (2000–2006) showing two sets of outcomes: an upper case and a lower case;
- extrapolations (2007–2052) showing two sets of outcomes: an upper case and a lower case;

The minimum levels of operating cash flows were fixed sterling amounts (see Exhibit 5) set at or slightly below the operating results projected under the downside case for 1997-99 and the lower case thereafter. The Operating Cash Flow Covenant would be breached in the event of failure to achieve the specified minimum cash flows in any three successive half years after the Effective Date. A breach of the covenant would be an event of default which would allow the lenders to demand early repayment of their loans, enforce their security or seek to effect substitution. The right of substitution derived from the Concession Agreement allows the lenders to take over the Concession in case of default thereby assuming management of the project and enabling them to apply revenues in satisfaction of sums dues to them.

Stabilisation Facility Limits and Debt Service Covenants

The restructuring proposals allowed Eurotunnel to use an additional credit facility known as the Stabilisation Facility if it was unable to pay interest in cash during the Stabilisation Period which is scheduled to end in 2005. There are covenants regarding the maximum amount of Stabilisation Facility that can be availed of both on an annual basis and a cumulative basis during this period. There are also covenants regarding the minimum debt service coverage ratio and interest coverage ratio after the Stabilisation Period. These covenants are described in Exhibit 5. A breach of any of these covenants

5.22 | Derivatives and Risk Management

would be an event of default which would allow the lenders to demand early repayment of their loans, enforce their security or seek to effect substitution.

Sensitivity Analysis by Eurotunnel

As part of the restructuring proposals, Eurotunnel provided the results of sensitivity analysis identifying the variation of key assumptions which would result in a breach of the Stabilization Facility Limits and the Debt Service Covenants. The results of this analysis was encouraging. For example, the stabilisation facility limit would be breached in 2000 only if operating revenues were 22% below the lower case or operating expenditure were 26% higher than the lower case in each year from 1997 onwards.

In its introduction to these sensitivity results, Eurotunnel stated as follows:

"The Group's financial results may be positively or negatively affected by fluctuations in the value of the French franc against sterling, depending on the Group's mix of revenue and costs at the time and the impact, if any, of exchange rate on travel patterns. The Directors believe that this potential exposure can be managed, other than for extreme movements in FRF:£ exchange rates, by a combination of, *inter alia*, yield management, hedging and cost control such that the Eurotunnel will continue to comply with the Operating Cash Flow Covenant.

The results summarized in the tables which follow have been calculated in each case by reference to the Stabilisation Facility Limit and where appropriate, the Debt Service Covenants. However, any sustained adverse variation, even if small, in a key assumption underlying the Downside Case and the Lower Case would cause a breach of the Operating Cash Flow Covenant possibly in early 2000, assuming that the Effective Date occurs between 1 January 1998 and 30 June 1998."

Required

- 1. What are the foreign exchange risks that Eurotunnel faces?
- 2. What should be its policy regarding foreign exchange risk management?
- 3. What specific strategies should it follow to manage foreign exchange risks?

	31 Dec 1994	31 Dec 1995	31 Dec 1996
	£m	£m	£m
Assets			
Tangible fixed assets			
Concession fixed assets	9,452.9	9,395.2	9,279.9
Other fixed assets	4.4	4.2	4.2
Total tangible fixed assets	9,457.3	9,399.4	9,284.1
Financial fixed assets			
Shares	0.0	0.1	0.1
Others	3.9	0.3	0.2
Total fixed assets	9,461.2	9,399.8	9,284.4
Stocks	5.5	13.2	12.5

Exhibit 1 Eurotunnel Group: Combined Balance Sheet

Trade debtors	0.0	39.7	28.2
Other debtors	0.0	15.5	32.5
Debtors	61.6	0.0	0.0
Investments and liquid funds	302.1	94.5	130.7
Total current assets	369.2	162.9	203.9
Prepaid expenses	22.9	6.7	5.5
Total assets	9,853.3	9,569.5	9,493.9
Shareholders' funds and liabilities			
Issued share capital	1,318.5	1,370.9	1,370.9
Share premium account	1,204.1	1,207.6	1,207.6
Profit and loss account brought forward	(2.2)	(389.1)	(1,314.0)
Loss for the year	(386.9)	(924.9)	(685.4)
Exchange adjustment reserve	(393.6)	(798.0)	(85.1)
Total shareholders' funds	1,739.8	466.5	494.0
Provisions	0.1	5.0	26.9
Loan notes	7.1	7.1	7.9
Bank loans and overdrafts	8,015.7	8,891.6	8,791.2
Other creditors	87.1	189.6	161.6
Total creditors	8,110.0	9,088.4	8,960.7
Deferred income	3.3	9.6	12.3
Total shareholders' funds and liabilities	9,853.3	9,569.5	9,493.9

Exhibit 2 Eurotunnel Group: Combined Profit and Loss Account

·	31 Dec 1994	31 Dec 1995	31 Dec 1996
	£m	£m	£m
Turnover			
Turnover and other operating income	30.6	298.6	448.0
Own work capitalised	715.8	0.0	0.0
Deferred expenses and recharges	51.5	5.3	35.4
Total turnover	798.0	303.9	483.5
Operating expenditure			
Materials and services(net)	336.1	277.5	265.1
Staff costs	65.6	90.1	88.2
Depreciation	145.7	130.4	136.8
Provisions	0.0	6.0	25.5
Other operating charges	0.2	0.3	0.6
Total operating expenditure	547.5	504.3	516.2
Operating profit/(loss)	250.4	(200.4)	(32.8)
Financial income			
Income from investments	10.0	4.2	2.6

5.24 | Derivatives and Risk Management

Interest receivable and similar income	6.3	8.3	2.4
Profit on disposal investments	0.5	0.6	0.3
Exchange differences	44.7	89.4	90.7
Total financial income	61.6	102.5	96.0
Financial charges			
Interest payable and similar charges	649.2	768.2	620.6
Exchange differences	49.7	58.8	117.9
Total financial charges	698.9	827.0	738.4
Financial result/(loss)	(637.3)	(724.5)	(642.5)
Net amount written off fixed assets	0.0	0.0	(10.1)
axation	0.0	0.0	0.0
Loss for the year	(386.9)	(924.9)	(685.4)
Loss per unit	(53.2p)	(103.1p)	(74.5p)

Exhibit 3 Eurotunnel Group: Combined Cash Flow Statements

	31 Dec 1994	31 Dec 1995	31 Dec 1996
	£m	£m	£m
Net cash (outflow)/ inflow from operating			
activities	(333.1)	101.3	115.5
Returns on investments and servicing of			
finance			
Interest received	15.8	14.2	5.4
Interest paid	(591.3)	(611.8)	(34.2)
Bank fees and expenses paid	(81.4)	(59.8)	(6.6)
Net cash outflow from returns on investments			
and servicing of finance	(656.9)	(657.5)	(35.4)
Investing activities			
Payments to acquire tangible fixed assets	(457.6)	(68.9)	(44.4)
Sale of tangible fixed assets	4.3	3.0	1.1
Net cash outflow from investing activities	(453.2)	(65.9)	(43.3)
Net cash (outflow) /inflow before financing	(1,443.2)	(622.1)	36.7
Financing			
Issue of ordinary share capital	908.7	55.9	0.0
Net increase/(decrease) in bank loans	795.2	345.5	(1.5)
Net cash inflow/ (outflow) from financing	1,703.9	401.4	(1.5)
Exceptional activities			
Rolling stock re-instalment insurance receipts	0.0	0.0	4.5
Net cash inflow from exceptional activities	0.0	0.0	4.5
Increase/ (decrease)in cash and cash			
equivalents	260.7	(220.7)	39.8

	The MCC	<mark>5raw-Hill</mark> Companies	s			
		Proj	Exhibit 4 jected Cash F	lows		
Year	Base Operating profit before depreciation	Case / Upper Case Net cash inflow from operating activities	Net capital expenditure	Downs Operating profit before depreciation	ide Case/ Lower Case Net cash inflow from operating activities	Net capital expenditure
1996						
(actuals)	130	115	43	130	115	43
1997	203	174	53	198	169	53
1998	277	258	33	272	252	46
1999	338	312	12	296	294	38
2000	381	322	40	298	287	67
2001	408	380	101	364	335	105
2002	448	436	66	398	376	113
2003	492	486	123	426	409	102
2004	551	543	96	463	450	88
2005	625	617	86	517	502	86
2006	710	200	20	568	553	53
2007	773			610		
2017	1240			994		
2027	1778			1451		
2037	2371			2018		
2047	3064			2617		
2051	3389			2896		

5.26 | Derivatives and Risk Management

Exhibit 5 DEBT COVENANTS AND DEFAULT

A. Operating Cash Flow Covenant

The operating cash flow covenant would be breached and an event of default would occur if operating Cash Flow (before deducting leasing payments) in any three successive six month periods during the period commencing on the first 1 January or 1 July falling after the effective date and ending on 31 December 2005 are less than the relevant amounts, which are fixed in sterling terms, set out below in each of those six month periods

	Operating Cash Flow (£ m) for Six Month Period Ended	
	30 June	31 Dec
1998	125	125
1999	145	145
2000	143	143
2001	165	165
2002	185	185
2003	200	200
2004	225	225
2005	250	250

B. Stabilisation Facility Limit

The Stabilisation Facility Limit would be breached and an event of default would occur if the short fall between interest payable on advances of Residual Junior Debt, the Resettable Facility and the Equity Notes and the interest actually paid in cash exceeds the periodic or cumulative limits set out in the table below.

Relev	ant periods	Period limit £ m (equivalent)	Cumulative limit £ m (equivalent)
15 October 1996	– 31 January 1998	520	520
1 February 1998	– 31 January 1999	390	850
1 February 1999	- 31 January 2000	280	1100
1 February 2000	- 31 January 2001	280	1300
1 February 2001	- 31 January 2002	280	1500
1 February 2002	- 31 January 2003	220	1650
1 February 2003	- 31 January 2004	200	1800
1 February 2004	- 31 January 2005	150	1850
1 February 2005	- 31 January 2006	75	1850

The cumulative limits are set independently of the periodic limits and each will be specified as separate sterling and French franc amounts in the same ratio as the estimated split of the debt service costs of the Residual Junior Debt, the Resettable Facility and the Equity Notes each to be determined on the basis of exchange rates prevailing on the Strike Date.

C. Debt Service Covenant

These covenants apply after the end of the Stabilisation Period. An event of default would occur if

- The ratio of operating cash flow to debt service costs plus capital expenditure is less than 1.0 for any financial year in the period beginning 1 January 2006 and ending 31 December 2011 and 1.2 for any financial year in the period beginning 1 January 2012 and ending 31 December 2025.
- The ratio of earnings before interest and tax to total interest service costs is less than 1.0 for any financial year in the period beginning 1 January 2008 and ending 31 December 2011 and 1.5 for any financial year in the period beginning 1 January 2012 and ending 31 December 2025.

Exhibit 6

Note on Operation of the Eurotunnel¹

The Eurotunnel Group of which Eurotunnel PLC (incorporated in the UK) and Eurotunnel SA. (incorporated in France) are the ultimate holding companies has a 65 year Concession from the French and UK governments to carry out the development, financing, construction and operation of the Channel tunnels. This note describes the Eurotunnel system and the operations of the Eurotunnel Group and draws heavily on documents published by Eurotunnel.

The Eurotunnel System

The Eurotunnel System plays a unique role in Europe's transport network. It is the sole service offering a direct link between the motorway networks of the UK and Continental Europe, also permitting passenger and through freight trains to travel directly between the UK and the Continental European destinations.

Four separate train services operates through the tunnel. Eurotunnel operates the tourist and freight vehicle shuttle services under the Le Shuttle name between the UK terminal located near Folkestone and the French terminal in Coquelles. The Eurostar passenger services and through freight train services are not operated by Eurotunnel, but their operators make payments to Eurotunnel for the use of the Tunnel. Eurotunnel generates additional revenue from the provision of duty-free, tax-free and tax-paid shopping and catering facilities to customers at the terminals, as well as from other ancillary services such as telecommunications and property development. The Group holds the Concession to operate the Eurotunnel System until 2052. In 1997, the Group was in discussions with the governments of the UK and France for a significant extension to the current Concession to be granted.

The Eurotunnel System linking the UK and France comprises three tunnels, each approximately 50 kilometres in length, which run under the English Channel between the two terminals. The two single-track rail tunnels are normally used for trains running in one direction. The third tunnel, which lies for most of its length between the two rail tunnels, is used for maintenance and safety purposes. In addition, there are four cross-over points which allow trains to switch between the two rail tunnels while maintenance work is being carried out on isolated sections of the tunnels.

The two terminals have direct access to the motorway networks in both the UK and France. They are the arrival points for road vehicles utilising Eurotunnel's Le Shuttle services. Each terminal offers a range of services to customers, including duty-free sales, shops and catering facilities. Outward and

Note prepared by Professors Jayanth R. Varma and N Venkisteswaran, Indian Institute of Management, Ahmedabad.

5.28 | Derivatives and Risk Management

inward frontier controls and security clearance take place at the departure terminal, avoiding the need for such systematic controls on arrival. The movement of all rail traffic through the System is controlled from the main railway control centre located within the UK terminal. A standby control centre is located within the French terminal.

Services and Revenues

Le Shuttle Tourist

The Le Shuttle Tourist service carries mainly cars, coaches, minibuses and motorcycles in shuttles between the UK and the French terminals. Customers drive to the terminals where they can either first visit the passenger terminal building and make use of its facilities or head directly for the first available shuttle departure. Tickets can be purchased in advance from travel agents or by telephone from the customer service centre or on arrival at the tolls. Customers accompany their vehicles throughout the journey which, from platform to platform, normally takes approximately 35 minutes.

In 1996, there was a total of 36,045 departures, compared with 26,687 in 1995, an increase of 35 percent. A total of 2,155,117 tourist vehicles were carried through the Tunnel in 1996. This represents an increase of 73 percent over 1995.

Le Shuttle Freight

Le Shuttle Freight service carries HGVs in shuttles between the UK and the French terminals. Customers can either negotiate rates, or buy tickets on arrival at the terminals. HGV drivers have dedicated tolls and customs, security clearance and facilities in both terminals (including duty-free services) as well as additional truck stop facilities close to the terminals. Unlike the Le Shuttle Tourist service, HGV drivers and their passengers travel separately from their vehicles in a "Club Car" wagon in which a meal is served.

In late 1996, the operations of Le Shuttle Freight was suspended following a fire in an HGV onboard a freight shuttle. Despite this, a total of 519,003 HGVs were transported through the Tunnel during 1996, an increase of 33 percent over 1995. In 1996, there was a total of 32,264 departures, an increase of 8 percent over 1995. Freight shuttle services were reintroduced in May 1997 after a safety inquiry.

Eurostar

Eurotunnel derives revenue from the use made of the Tunnel by Eurostar passenger trains. Eurotunnel does not operate these services but ensures their safe and efficient passage through the System. Under the terms of the contract, Eurotunnel is guaranteed a minimum level of revenue until 2006.

The Eurostar high-speed passenger service is operated jointly by London & Continental Railways, SNCF and SNCB. In normal operation, the service carries passengers between London and Paris in a scheduled direct time of about three hours and between London and Brussels in a scheduled direct time of about three hours and ten minutes. Following the commencement of operations on 14 November 1994, the frequency of the Eurostar service increased up to a minimum of 14 scheduled services per day in each direction between London and Paris and eight per day in each direction between London and Paris ran between London and Paris, an increase of 77 percent compared with 1995, and a total of 4,636 trains ran between London and Brussels, an increase of 51 percent compared with 1995. In 1996, a total of 3,672,015 passengers were carried by Eurostar services
How and Why do Firms Hedge? | 5.29

between London and Paris, an increase of 67 percent compared with 1995, and a total of 1,194,551 passengers between London and Brussels, an increase of 66 percent compared with 1995.

Through freight services

Eurotunnel derives revenue from the use made of the Tunnel by long distance through freight trains. Eurotunnel does not operate these services but ensures their safe and efficient passage through the System. Under the terms of the contract, Eurotunnel is guaranteed a minimum level of revenue until 2006.

Through freight services are operated by SNCF and Railfreight Distribution, which have formed a common management body for international freight services on both networks. The level of through freight traffic averaged approximately 140 trains per week during 1996, an increase of 36 percent over 1995. This represented 7,191 trains in total for 1996 and an average weekly volume of approximately 45,000 gross tonnes over this period, an increase of approximately 75 percent over 1995. In 1996, 2.4 million tonnes of freight were transported through the Tunnel on through freight trains.

There are three types of through freight traffic which pass through the Tunnel : (I) freight containers and swap bodies ("intermodal traffic"); (ii) conventional rail wagon services for a wide variety of goods; and (iii) automotive traffic comprising the transport of new vehicles and car components. Of the trains running through the Tunnel in 1996, approximately 67 percent were intermodal, 19 percent were conventional rail wagon and 14 percent were automotive.

Through freight services are principally intermodal services between the UK and Italy, France, Spain and Belgium. Automotive trains carry new vehicles from the UK to Italy and new vehicles are transported to the UK from France, Italy and Belgium.

Retail

Eurotunnel generates a significant proportion of its revenues from retail activities, particularly dutyfree sales. The retail activities include shops providing duty-free, tax-free and tax-paid goods and catering facilities, located at both the UK and French terminals. These retail activities are available for use by travelling customers. Eurotunnel generates revenue form these retail activities in a number of ways :

- (i) duty-free activities are owned by the Group, although they are operated under a management contract by Aer Rianta, the Irish airport operator and retail services provider;
- (ii) retail concessions pay Eurotunnel a rent, either on a fixed or percentage of turnover basis;
- (iii) outlets known as "Le Shop" which sell "Le Shuttle" merchandise and which are located in both terminals are owned and managed by the Group;
- (iv) fuel retailers located in the terminals pay the Group a fee based on the volume of fuel and other goods sold; and
- (v) advertising revenue is generated from advertising space available in the terminals and on shuttles.

Telecommunications

A small proportion of its revenues is derived from selling fixed-line telecommunications capacity installed in the Tunnel. Eurotunnel upgraded its existing telecommunications infrastructure in the Tunnel and is now equipped to offer international telecommunications services over its network in competition with the existing public telecommunications operators. By the end of 1996, Eurotunnel had obtained international telecommunications operator licences in the UK and France. These licences allow Eurotunnel to offer international telecommunications services to the public on a commercial basis.

5.30 | Derivatives and Risk Management

Property development

A further small proportion of its revenues is derived from property development activities. The Eurotunnel Group includes two wholly-owned companies, Eurotunnel Developments Limited ("EDL"), which holds certain property interests, and Eurotunnel Developments SA. ("EDSA").

Eurotunnel currently manages and controls approximately 207 hectares of land close to the terminals. There is scope both in the UK and France for this land to be developed. Eurotunnel intends to carry out development by selling or leasing land and buildings where the proposed development is likely to enhance the status of the area close to the terminals and where there is potential for increasing use of the Tunnel.

Market Growth and Market Share

Le Shuttle

The target market for Le Shuttle services is car, coach and accompanied roll-on roll-off ("RoRo") HGV traffic between continental Europe and UK ports between Plymouth and Newcastle. The Introduction of the Le Shuttle services has expanded the cross-Channel market and has diverted traffic away from other sea crossing routes to the Dover/Folkestone to Calais("DoCa") route as may be seen from Table 1.

		c growin a		a silai Co		
	Average annual growth of total	Share channe	Share of DoCa in total channel traffic			
	cross channel traffic 1994-96	1994	1995	1996		
Car	13%	54%	60%	68%		
Coach	9%	74%	78%	84%		
Accompanied HGV	5%	48%	58%	66%		

Table 1 Total cross-Channel traffic growth and DoCa shares

Eurotunnel's estimates of its annual market share of traffic on the DoCa route for cars, coaches and accompanied HGVs are set out in Table 2.

	19	1994		1995		96
	Vehicles '000	Market share (%)	Vehicles '000	Market share (%)	Vehicles '000	Market share (%)
Car	82	3%	1,224	30%	2,097	41%
Coach	24	13%	58	27%		
Accompanied HGV	65	8%	391	36%	519	41%

	T 01 //1	1	an a	
Table 2	Le Shuttle	share of	f DoCa	market

The Le Shuttle service has the following competitive advantages over conventional ferry services :

- *Speed:* The normal transit time between motorway in the UK and autoroute in France is generally shorter by Le Shuttle than the normal time by any competing cross-Channel sea service, with its quick loading and unloading of tourist vehicles.
- *Frequency of service:* The frequency of service both for Le Shuttle Tourist and Le Shuttle Freight is greater than that offered by each of Eurotunnel's competitors.

How and Why do Firms Hedge? | 5.31

• *Dependability:* Le Shuttle services are not affected by sea conditions and are less vulnerable to adverse weather conditions than ferries.

Eurostar

The market for Eurostar services is business and leisure passengers travelling in both directions between the UK and continental Europe. The market is varied, including business passengers travelling within the London/Brussels/Paris triangle, passengers travelling between the UK and France, Belgium, Holland and Germany, as well as some long distance travellers to or from destinations such as Spain, Italy and Eastern Europe.

The London-Paris and London-Brussels air passenger market is Eurostar's prime target. Eurostar services connect London with the major centres of Paris and Brussels with a service that competes on duration, comfort and price with airline services.

The growth of Eurostar's share of the city to city passenger market between London and Paris and London and Brussels is set out below

	1994		19	95	1996		
	Passengers '000	Market Share (%)	Passengers '000	Market Share (%)	Passengers '000	Market Share (%)	
London- Paris	108	3%	2,219	40%	3,672	56%	
London- Brussels	47	4%	700	37%	1,195	50%	
Total	155	3%	2,919	39%	4,867	54%	

Table 3 Eı	ırostar	market	share
------------	---------	--------	-------

Through freight services

Through freight services compete with most modes of sea and rail freight transport between the UK and continental Europe. Intermodal trains compete with load-on load-off ("LoLo") container services and with services for accompanied and unaccompanied freight offered by ferries and accompanied freight offered by Le Shuttle. Conventional rail wagon services compete with other freight modes, but often enjoy an advantage in the case of very heavy or large cargoes which can be loaded directly from a production facility. In the automotive market, competition is from combinations of road, sea and rail modes in the UK and in continental Europe which are used to deliver cars from factories to retail outlets.

The volume of freight transported by through freight services as a proportion of the target freight market is set out in Table 4 below.

rasic fillough height hiur her bhur e	Гable	Through freight market share	Ļ
---------------------------------------	-------	------------------------------	---

	1994		1995		1996	
	Million tonnes	Market share (%)	Million tones	Market share (%)	Million tones	Market share (%)
Through freight	1.1	2.3%	1.9	3.6%	2.4	4.4%

The share achieved by through freight services of its target market has remained below 5 percent, although it has continued to increase. Through freight services have, however, benefited from the

5.32 | Derivatives and Risk Management

transfer of traffic from train-ferry services since their termination in December 1995. They have also succeeded in penetrating part of the LoLo container market. Penetration of RoRo markets and the automotive market has, however, been lower than Eurotunnel's expectations. Existing road operators have sought to retain their business wherever possible and have been assisted by large tariff reductions on cross-Channel ferry and Le Shuttle Freight services. These tariff reductions have been facilitated by the current over-capacity on the DoCa route. There has been little development of through freight services since the start of the privatisation process of Railfreight Distribution in July 1996.

Traffic and Revenue Projections

The traffic and revenue projections made by Eurotunnel for 1997-99 are given in Tables 5 to 8 below.

	1996	1997	1998	1999				
	Actual							
Projected total traffic volumes on the DoCa route								
Car (million vehicles)	5.12	5.13	5.55	5.33				
Coach (million vehicles)	0.21	0.21	0.22	0.22				
Accompanied HGV (million vehicles)	1.28	1.50	1.63	1.77				
Total	6.61	6.84	7.40	7.32				
Projected Le Shuttle market share on the DoCa route								
Car	41%	49%	59%	62%				
Coach	29%	38%	45%	45%				
Accompanied HGV	41%	17%	45%	43%				
Total	41%	42%	56%	57%				

N N

 Table 6
 Projected total passenger and freight traffic volumes

	1996 Actual	1997	1998	1999					
Passenger Traffic									
Total passenger traffic(1) (million passengers	89	95	101	104					
London to Paris/Brussels passenger traffic(2) (million passengers)	9.0	10.6	13.4	14.6					
Freight Traffic									
RoRo accompanied (million tonnes)	25	27	29	31					
RoRo unaccompanied (million tonnes)	17	18	20	21					
Other freight (million tonnes)	13	14	15	16					
Total(3)	55	59	64	68					

(1) Total passenger traffic, travelling by air, Eurostar, or sea, between the UK and continental Europe.

(2) Total Eurostar passenger traffic, and air traffic between London and Paris or London and Brussels

(3) Total unitised freight and new vehicles traffic, transported by rail or sea, between the UK and continental Europe.

How and Why do Firms Hedge? | 5.33

	Actual	tual Base Case			Do	wnside C	ase		
	1996	1997	1998	1999	1997	1998	1999		
Eurostar ⁽¹⁾	54%	63%	72%	73%	63%	66%	64%		
Through freight ⁽²⁾	4.4%	4.7%	5.0%	5.1%	3.4%	3.3%	3.4%		

 Table 7
 Projected Eurostar and through freight market shares

(1) Market share of London to Paris/Brussels passenger traffic

(2) Market share of total unitised freight and new vehicles traffic, transported by rail or sea, between the UK and continental Europe.

	Actual		Base Case		Downside Case		
	1996	1997	1998	1999	1997	1998	1999
	£ m	£ m	£ m	£ m	£ m	£ m	£m
Le Shuttle	145	132	248	296	127	236	272
Railway Services	198	205	216	240	205	216	221
Retail and Other(1)	140	230	185	118	229	184	117
Total turnover	483	567	649	654	561	636	610

Table 8 Projected Group revenue

(1) Includes insurance receipts of £ 33 million in 1996 and £ 95 million in 1997.

Chapter Six

Options and their Payoffs

Objectives

Options are instruments that give the holder the right to trade without an obligation to do so. In this chapter, call options (right to buy) and put options (right to sell) will be described. The principal tool used in this chapter is a diagram showing the option value at maturity. Our analysis allows us to draw an analogy between options and insurance. This chapter does not deal with how options are valued, but analyses the determinants of option value, and also derives some important upper and lower bounds on an option value.

6.1 EUROPEAN AND AMERICAN CALLS AND PUTS

Options are instruments that give the holder the right to trade without an obligation to do so. Since a trade involves either buying or selling, there are two principal types of options. A call option is the right to buy an asset at a specified price, while a put option is a right to sell an asset at a specified price. In neither case does the holder have an obligation to buy or sell. The specified price at which the option holder has the right to trade is known as the strike price or the exercise price.

To complete the definition of the option, we need to specify when the holder can exercise his right to trade. All options have a maturity or expiry date after which the option to trade cannot be exercised. In American options, the holder can choose to trade at any time before the maturity date. In European options, the holder can trade only on the maturity date— neither before nor after it.

Thus, we have four types of options:

	Only on maturity date	On or before maturity date
Right to buy an asset	European Call Option	American Call Option
Right to sell an asset	European Put Option	American Put Option

Another point to be noted is that an option is itself an asset that can be bought and sold. Just as one can be a buyer or a seller of gold, one can be a buyer or a seller of a gold call option. It is the holder (buyer) of the option who decides whether to exercise it or not. The seller simply has to accept whatever decision the buyer takes and perform his side of the contract. Thus, the seller of an option has an obligation to trade when the buyer so demands, but he himself has no right to trade.

Since there are two kinds of options (put and call) and one can either be a buyer or a seller, there are four important possibilities.

Buyer		Seller
Call Option	Right to buy at the exercise price	Obligation to sell at the exercise price
Put Option	Right to sell at the exercise price	Obligation to buy at the exercise price

Each of these four possibilities must be distinguished from the situation of the buyer or seller of a forward contract.

6.2 | Derivatives and Risk Management

	Buyer	Seller
Forward Contract	Right and obligation to buy at the	Right and obligation to sell at the
	forward price	forward price

We now proceed to analyse and compare each of these possibilities to obtain a better understanding of what option contracts are all about.

6.2 PAYOFF AND PROFIT DIAGRAMS

We will now compare call and put options with forward contracts. Let us assume that the forward contracts and the options are on Infosys shares. Let the forward price be Rs 1,500 and let the strike prices of the options also be the same. Instead of trying to value these instruments, we will look at the position at the time of maturity of the contracts.

Consider a person who is long the forward contract. If the price of Infosys on the maturity date is exactly Rs 1,500, the forward contract makes no difference since the price under the forward contract is the same as the market price. If the price at maturity is higher, the long receives a positive payoff from the forward contract. For example, if the price is Rs 2,500, the payoff is +1,000 because he can buy at Rs 1,500 under the forward contract and sell in the spot market at 2,500. If the price at maturity is Rs 1,000 his payoff is -500 because he has to buy at Rs 1,500 stock that is worth only Rs 1,000. Every rupee rise in the price of Infosys at maturity translates into an additional rupee of payoff for the forward buyer and every rupee fall translates into a one rupee reduction in the payoff. This can be shown in a diagram as follows:



Figure 6.1 The payoff of a long forward at maturity is a straight line at 45°, indicating that a rise in the price of the underlying causes an equal gain in the forward, while a decline in price causes an equal loss

The position of a person who is short the forward contract is exactly the opposite of that of the forward buyer. If the price of Infosys on the maturity date is exactly Rs 1,500, the payoff is zero. Every rupee fall in the price of Infosys at maturity, translates into an additional rupee of payoff for the forward buyer and every rupee rise translates into a one rupee reduction in the payoff. The diagram shows this as follows:

Options and their Payoffs | 6.3



Figure 6.2 The payoff of short forward at maturity is a downward sloping straight line at 45°, indicating that a rise in the price of the underlying causes an equal *loss* in the forward and a decline in price causes an equal *gain*

When compared with the position of the buyer of a call option, the price of Infosys on the date of maturity is exactly Rs 1,500 and the payoff is zero as before. The payoff is zero, also if the price is below Rs 1,500 because in such a situation, the option holder has no obligation to exercise the call. He will simply let the option lapse and the payoff will be zero. If the maturity price is above Rs 1,500, the call buyer gets a positive payoff by exercising the option. If for example, the price is Rs 2,500, he will exercise the option at Rs 1,500, sell the share for Rs 2,500, and receive a payoff of Rs 1,000. Every rupee rise in the price of Infosys at maturity beyond Rs 1,500 translates into an additional rupee of payoff, but a fall below Rs 1,500 has no impact whatsovever. The diagram thus seems to be the payoff of a long forward contract, except that the negative part looks as if it has been chopped off and replaced by zero.



Figure 6.3 The payoff at maturity of a long call is zero below the strike price. The payoff of a long call thus has the upside of a forward without its downside

If a person has sold a call option, his payoff is zero if the price at maturity is less than or equal to Rs 1,500 because the buyer will not exercise the option. Every rupee rise in the price beyond this level, produces a negative payoff because he has to buy from the market at a high price and sell it to the option holder at the fixed price of Rs 1,500. The diagram thus seems to be the payoff of a short forward contract, except that the *positive* part looks as if it has been chopped off and replaced by zero:

The McGraw·Hill Companies

6.4 | Derivatives and Risk Management



Figure 6.4 Payoff at maturity of short call is zero below the strike price. Above the strike price, the payoff is like that of a short forward. The payoff of a short call thus has the downside of a short forward without its upside

We observe that the call buyer's payoff can never be negative; it is either zero or it is positive. On the other hand, the call seller's payoff can never be positive; it is either zero or it is negative. Why would one sell a call option if it can only lead to losses? Why would people not buy options if this can only lead to profits? The answer is that what we have drawn so far, are the payoffs of the option, and not really the complete profit or loss.

People sell options only if they are paid to do so. The option has a price or premium and the seller receives it when the option is sold. When the option payoff is zero, the seller actually makes a profit equal to the premium earned while writing the option. Even when the payoff becomes negative, losses start only when the negative payoff exceeds the premium. Similarly, whenever the option is not exercised the option has a positive payoff, there is a profit only after the premium has been covered. That is why all the diagrams discussed so far were called payoff diagrams and not profit diagrams. Let us now analyse the profit diagrams of the call buyer and seller, assuming that the call option cost Rs 150 at its inception:



Figure 6.5 If the underlying closes below the strike, the long call has a loss equal to the option premium. Above the strike, the profit diagram slopes upward at 45° and the position soon becomes profitable

Options and their Payoffs | 6.5

Turning now, to payoff and profit diagrams for put options, it can be recalled that the payoffs of call options were obtained from the payoffs of forward contracts by chopping off the positive or negative parts of their payoffs. The payoff of put options is obtained the same way, except that for long puts, one must start with short forward contracts and for short puts one must start with long forward contracts. This is opposite of what was done for call options.



Figure 6.6 If the underlying closes below the strike, the short call has a profit equal to the option premium. Above the strike, the profit diagram slopes downward at 45° and the position soon loses money

The payoff diagram of a *long put* looks like the payoff of a *short* forward contract, except that the *negative* part has been chopped off and replaced by zero. The payoff diagram of a *short put* looks like the payoff of a *long* forward contract, except that the *positive* part has been chopped off and replaced by zero. To move from payoffs to profits, we must add the put option premium in the case of short puts and subtract it for long puts. Assuming that the put option premium at inception was Rs 100, the payoff and profit diagrams are as follows:





The McGraw·Hill Companies

6.6 | Derivatives and Risk Management



Figure 6.8 If the underlying closes above the strike, the short put has a payoff of zero and a profit equal to the option premium. Below the strike, the profit diagram is similar to that of a short forward. The short put has the downside of a forward without its upside. The option premium is the price of bearing the whole risk

6.3 OPTIONS AS INSURANCE

The very term 'option premium' suggests a similarity to an insurance premium. We now explore this similarity further. In an insurance contract also, the insurance company collects a premium while selling the insurance and thereafter it cannot receive a positive payoff. It has a zero payoff if the insured risk does not materialize and suffers a negative payoff if the risk materializes. This is exactly the same situation as that of an option seller. Similarly, the position of an insurance buyer is akin to that of an option buyer—apart from the initial premium, there is no possibility of a negative payoff. If the risk materializes, there is a positive payoff, else, there is none.

Some greater similarities can be observed. If a person insures his beautiful Persian carpet against losses due to flood and, if a flood does happen, the insurance company will pay the assessed value of the loss. Alternatively, it may simply pay the full insured value of the carpet, take away the damaged carpet, and sell it at whatever price it can realize. In this case, the situation is as if the owner of the carpet had a put option to sell the carpet to the insurance company at its insured value.

A put option can thus be regarded as an insurance against a fall in the value of the asset. There are however, some significant differences. An insurance contract typically covers losses only due to some specific risks like floods, earthquake, etc., while a put option protects the holder against a reduction in the price of the asset, regardless of the cause of reduction. Insurance companies sometimes, have to deal with problems of moral hazard—for example, the Persian carpet may be ruined by someone who carelessly spilled coffee on it. Insurance companies therefore, restrict their coverage to well–defined risks that are largely outside the control of the insurance buyer. Options are generally written on assets whose values are not in general influenced by the actions of the option buyer. Option contracts, therefore, normally cover a loss of value regardless of the cause of loss.

Another difference is that option contracts can also protect a person from a rise in value. A person who does not own a Persian carpet may want to buy one and show interest in a call option that would allow him to buy the carpet a year later at a specified price. Thus option contracts can not only provide insurance against a fall in value, but also against a rise in value.

Nevertheless, the similarities between insurance and option price are sufficiently strong so as to list down the various determinants of the latter by first looking at the determinants of an insurance premium:

Options and their Payoffs | 6.7

- 1. The insurance premium rises with the probability of the risk involved. Premium for flood insurance will be high near a river bank and low on top of a mountain. Premium for fire insurance will be higher for a house made of wood, than for one made of bricks.
- 2. The insurance premium will be higher if the magnitude of the loss is higher. It will cost more to insure a home for Rs 2 million than to insure the same home for Rs 1 million.
- 3. The insurance premium for a one year policy would be more than that of a six months policy. In most cases, the premium for one year would be double than that of six months.
- 4. Finally and least obviously, the premium also depends on the interest rate. The insurance company collects the premium upfront and pays the claim much later. In the meantime, it earns interest on the premium collected. In life insurance, the impact of the interest rate is very high. In general insurance, the impact is lower, but it is still quite common for insurance companies to have a claim payout greater than the premium and make profits only from the investment income.

We shall now see how we can apply these same insights to option pricing and enumerate the determinants of the option premium.

6.4 DETERMINANTS OF OPTION PRICE

There are six important determinants of the option price:

- 1. Volatility corresponds closely to the probability of risk in insurance. In options, the risk covered is the fluctuation in the price of the asset. A call option protects against upward movement of the price and a put option protects against downward movements of the price. In either case, the extent of fluctuation of the share price determines the risk. This is measured by the volatility which is essentially the standard deviation of price changes. Volatility will be defined more precisely later. Higher the volatility, higher is the option premium.
- 2. Time to maturity is clearly a determinant of the option premium as in the case of insurance. It is easier for an asset price to drop by 50 percent in one year than in six months. Normally, the longer the time taken for maturity, the higher is the option premium. For American options, this is obvious as the holder of a one year American option can do everything that the holder of a six month option can. The situation is less clear cut for European options. In this case, the holder of a one year option cannot exercise at the end of six months and exceptions do arise to the intuitive idea that longer maturities require larger premiums. Firstly, a European call option expiring before a large dividend is paid out, can be worth more than a call expiring after the dividend is paid out. This is because the dividend reduces the value of the call as discussed below. The second exception is that a short maturity European put option can be worth more than a longer maturity European put option. This is also discussed later.
- 3. The exercise price of a put option is like the insured value of the asset. The higher this is, the higher the option premium would be. Another way of looking at this is that the higher the exercise price, the easier it is for the asset price to fall below this level. The risk of loss is higher and the option seller would, therefore, demand a higher premium. The situation is the reverse for a call option. Here, the higher the exercise price, the lesser the risk that the asset price would climb above that level. It therefore follows that there is a lower the probability that the option would be exercised. Thus the option premium is positively related to the exercise price for a put option and negatively related to a call option.
- 4. The current asset price is equally important as it is the reference point against which the exercise price must be compared. The higher the asset price, the lower is the chance that it will fall below

6.8 | Derivatives and Risk Management

the exercise price and therefore, the lower the premium for a put option. On the other hand, a high asset price makes it more likely that the asset price will move above the exercise price and therefore, the higher the premium for a call option.

- 5. The interest rate is also a determinant of the option premium. In the case of insurance, we saw that high interest rates help reduce the premium. This is true for put options but not for call options. The relationship is complicated by the fact that the expected movement of the asset price is also influenced by the interest rate. When the interest rate is high, the expected return on all assets has to be commensurately higher. If we assume that the current stock price is unchanged, then future stock prices are likely to be higher. For put options, this reduces the chance that the asset price will drop below the exercise price, and thus reduces the value of the put option. For call options, the expected higher return increases the chance that the asset price will rise above the exercise price, and thus it increases the value of the call option.
- 6 Dividends that are expected during the life of the option are another determinant of the option price. When the dividend is paid, the stock price falls. This makes it more likely that the stock would fall below the exercise price and thus raises the put option premium. In the case of the call options, the dividend reduces the chance that the asset price will rise above the exercise price, and thus decreases the value of the call option.

In Chapter 10, when we discuss the Black–Scholes option pricing model, we would also be able to quantify the effect of these determinants of option prices accurately. What we shall do now, is to derive some bounds within which the option price should lie.

6.5 BOUNDS ON OPTION PRICE

6.5.1 Formulas for Option Payoffs

The diagrams for the payoffs of the call option, describe the payoff by the mathematical formula max $(0, S_T - X)$ where S_T is the price of the underlying asset at maturity and X is the exercise price. What this formula says is that when $S_T > X$ or $S_T - X > 0$ the option would be exercised and the payoff is $S_T - X$. On the other hand if $S_T > X$ or $S_T - X < 0$ the option would not be exercised and the payoff is 0. In the former case, max $(0, S_T - X)$ is $S_T - X$ while in the latter case, it is 0. Thus, in both cases the formula gives the right answer.

The payoff for the call option seller is the negative of the option buyer's payoff and can be written as $-\max(0, S_T - X) = \min(0, X - S_T)$. Similarly, we can see that the payoff for the put option buyer is max $(0, X - S_T)$ while the payoff for the seller is $-\max(0, X - S_T) = \min(0, S_T - X)$.

This is summarized in the following table:

	Buyer	Seller
Call Option	$\max\left(0,S_T\!-\!X\right)$	$\min\left(0,X-S_{T}\right)$
Put Option	$\max\left(0,X-S_{T}\right)$	$\min\left(0,S_T-X\right)$

6.5.2 Lower Bound for Call Option Price

Consider now the current stock price S instead of the price S_T at maturity. It is evident that max (0, S - X) is a lower bound on the current value of an American call option because such an option can be exercised immediately to attain this value. If we choose not to exercise the option, it must be because an even higher value is expected in future. Max (0, S - X) is a lower bound on the option value even when

Options and their Payoffs | 6.9

the option is a European call, assuming that the stock or other underlying for the option does not pay any dividends. If max (0, S - X) = 0, it is clearly a lower bond because the option has only a right and no obligation. One can always throw away the option and so it can never have a negative value. Now consider the situation where S > X and max (0, S - X) = S - X. By exercising the option now, the holder buys the stock immediately for a price of X. It he holds on to the option, he could buy the stock for X or lower. If the market price at maturity is lower than X, the option could be discarded and bought at that lower price. Otherwise, the option could be used to buy at X. Thus there are two advantages in keeping the option instead of exercising it immediately:

- 1. By waiting till maturity, the holder may be able to buy the stock at a price less than X. By exercising the option, X has to be paid.
- 2. By exercising the option immediately, *X* is paid immediately, while by holding on to the option, the payment is deferred till maturity of the option. Thus the interest on *X* is saved till the maturity date.

This analysis points to the following:

- 1. An American call option should never be exercised before maturity; it should be treated as a European option.
- 2. The value of the option is always greater than or equal to max (0, S X).

The quantity max (0, S - X) is known as the intrinsic value of the call option. The assertion can thus be rephrased as 'a call option is worth at least its intrinsic value'.

We can get an even better lower bound by considering the interest loss on the exercise price. Consider the following two alternatives:

- 1. Buy the stock now by paying its price *S*.
- 2. Buy the option and invest enough money to ensure that the principal plus interest will be equal to X at the maturity of the option. The investment required is the present value of X or $X e^{-rt}$ where r is the rate of interest continuously compounded and t is the time to maturity. If C is the price of the call option, then the total investment required is $C + X e^{-rt}$.

The second alternative will ensure that one is able to buy the stock at the maturity of the option without further expenditure of money. If the stock price at maturity is greater than X, the option to buy the stock at the price X is used, which would be paid for by selling the bond which was invested in $X e^{-rt}$. If the stock price at maturity is less than X, the stock can be bought from the market at a price less than the value X which is obtained by selling the bond. In this case, one has some money left over after buying the stock.

The second alternative is thus better than the first. In both alternatives, the stock remains at the maturity of the option, but in the second case, there is a chance that there is surplus cash as well. Therefore the cost of the second alternative $C + X e^{-rt}$ must be greater than the cost of the first alternative which is S. This implies that $C + X e^{-rt} \ge S$ or

$$C \ge S - X \,\mathrm{e}^{-rt} \tag{6.1}$$

This gives another lower bound—the price of the call option is greater than $S - X e^{-rt}$. In addition, the option price can never be negative as already argued. This gives us a better lower bound than the earlier lower bound:

$$C \ge \max(0, S - X e^{-rt})$$
 (6.2)

6.10 | Derivatives and Risk Management

6.5.3 Lower Bound for Put Option

To get the lower bound for the price of a put option, it is assumed that one is currently holding the stock and comparing two alternatives:

- 1. Sell the stock now for its price S.
- 2. Buy the put option and borrow enough money to ensure that the repayment (principal plus interest) at the maturity of the option will be equal to X. The borrowing required is the present value of X or $X e^{-rt}$ where r is the rate of interest continuously compounded and t is the time to maturity. If P is the price of the call option then the total cash inflow now is $X e^{-rt} P$.

The second alternative ensures that one is able to repay the borrowing by selling the stock at the maturity of the option without any further expenditure of money. If the stock price at maturity is less than X, the option to sell the stock at the price X is used and the proceeds repay the loan. If the stock price at maturity is greater than X, the stock is sold in the market at a price greater than the value X which is required to repay the loan. In this case, some money would be left over after repaying the loan.

The second alternative is therefore, better than the first. In the first alternative, there will be no cash inflow in future, but in the second case, there is a chance that there will be surplus cash at the time of maturity of the option. This possibility of an additional cash inflow in future has to be compensated by a lower cash inflow now. Therefore, the immediate cash inflow under the second alternative $X e^{-rt} - P$ must be less than the immediate cash inflow under the first alternative, which is *S*. This implies that $X e^{-rt} - P \le S$ or $P \ge X e^{-rt} - S$

An option price can never be negative as it can always be discarded (it is only a right and has no obligations), one can, therefore, state the lower bound for a put options as:

$$P \ge \max(0, X e^{-rt} - S)$$
 (6.3)

As in the case of call options, one can also say that an American put option is worth at least max (0, X - S) because this value can be obtained by immediate exercise. If one chooses not to exercise, it must be because one hopes to extract more from it in future.

For a European put, however, one has to be content with the weaker lower bound $\max(0, X e^{-rt} - S)$. The lower bound $\max(0, X - S)$ is not in general valid for European put options. In other words, a European put option can be worth less than its intrinsic value and is therefore, less valuable than an American option. This implies that an American put option may sometimes be exercised before maturity. The reason is that in the case of the put, immediate exercise has a benefit in terms of receiving the intrinsic value immediately and there is interest to be gained in doing so.

6.5.4 Upper Bound for Call Option

A call option is always of less value than the stock price itself, unless the exercise price is zero. This is because if a person buys the option, then at maturity he has to spend an additional amount to buy the stock. This additional amount is either the exercise price or the stock price at maturity, whichever is lower. One thus obtains the upper bound $C \le S$. Combining this with the previous lower bounds and the known equality of the American and European call options, the following is obtained:

$$\max(0, S - X) \le \max(0, S - X e^{-rt}) \le C_E = C_A \le S$$
(6.4)

where C_E is the price of a European call option and C_A is the price of an American call option. As already stated, the lower bound here assumes that the stock or other underlying for the option does not pay any dividends.

6.5.5 Upper Bound for Put Option

The American put option is always of less worth than the exercise price, unless the stock price drops to zero. This gives the upper bound $P_A = X$ where P_A is the price of an American put option. Combining this with the intrinsic value lower bound which is valid for the American option, we obtain the bounds:

$$\max(0, X - S) \le P_A \le X.$$

The European put option is always of less worth than the present value of the exercise price, unless the stock price is guaranted to be zero at maturity of the option:

$$P_F \leq X e^{-r}$$

Combining this with the lower bound for the European put, the following is obtained:

$$\max(0, X e^{-rt} - S) \le P_E \le X e^{-rt}$$
(6.5)

6.6 PUT CALL PARITY

It has already been discussed that the payoff diagram of a long call option is just the positive part of the payoff of the long forward contract, with the negative part chopped off. Similarly, the payoff diagram of a short put option is just the negative part of the payoff of the long forward contract, with the positive part chopped off. So, if these are combined, the payoff of a long forward contract is obtained. Another way of looking at it is that the long call gives one the right to buy (but there is no obligation to buy) while the short put gives one the obligation to buy (but not the right to do so). If the two are combined, the right and obligation to buy are obtained. This is the same as a forward contract to buy the asset at the specified price. This is one version of the put call parity.

The put call parity can now be restated in terms of the current stock prices, instead of a forward contract. Consider buying the forward contract and simultaneously investing enough money to ensure that the principal plus interest would be equal to the forward price at the maturity of the forward contract. The investment required is the present value of X or $X e^{-rt}$ where r is the rate of interest continuously compounded and t is the time to maturity. This combination of a forward contract and bonds of $X e^{-rt}$ is equivalent to buying the stock instantly if the stock pays no dividend. This is because at maturity, the bonds can be sold to pay for the price to be paid under the forward contract and one ends up with the stock without any further payment of cash.

Combining these two ideas, we conclude that buying the stock is equivalent to the following combination:

- 1. Long a call at an exercise price X
- 2. Short a call at an exercise price X
- 3. Long bonds worth $X e^{-rt}$

We can thus write put call parity as:

$$S = C - P + X e^{-rt}$$
 or as $S + P = C + X e^{-rt}$ (6.6)

where S is the current price of the stock, C is the price of the call option, P is the price of the put option, X is their common exercise price, r is the continuously compounded rate of interest and t is the time to maturity.

This can be interpreted as follows. The left hand side is buying a stock and protecting it against downside risk with a put option. This leaves us with the upside of the stock. The right hand side contains

6.12 | Derivatives and Risk Management

the call option which gives us the upside of the stock and it also contains enough bonds to pay the exercise price of the call option. To see that these two are equivalent, consider the two possibilities of the share price at maturity:

- 1. The stock price at maturity S_T is less than X. In this case, on the left hand side, the put option would be exercised to receive X. On the right hand side, the call option would not be exercised while selling the bonds gives X. Thus in this case, the two sides are equal.
- 2. The stock price at maturity S_T is more than X. In this case, on the left hand side, the put option would not be exercised and the stock with a value S_T would continue to be held. On the right hand side, the call option would be exercised by selling the bonds to pay for the exercise price. Again, the result is a holding of stock with a value S_T . Thus the two sides are equal in this case too.

Chapter Summary

Options are instruments that give the holder the right to trade without an obligation to do so. In this chapter, we have discussed four types of options:

	Only on maturity date	On or before maturity date
Right to buy an asset	European Call Option	American Call Option
Right to sell an asset	European Put Option	American Put Option

Correspondingly, the seller of an option has an obligation to trade without any right to do so.

Diagrams showing the payoffs at maturity from call and put options have already been discussed. These payoffs can be described by the following formulas:

	Buyer	Seller
Call Option	$\max\left(0,S_T-X\right)$	$\min\left(0,X-S_{T}\right)$
Put Option	$\max\left(0, X - S_T\right)$	$\min\left(0,S_T\!-\!X\right)$

Payoff is never negative for the option buyer and never positive for the option seller. Payoff diagrams must be distinguished from profit diagrams which include the effect of the initial option price that is paid by the buyer and received by the seller.

Options have great similarities to insurance contracts. A put option can be regarded as an insurance against a fall in the value of the asset and a call option as an insurance against a rise in the value of the asset.

The most important determinants of the value of an option are: the current price of the underlying, the exercise price, the rate of interest, the volatility of the underlying and the time to maturity.

The option value must be within the following bounds:

$$\max (0, S - X) \le \max (0, S - X e^{-rt}) \le C_E = C_A \le S$$
$$\max (0, X e^{-rt} - S) \le P_E \le X e^{-rt}$$
$$\max (0, X - S) \le P_A \le X$$

where C_E is the price of a European call option, C_A is the price of an American call option, P_E is the price of a European put option, P_A is the price of an American put option, X is the exercise price, S is the current price of the underlying, r is the interest rate and t is the time to maturity. The lower bounds for the call option assume that the stock or other underlying for the option does not pay any dividends.

Put call parity is a relationship between the prices of put and call options and the price of the underlying.

$$S + P = C + X e^{-r}$$

where S is the current price of the stock, C is the price of the call option, P is the price of the put option, X is their common exercise price, r is the continuously compounded rate of interest, and t is the time to maturity.

Put call parity says that buying a stock and protecting it against downside risk with a put option is the same as buying the upside of the stock with a call option, and buying enough bonds to pay the exercise price of the call option at maturity.

Suggestions for Further Reading

The discussion in this chapter is continued in Chapters 9 and 10. Suggestions for Further Reading are given at the end of those chapters.

Problems and Questions

- 1. A person pays Rs 1.50 for a call option with a strike of 109. What should the price of the underlying be at maturity for her to break even?
- 2. A person who has bought an option calculates that he would make a loss of Rs 2 if the price of the underlying is Rs 50 at maturity but will gain Rs 2 if the price is Rs 55. What is the strike price of the option and what is the price paid for it? Is the option a call or a put?
- 3. A person who has bought an option calculates that she would make a gain of Rs 3 if the price of the underlying is Rs 100 at maturity but will lose Rs 5 if the price is Rs 110. What is the strike price of the option and what is the price paid for it? Is the option a call or put?
- 4. Which of the following one-month options is incorrectly priced if the interest rate is 12% p.a. on simple interest basis and the price of the underlying is 100?

Option type	Strike price	Option price
American Call	110	15.00
American Put	110	9.50
European Put	110	9.50
European Call	90	10.50

- 5. When comparing two call options on the same underlying at the same maturity but with different strikes, you find that both options have the profit if the underlying is at 100.25 at maturity. Which of the following statements is true:
 - (a) The higher priced option has a strike below 100.25 and the lower priced option has a strike above 100.25
 - (b) The lower priced option has a strike below 100.25 and the higher priced option has a strike above 100.25
 - (c) Both strikes are at or below 100.25
 - (d) Both strikes are at or above 100.25
- 6. What would your answer be to the previous problem if both options were puts instead of calls?
- 7. Rank the following options from the most expensive to the least expensive:
 - (a) European put with a strike of 100 and maturity of three months
 - (b) European put with a strike of 110 and maturity of one month
 - (c) European put with a strike of 105 and maturity of one month
 - (d) American put with a strike of 110 and maturity of three months

Chapter Seven

Option Markets

Objectives

This chapter provides information about the important options markets worldwide in commodities, equities, bonds, interest rates and currencies. It discusses the issues in the design of option contracts, the choices regarding option exercise and settlement, and the listing of different option strikes. It also provides information about the size of different option markets.

7.1 EXCHANGE TRADED OPTIONS AND OTC OPTIONS

In Chapter 2, a key difference discussed between futures and forwards, was that futures are traded in exchanges while forwards are either bilaterally negotiated, or traded in OTC markets. In the case of options, the same term 'option' is used to refer to both exchange traded contracts and bilateral or OTC contracts. However, some of the key differences discussed in Chapter 2 between exchange traded and OTC contracts are equally valid for options.

- **Transparency:** In case of exchange traded options, there is a great deal of transparency about past trades, prices and volumes.
- **Standardization:** Exchange traded options are standardized in many respects. First, expiry dates are standardized. Second, contracts are also standardized in terms of the type and quality that is to be delivered. Strike prices are also standardized.
- Initial Margins and Mark to Market: An option buyer can at worst, lose the premium that has already been paid. Once the option seller has collected the premium, he does not have to worry about any further credit risk. The option buyer however faces credit risk. Exchanges deal with this risk by imposing mark to market requirements on option sellers similar to that imposed on futures traders. Option buyers may also choose to subject their entire option portfolio to mark to market obligations, particularly when they have sold some options and bought some others.
- Novation: Novation works for options, just as it does for futures: the exchange becomes the counter party for all trades.
- Anonymous Trading: One advantage of novation is that it allows anonymous trading.
- Easy Close Out: Another advantage of novation is that it makes it very easy to unwind an options transaction.
- **Delivery Rare:** One important consequence of standardization and ease of unwinding is that most exchange traded transactions are unwound prior to maturity.

7.2 OPTIONS MARKET CONTRACT DESIGN

An exchange that seeks to introduce an options market on any underlying has to make several choices in designing the contact. Exchange traded contracts are highly standardized and it is up to the exchange to achieve this standardization.

7.2 | Derivatives and Risk Management

Some of the design choices are the same as in futures markets and the reader may refer to Chapter 2 for a more detailed discussion of these choices:

- **Underlying:** The exchange has to choose the precise quality and grade of the asset on which the option contract would trade.
- Contract Size: The contract size is the minimum unit in which trading can take place.
- **Contract Months:** The maturity date is usually standardized with only one contract for each maturity month. The number of such contract months that trade at the same time is a choice that has to be made.
- **Position Limits:** To prevent one or two very large players from manipulating the markets, many exchanges place limits on the maximum position that can be taken by a single person. This varies from contract to contract-based on the size of the overall market, and the ease of manipulating the market.
- Price Quotes: Options markets typically quote prices as the option premium per unit of the underlying. The price per contract is then obtained by multiplying the quoted price by the contract size. In some cases, however, the price is quoted in units of 0.01 or 0.001. For example, the currency options on the euro at the CME in the US are stated in terms of 0.001. A quote of 20.30 really means an option premium of 0.023 dollars or 2.3 cents per euro. As the contract is for €125,000 the premium per contract is \$2,537.50.
- **Price Limits:** In some options markets, there is a limit on the maximum price movement that is permitted in a single day.
- Price Increments: The CME currency options on the euro discussed earlier has a minimum price increment of \$0.0001. Since the contract size is €125,000, the minimum price increment translates into \$12.50 per contract.
- Exercise and Assignment: The precise mechanisms for exercise have to be specified. This is discussed in detail in the next section.
- Strike Price Intervals: Options have an exercise price or strike price. Thus, for the same underlying and expiration date, an infinite number of different option contracts can be created by simply varying the strike price. Exchanges have to decide on the number of strikes to be included at inception and also lay down rules for the introduction of new strikes when the underlying drifts away from existing strikes.

7.3 OPTION EXERCISE, SETTLEMENT, AND ASSIGNMENT

7.3.1 Exercise Style

In terms of exercise, options can be either European or American. Typically, single stock options tend to be American and index options tend to be European. In general, options on most other underlyings tend to be American as well.

7.3.2 Settlement

There are three choices when it comes to settlement:

1. **Physical settlement:** This is the most intuitive method in which option exercise leads to giving or taking delivery of the underlying asset. Most options on individual stocks are physically settled.

Option Markets | 7.3

The McGraw·Hill Companies

- 2. **Cash settlement:** In this, there is no delivery of the underlying, but a cash settlement is made on the basis of the difference between the strike price and the current market price of the underlying. For example, if a call option with a strike of 120 is exercised when the underlying is traded at 135, a cash settlement of 15 has to be made to the option holder by the option seller. The option holder then needs 120 of his own funds (in addition to this 15) to buy the underlying in the market. Physical settlement would also have required the holder to spend 120 of his own funds to exercise the option and receive delivery of the underlying. Most index option contracts are cash-settled.
- 3 Options on futures: Most exchange traded option contracts are not on the underlying itself but on futures on the underlying. For example, the Chicago Board of Trade offers American style options on its corn futures contracts. In this case, when the corn option is exercised, the option seller gives or takes delivery of a corn futures contract. This is physical settlement in a strict sense, as the underlying for the option contract is corn futures and that is exactly what is delivered. However, if the option is loosely regarded as an option on corn itself, then delivery of corn futures can be regarded as another form of settlement. This could potentially avoid some of the difficulties of actually delivering 5,000 bushels of corn. However, the expiry of the option contract is usually (but not always) also the expiry of the underlying futures. Therefore if the corn futures option is exercised close to expiration, the holder may end up having to go through the physical settlement of the underlying futures.

7.3.3 Assignment

In a bilateral option contract between a buyer and a seller, exercise of an option is a straightforward matter of the option seller performing his obligation to the option buyer. But in an exchange traded contract, if one option holder exercises the option, the exchange needs to decide which of the thousands of option sellers should be called upon, to fulfil the obligation under the option. The problem arises because not all option holders may choose to exercise the option at any given point of time. Even if all option holders are perfectly rational, it is not necessary that when one holder exercises, the other holders should also choose to do the same. Different holders may face different transaction costs of buying and selling the underlying, and may also face different borrowing and lending rates for both cash and securities.

Consider a very simple example where the strike price of a call option is 101. At expiry, the underlying closes at a mid price of 101.05 so that the option is in the money. Does this mean that all option holders should exercise the option? To analyse this situation completely, assume that the bid and ask prices at expiry are 100.95 and 101.15 respectively. An option holder who needs the underlying would reason as follows—if he lets the option lapse and buys the underlying in the market, it would cost him 101.15. If he exercises the option, he would acquire the underlying at a cost of 101. It is therefore, cheaper to exercise the option.

Consider now a holder who does not actually need the underlying. He will reason as follows—If he exercises the option, it will cost him 101 and when he turns around to sell the unwanted underlying, he will realize only 100.95. He is therefore, better off letting the option lapse. Even if all holders are rational, some may choose to exercise and some may choose to let the option lapse.

What does the exchange do under this situation? When some holders exercise and some do not, the exchange usually randomly assigns the exercise to some option sellers. The randomly chosen sellers have to fulfil their option obligations, while the rest continue to have a short position in the option.

7.4 | Derivatives and Risk Management

Option assignment creates problems for option sellers, particularly in the case of cash-settled American options. Consider a put option seller who has hedged her short option position by taking a short position in the underlying. If she is assigned an option exercise, she no longer has a short option position and therefore does not need the hedge. In practice, she would be able to lift the hedge only the next day and would be left with an overnight exposure on the unwanted hedge. The problem is mitigated when the option is physically settled as in this case, she receives the underlying as part of the option exercise. This creates a long position in the underlying, which partially offsets the short position that existed in the hedge. Most American options for this reason, are physically settled or are options on futures.

7.4 OPTIONS MARKETS

7.4.1 Understanding Options Markets Quotes

The end of day summary for an options contract (the S&P 500 March 2007 1440 call contract) is shown in Table 7.1. The first few items do not need much explanation as they are the familiar open-high-low-close prices that are published for stocks and other securities.

The settlement price is used to compute the mark to market margins and is therefore very important. Unlike in futures markets, the settlement price for options is neither the last traded price, nor an average of the last few minutes of trading. This is because options are much less liquid than futures and it is quite possible that the underlying has moved significantly since the last option trade.

Suppose that an option with a strike of 1300 had its last trade of the day at 140 when the underlying was at 1430. At this point, the option had an intrinsic value of 130 (1430-1300) and the additional 10 represents the time value of the option. Suppose thereafter, the underlying continued to trade and ended the day at 1440. At this level of the underlying, the 1300 call has an intrinsic value of 140. Since the option has some intrinsic value, the last traded price of the option of 140 is clearly a stale price. Most option exchanges use a theoretical option price model to arrive at the settlement price that reflects the closing price of the underlying, as well as the traded price of other options on the same underlying that might have traded more recently than the 1300 call. From a risk management point of view, it is imperative for the exchange to do this. The settlement price may not be perfect, but this does not matter, as errors are likely to be reversed in the near future, and at any event at expiry, or at exercise.

The 'price change' is the change in the settlement price. The change of 320 points or 3.20 implies that the settlement price on the previous day was 3.60. Every short (seller) would face a mark to market outflow because of the price rise. Longs who subject their portfolios to mark to market, would receive a mark to market inflow for the same reason. The S&P 500 Call is actually an option on the S&P 500 futures contract and its price is quoted in index points in the same way as the future. As explained in Chapter 2, the S&P 500 index contract has a multiplier of \$250 and therefore the mark to market for a long holding one contract on this day would be \$250 times 3.20 or \$800.

The next item (open interest) is similar to that in futures markets as discussed in Chapter 2. The open interest is defined as the aggregate of all long positions or equivalently, as the aggregate of all short positions. The total of all long positions is always equal to the total of all short positions, so that the net stock of options positions (long minus short) is always zero. The open interest is thus a measure of how much hedging or speculative interest there is in that contract. In Table 7.1, the open interest is 4,563 contracts up six contracts from the previous day.

In stock markets, people often track the 52-week high-low of a stock. In the derivatives market, it is customary to track the high-low during the life of the contract. This is because many contracts have a life

of only a few months and therefore, the 52-week high-low is meaningless. In this example, the contract has traded in a range of 1.50 to 17.00 during its life.

Another quantity which is often mentioned in this context, is the notional principal value of the contract. Notional value of options contracts is defined in terms of the price of the underlying and not the price of the option itself. If the settlement price of the S&P 500 March 2007 futures was 1440.80, then based on the index multiplier of \$250, the notional principal value of the 1440 call option contract is:

$$1440.80 \times \$250 = \$360,200$$

Thus, the open interest of 4,563 contracts represents a notional principal value of 4,563 times \$360,200 or \$1.6 billion. Similarly, the traded quantity of 845 contracts represents a traded notional value of 845 times \$360,200 or \$0.3 billion. However, based on the option price of 6.80, the option premium that changed hands while trading these 845 contracts is much lower. It is only $845 \times 6.80 \times 250 = 1.4 million

S&P 500 March 2007 Call	s
Strike	1440
Open Range	3.00
High	7.20
Low	3.00
Close Range	7.20
Settlement	6.80
Price Change (points)	320
Delta	0.518
Exercises	
Total Volume	845
Open Interest	4,563
Change in Open Interest	+6
Contract High	17
Contract Low	1.5

Table 7.1S&P 500 March 2007 1440 Call at the CME on 12, January 2007. This table
gives price, volume, and open interest data for the contract

7.4.2 Global Options Market Size

At the end of 2005, the world's financial options markets had nearly \$100 trillion of notional value of options outstanding. This represented well over twice the world GDP. Nearly three-fifths of this value was in the derivative exchanges and the balance in the OTC options market. As in the case of the futures market, interest rate products dominate the options markets too. More information is provided in Figure 7.1.

Information about the number or value of options traded during the year is not available for OTC options. Table 7.2, therefore, deals only with exchange traded options. This table shows that the composition of the market changes dramatically when size is measured by the number of contracts instead of notional value. This is because single stock options tend to be of a smaller size. In the US, these option contracts are typically for a hundred shares representing a notional value of a couple of thousand dollars. By contrast, the short-term interest rate futures at the CME have a notional value of a

The McGraw·Hill Companies

7.6 | Derivatives and Risk Management



Figure 7.1 Composition of Open Interest in the world options markets. The open interest in the exchanges is about 40 percent larger than that in the OTC markets. Interest rate options dominate both markets. Currency options trade almost entirely in the OTC markets. *Source:* Bank for International Settlements, BIS Quarterly Review, June 2006 supplemented by partial data on the notional value of single equity options from the World Federation of Exchanges

million dollars, and most other option contracts tend to have notional values close to a hundred thousand dollars.

Table 7.2 Composition of exchange traded options markets. Interest rate productsdominate the market by traded value and notional value of open interest. In terms of thenumber of the contracts, equity products dominate because single stock equity contractstend to be of a smaller size (typically only a couple of thousand dollars). Commoditiesaccount for a small fraction of the market

	Millions of Contracts		Notional Value \$ Trillion	
	Open Interest	Turnover	Open Interest	Turnover
Interest rate	32	431	48.0	328.8
Currency	1	19	0.1	0.9
Equity	516	5,587	6.4	76.5
All financial	549	6,037	54.5	406.2
Commodities	11	809		

7.4.3 Indian Options Market Size

As of 2006, exchange traded options in India are confined to the equity markets. While futures have been permitted in commodities, options are not allowed in this segment. Interest and currency derivatives are not traded in exchanges. There is a significant OTC market for currency options. Indian companies have access to various global option markets to hedge commodity and interest rate risk.

	-		
	Index Options	Stock Options	All Options
Contracts traded in 2006 (millions)	18.702	5.214	23.916
Traded Value in 2006 (rupees billion)	6,381	2,011	8,392
Open Interest end 2006 (millions of contracts)	0.419	0.094	0.514
Open Interest end 2006 (rupees billion)	167	34	200

Table 7.3 Size of the Indian Equity Options Market in 2006. The index options marketis over three times bigger than the stock options market in terms of trading,
and over four times bigger in terms of open interest

7.5 COMMODITY OPTIONS

In India, as of 2006, options were not permitted on commodities, but they are likely to be permitted in the near future. Globally, however, options are traded on three broad groups of commodities:

- Agricultural: The CBOT trades options on food grains (wheat, rice and corn), oilseeds (soya), dairy products (milk and eggs), and livestock products (live cattle). The New York Board of Trade (NYBOT) trades contracts on food products like coffee, sugar and cocoa. In Europe, Euronext.liffe trades contracts on wheat, corn, coffee, sugar and cocoa.
- **Energy:** The New York Mercantile Exchange (Nymex) in the United States and ICE Futures (previously known as the International Petroleum Exchange) in London trade contracts on crude oil, natural gas, and electricity.
- **Metals:** The London Metal Exchange trades contracts on most important non-ferrous metals like aluminium, copper, tin, nickel, lead, and zinc as well as some important plastics. Precious metals like gold and silver are traded at the CBOT and the Nymex.

The US exchanges as well as Euronext, trade options on commodity futures while the London Metal Exchange trades options on the commodity itself. Most of these options have American style exercise and trade contracts up to a maturity of less than a year.

7.6 STOCK INDEX OPTIONS

Stock index options are among the most widespread option markets in the world. Most developed equity markets the world over, trade options on their equity indices. Like index futures, index options are also cash-settled. The principal exception is the index options at the CME which are settled by delivery of the corresponding index futures contract.

Most index options tend to be European, but there are some American options as well. All the issues discussed in Chapter 2 about the mechanics of index futures, particularly the treatment of dividends, are relevant to index options as well.

Many exchanges trade index options to maturities of one year or less, but US exchanges also provide options known as Long-term Equity AnticiPation Securities or LEAPS, with maturities of two to three years. In Europe, Eurex provide options with even longer maturities. Even in these cases, liquidity tends to be concentrated in near-term contracts.

Exchanges try to provide five to seven strike prices at inception with two to three prices being out of the money, two to three prices being in the money, and one strike being close to the money. As the

7.8 | Derivatives and Risk Management

underlying trades away from the original range of strikes, new strikes are added. This means that after a period of time, a wide range of strikes are available. Some exchanges follow a method of listing a large range of strikes at inception. Some exchanges also list options with a near-zero strike price (one tick away from zero). In all cases, trading tends to be concentrated in strikes that are slightly out of the money. Index options like index futures often trade without any position limits.

7.7 INDIVIDUAL EQUITY OPTIONS

Individual equity options are also quite widespread around the world. Retail investors are particularly attracted to these options. While most other options are dominated by institutional investors, companies, and financial institutions, equity options in many markets are dominated by retail investors. This is reflected in their small contract size.

Equity options globally tend to be physically settled American-type options, though there are some exceptions. In India, for example, equity options (like other equity derivatives) are cash-settled.

At inception, exchanges list contracts with in-the-money, at-the-money, and out-of-the-money strike prices. New series are generally added when the underlying trades through the highest or lowest strike price available. This means that after a period of time, a wide range of strikes are available. Trading tends to be concentrated in strikes that are slightly out of the money.

Equity options are usually subject to stringent position limits, related to the stocks' market capitalization and trading volume. In some countries like India, equity options are permitted only on the large, liquid stocks.

7.8 CURRENCY OPTIONS

Though currency options trade on the exchanges, particularly the CME and Euronext.liffe, the bulk of the trading in currency products takes place in the over-the-counter market.

7.9 INTEREST RATE AND BOND OPTIONS

Almost all the interest rate and bond options are options on the respective futures contracts. Options are traded on all the futures contracts discussed in Chapter 2.

- The Chicago Board of Trade offers contracts on US Treasury bonds and notes.
- The Chicago Mercantile Exchange offers options on Eurodollar futures and other euro-currency interest rate futures.
- Eurex trades contracts on German government bonds (Schatz, Bobl and Bund) and on the Eurozone money market rates (Eonia and Euribor).
- Euronext.liffe trades options on UK government bonds, on short-term interest rates, and on swap rates.

Most of these option contracts are American-style exercise. It is common to list a relatively large number of strikes at the inception of each contract. Options are listed with maturities up to two years or beyond.

Chapter Summary

Option contracts can be exchange traded or over the counter OTC contracts. Currency options trade in OTC markets. In other underlyings, the exchange traded options market is bigger in size, but the OTC market is quite significant.

Successful option contract design hinges on the correct choice of the underlying and the settlement mechanism, as well as other details like the contract size, contract months, position limits, and margins. Options can be physically settled, cash settled or settled, by delivery of a futures contract on the same underlying.

The open interest and traded value of options is usually defined in terms of the notional value and not in terms of the actual price or option premium. The outstanding open interest is about \$100 trillion or more than twice the world GDP. Financial products (particularly interest rates and equity index options) dominate the global options market today. In India, exchange traded options are confined to equities and equity indices.

Suggestions for Further Reading

The websites of the major derivative exchanges provide detailed information about contract specifications of each futures contract. For example:

http://www.cboe.com/ for the Chicago Board Options Exchange (CBOE) in the United States http://www.iseoptions.com/ for the International Securities Exchange in the United States http://www.cme.com/ for the Chicago Mercantile Exchange (CME) in the United States http://www.nseindia.com/ for the National Stock Exchange in India

Problem and Questions

- 1 Many exchanges permit a trader to trade or to request for an option with a strike that is not listed but is within the range of strikes that are listed. The trade or quote request must be for a large quantity (typically 50 times the standard contract size). How do the parties that trade this "flexible option contract" gain by listing it on the exchange instead of doing an OTC trade between themselves?
- 2. How do the flexible options described in the previous problem benefit or hurt small traders? How do they benefit the exchange?
- 3. Would you agree with the following statement: "The principal reason why exchanges design option contracts on the future instead of directly on the underlying is to attract the hedging volume of the option traders into their futures contracts instead of into the spot market for the underlying"?
- 4. Would you agree with the following statement: "American options are more attractive when there is physical settlement than when there is cash settlement"?

Chapter **Eight**

Risk Neutral Valuation

In a world of certainty, valuation is simply an adjustment for the time value of money—all cash flows have to be discounted at the prevailing rate of interest. This is an uninteresting world in which there would be no need for derivative markets or for risk management.

When we bring uncertainty into the picture, we have to adjust for risk. Most people do not like to take risk and would prefer low-risk investments. The valuation of uncertain cash flows must reflect this reality of human behaviour and somehow reduce the value of risky cash flows. There are many ways of doing this.

The method that is most useful for derivative valuation is that of risk-neutral probabilities. This chapter discusses this method and shows how it is related to other better known methods. It also shows how this method can be justified using utility theory or by arbitrage arguments.

8.1 RISK ADJUSTED DISCOUNT RATES AND CERTAINTY EQUIVALENTS

The most common method for adjusting for risk in corporate finance and in investment analysis is by adjusting the discount rate. The higher the risk, the higher is the return that investors demand and the higher, therefore, is the Risk Adjusted Discount Rate (RADR) at which we perform the discounting. What we discount at the RADR is the expected value of the cash flows. For example, if there is a 40 percent chance that the cash flow in the first year would be 55 and a 60 percent chance that the cash flow would be 155, then the average or expected value is 115 ($0.40 \times 55 + 0.60 \times 155 = 22 + 93 = 115$). It is this expected cash flow of 115 that is discounted for one year at the RADR. Suppose, the risk free rate is 10 percent and the riskiness of the cash flows requires a 5 percent risk premium, the RADR is 15 percent

and the value of the uncertain cash flow is $\frac{115}{1.15} = 100$.

The concept of RADR immediately leads to the question as to why we are adjusting the discount rate and not the cash flows. The answer is that it is possible to do things that way too. Most finance text books have a small section on Certainty Equivalents (CE) which discusses this approach. The CE approach to the above example would start with the following question. Suppose there were a game that would give you 55 immediately with a 40 percent chance and 155 immediately with a 60 percent chance, what would you be willing to pay to be allowed to play this game? Since the cash flows are immediate, there is no scope for discounting here; we have to adjust the cash flows. Most people are risk averse and would say that they would pay less than the expected value of 115 to play this game. They may perhaps be prepared to pay 110. This value of 110 is the CE of the uncertain cash flow. Now suppose this uncertain cash flow (of 55 with 40 percent chance and 155 with 60 percent chance) will take place only after one year. The CE approach would say that this uncertain future cash flow can be valued by discounting its CE at the risk-free rate. In this case, 110 discounted at the risk free rate of 10 percent gives 100 which is the same as what was calculated in the RADR approach.

8.2 | Derivatives and Risk Management

The correct value is achieved by finding the correct RADR and discounting the expected cash flows at this RADR, or by finding the correct CE and discounting this CE at the risk-free rate. The approach that is chosen depends on whether it is easier to estimate the RADR or the CE. In most of elementary finance, it is easier to estimate the RADR by using the Capital Asset Pricing Model or another model of what the return on a risky security should be. But the CE method correctly implemented is equally correct, and there are situations where CE is superior to RADR.

8.2 **RISK-NEUTRAL PROBABILITIES**

While valuing derivatives, however, it is convenient to adjust for risk in a third way. Instead of adjusting the discount rate as in RADR or the cash flows as in CE, probabilities shall be adjusted. This can also be understood as a re-interpretation of the CE. In the previous example, a cash flow of 55 with a 40 percent chance and 155 with a 60 percent chance had a CE of 110 as opposed to the expected value of 115. The re-interpretation is that risk-averse people act as if they attach a higher probability to the bad outcome (cash flow of 55) than the true probability of 40 percent and correspondingly attach a lower probability to the good outcome than its true probability of 60 percent. Suppose we change these two probabilities to 45 percent and 55 percent. The expected value using these altered probabilities is the CE of 110 ($55 \times 0.45 + 155 \times 0.55 = 24.75 + 85.25 = 110$). These altered probabilities are known as risk-neutral probabilities. The CE of 110 can be then interpreted as the expected value using risk-neutral probabilities. The CE of 110 can be then interpreted as the expected value using risk-neutral probabilities. The value of a cash flow of one year hence of 55 with a 40 percent chance and 155 with a 60 percent chance can be valued by first changing the probabilities to make it 55 with a 45 percent chance and 155 with a 55 percent chance and then discounting the resulting expected values at the risk-free rate of 10

percent: $\left(\frac{55 \times 0.45 + 155 \times 0.55}{1.10} = \frac{24.75 + 85.25}{1.10} = \frac{110}{1.10} = 100\right)$. Thus, once the shift has been made to

risk-neutral probabilities, we can ignore risk completely and simply discount expected values at the risk free-rate. Risk adjustment is entirely subsumed in the process of distorting probabilities. This technique is known as risk-neutral valuation.

In the case of derivatives, the risk-neutral valuation is particularly convenient. As the word derivative suggests, the value of the derivative is derived from the value of an underlying asset. The derivative is valued, not in isolation but in relation to the known market price of the underlying. In this case, the known value of the underlying usually contains enough information to allow us to find the risk-neutral probabilities. Once we have found the risk-neutral probabilities for a specific underlying, all kinds of derivatives on that underlying can be valued simply by computing expected values and discounting them. The RADR approach, on the other hand, would first of all require a different RADR for each kind of derivative on the same underlying. Worse, for derivatives that depend in a non-linear manner on the underlying, the RADR would not remain the same as the price of the underlying changes. This means that the RADR approach would be computationally impractical to implement.

In this book, it is taken for granted that the principle of risk-neutral valuation is correct and a rigorous proof of this principle is therefore, not offered. The reader is asked to refer to books of a higher level for such proof. However, optional sections at the end of this chapter explain the principle of risk-neutral valuation using two different approaches.

One approach uses the theory of expected utility, according to which risk aversion arises from the principle of diminishing marginal utility. Investment in a risky security may produce a high return in

Risk Neutral Valuation | 8.3

some states of the world, but diminishing marginal utility causes the increase in utility to be less than proportionate to the increase in wealth. Also, when the risky investment produces negative returns, the loss of utility is much higher (again as a result of diminishing marginal utility operating in reverse). Thus, people who are simply trying to maximize their average (or expected) utility display risk aversion. It can be demonstrated that risk aversion can be captured by distorting probabilities. The risk-neutral probabilities of bad states of the world must be made higher than the actual probabilities and the riskneutral probabilities of good states of the world must be made lower than the actual probabilities. It is possible to derive an exact formula for adjustment to the probabilities on the basis of the marginal utilities in different states of the world. This approach is discussed further in section 8.4 that deal with the expected utility derivation of risk-neutral probabilities.

A second approach to risk-neutral valuation uses arbitrage arguments. An arbitrage opportunity is a method of turning nothing into something. More precisely, an arbitrage opportunity is a security or portfolio that has:

- 1. zero or negative cost
- 2. a positive probability of making money (producing positive cash flows in future)
- 3. no probability of losing money (producing negative cash flows in future)

Clearly, an arbitrage opportunity is inconsistent with any rational or efficient market. This is because everyone would buy an arbitrage portfolio at zero cost and make money out of nothing. If everyone tries to buy this portfolio, its price would rise until it becomes positive and the arbitrage opportunity disappears. The least that we expect out of a rational market is that there be no arbitrage opportunities.

Risk-neutral probabilities are intimately related to the absence of arbitrage. The proof of this relationship involves very deep and important mathematical theorems that we shall not get into in this book, but the relationship itself is very simple. A market is free of arbitrage opportunities, if and only if there is a set of risk-neutral probabilities under which all assets can be valued using risk-neutral valuation.

This result only tells us that a set of risk-neutral probabilities exists. The important task of finding these probabilities remains. Fortunately, for most of the derivatives considered in this book, there is a simple answer to this problem. The risk-neutral probabilities are uniquely determined by the requirement that the underlying asset earns the risk free rate of return under the risk-neutral probabilities. In other words, while the expected return undergoes a change, the variance and other characteristics of the distribution are unchanged, while going to risk-neutral probabilities.

The optional sections 8.5, 8.6, and 8.7 of this chapter provide further material on the 'no arbitrage' approach to risk-neutral probabilities.

8.3 COST OF CARRY MODEL REVISITED

The cost of carry model for futures was developed in Chapter 3, using arbitrage considerations. Many people find it puzzling at first sight that the cost of carry model uses the risk-free rate in the relationship between spot and futures prices. Intuitively, they expect to see the expected return on the asset in that relationship.

Use of risk-neutral probabilities clarifies this very quickly. Under risk-neutral valuation the current spot price of the underlying is equal to the expected spot price at maturity, discounted at the risk-free rate:

8.4 | Derivatives and Risk Management

 $S_0 = e^{-rT} E^*[S_T]$. The forward price must be simply equal to the expected spot price at maturity—the forward price is also paid only at maturity, no discounting is required here¹:

 $F = E^*[S_T]$. Combining the two, $F = e^{rT}S_0$ is derived, which is the cost of carry model. The key point is that the forward price equals the expected spot price at maturity only under the risk-neutral probabilities, and not under the actual probabilities. This is why the discount rate is the risk-free rate and not the expected return on the underlying.

In Chapter 3, the cost of carry model was derived using an arbitrage argument. Here, the same result follows more easily, using risk-neutral probabilities. But as explained above, this is not a coincidence. The absence of arbitrage is equivalent to the existence of risk-neutral probabilities and therefore, at the end of the day, the two approaches are one and the same. Computationally, of course, risk-neutral valuation is much easier because arbitrage portfolios do not have to be created each time.

The option valuation models to be discussed later in this book were first derived over 30 years ago using complex arbitrage arguments. The use of risk-neutral valuation simplifies the analysis considerably and is the route that this book shall take.

8.4^{*} EXPECTED UTILITY DERIVATION OF RISK-NEUTRAL PROBABILITIES

Utility theory provides the foundation for the modern theory of decision making under uncertainty, not only in finance but also in a major part of microeconomics. Many pricing models like the Capital Asset Pricing Model are based on utility theory as are most models of portfolio choice and hedging behaviour. Utility theory also provides a succinct explanation of the phenomenon of risk aversion.

In this theory, people derive utility from the consumption of goods and services. Savings and investment are tools to shift consumption from one time period to another. Instead of consuming all their income, people save and invest a part of it and use savings to finance consumption at a future point of time.

If we consider consumption at only two points of time—time t = 0 and time t = 1 - a, a simple way of writing the utility function is:

$$u(c_0) + au(c_1)$$
 (8.1)

where c_0 is the consumption at time 0 and c_1 is the consumption at time 1. The fraction *a* (typically less than 1) represents the subjective rate of time preference—most people prefer consumption now to consumption later. Then why do people save at all? The answer is provided by the fact that u(c) is not a linear function.

It exhibits diminishing marginal utility, whereby successive additional units of consumption provide lower and lower levels of additional utility. Therefore during periods in which income is high, people would prefer to save a part of the income and consume it later, when income may be lower and the marginal utility of consumption, therefore, higher.

$$S_T - F$$
.

The zero initial price should equal the discounted risk-neutral expected payoff at maturity

$$0 = \mathrm{e}^{-rt} \left(E^* \left[S_T \right] - F \right)$$

This implies that $F = E^*[S_T]$ as stated.

⁶ Optional section. May be omitted without loss of continuity.

^{1.} Another way of seeing this is to observe that the price of the forward contract is zero at inception while the payoff at maturity is

Risk Neutral Valuation | 8.5

Diminishing marginal utility also explains risk aversion. It is clear that 10 units of consumption produces much less than twice the utility of 5 units. Because of diminishing marginal utility, the 6th unit has lower utility than any of the first 5 and the 7th, 8th, 9th and 10th units add even less. A 50 percent chance of having 10 units of consumption is therefore, worth less than having 5 units of consumption with certainty. Diminishing marginal utility means that people dislike fluctuations in consumption. They avoid fluctuations over time by savings, and they avoid fluctuations due to uncertainty by not getting into risky investments. They accept risky investments only if the return is large enough to overcome the effect of diminishing marginal utility.

Thus we postulate the investor as choosing savings and investment to maximize the expected utility:

$$E[u(c_0) + au(c_1)] = u(c_0) + aE[u(c_1)]$$
(8.2)

Suppose the investor has performed this maximization and has arrived at the optimum consumption levels c_0^* and c_1^* and the associated portfolio choices. Consider an asset which produces an uncertain cash flow x at time 1 and has a value (price) at time 0 of v. If the investor changes his portfolio by buying a small quantity h of this asset, his consumption at time 0 declines by hv and his consumption at time 1 rises by the uncertain amount hx. His utility changes by

$$-hv u'(\mathbf{c}_0^*) + ah E \left[u'(\mathbf{c}_1^*) x \right]$$
(8.3)

Where u' is the marginal utility of consumption.

This change in utility must be zero at the optimum portfolio choice because, if it were positive, the utility could be increased by buying a small amount of this asset and if it were negative the utility could be increased by selling a small amount of this asset.

For the change in utility to be zero, we must have:

$$v = E\left[a \frac{u'(c_1^*)}{u'(c_0^*)}x\right] = E\left[gx\right] \text{ where } g = a \frac{u'(c_1^*)}{u'(c_0^*)}$$
(8.4)

Now consider a discrete set of possibilities (states of the world) so that

$$v = E[gx] = \sum_{i=1}^{N} p_i g_i x_i$$
(8.5)

where the *i*'th state of the world has probability p_i and in this state, g takes the value g_i and x takes the value x_i .

Now we define the risk-neutral probabilities as follows:

$$p_i^* \equiv p_i g_i (1+r) \tag{8.6}$$

where *r* is the risk-free rate.

We obtain the risk-neutral valuation formula:

$$v = \sum_{i=1}^{N} p_i g_i x_i = \frac{1}{1+r} \sum_{i=1}^{N} (1+r) p_i g_i x_i = \frac{1}{1+r} \sum_{i=1}^{N} p_i^* x_i = \frac{1}{1+r} E^*[x]$$
(8.7)

We use E^* to denote the risk-neutral expectation (expectation using risk-neutral probabilities). We have just shown that the value of any asset is obtained by discounting the risk-neutral expectation at the risk-free rate.

8.6 | Derivatives and Risk Management

It may be noted that in bad states of the world, where consumption is low, the marginal utility of consumption will be high, g_i will be higher than 1 and the risk-neutral probabilities p_i^* will be higher than p_i . Thus the risk adjustment consists of over weighting the probabilities of bad outcomes and under weighting the probabilities of good outcomes.

8.5^{*} NO ARBITRAGE AND RISK-NEUTRAL PROBABILITIES—SINGLE PERIOD

In this section, no theories are presumed about how investors make decisions under uncertainty. The only assumption made is the absence of arbitrage which requires that investors prefer more money to less, when there is no risk involved.

An arbitrage opportunity is a security or portfolio that has:

- 1. zero or negative cost
- 2. a positive probability of making money (producing positive cash flows in future)
- 3. no probability of losing money (producing negative cash flows in future)

In other words, an arbitrage portfolio is a method of turning nothing into something.

When an arbitrage portfolio exists, everyone wants to buy it and as the price of the portfolio rises to positive levels, the arbitrage opportunity disappears. This means that a market with an arbitrage opportunity can never be in equilibrium. No arbitrage is thus a basic requirement of any sound market.

The fundamental result that we want to describe in this section is that there is no arbitrage, if and only if a set of risk-neutral probabilities exist. While we shall not prove this important fact, we shall provide an intuitive explanation as to why it must hold.

Consider a simple world in which S_1 , the stock price at time t = 1, can take only two values H (high) and L (low). Without loss of generality, assume that S_0 , the stock price at time t = 0, is 1. There is a bond that is worth $B_0 = 1$ at time t = 0 which will be worth $B_1 = (1+r)$ at time t = 1. This is a single period model where the risk-free rate is equal to r.

There are two 'no arbitrage' conditions:

- 1. H > B. Otherwise shorting the stock and putting the proceeds into the bond is an arbitrage opportunity. It has zero initial cost. The future cash flow from this is always positive (it is either B H or B L). It is thus a risk-free way of turning nothing into something.
- 2. L < B. Otherwise shorting the bond and putting the proceeds into the stock is an arbitrage opportunity. It has zero initial cost. Here, future cash flow is always positive (it is either H B or L B). It is thus, a risk-free way of turning nothing into something.

These two 'no arbitrage' conditions can be written together as L < B < H.

We now want to find risk-neutral probabilities such that the current stock price is the risk-neutral expectation discounted at the risk-free rate. Let us use the notation E^* to denote the expectation computed using risk-neutral probabilities. We then have

$$S_0 = E^* \left[S_1 / (1+r) \right] = E^* \left[S_1 / B_1 \right]$$
(8.8)

Using the fact that $S_0 = 1$, we obtain $E^*[S_1] = B_1$. The risk-neutral expected stock price at time 1 is equal to B_1 . Denote the risk-neutral probability by $p^* = P^*(S_1 = H)$. The requirement then is that $E^*[S_1] = p^*H + (1-p^*)L = B_1$.

^{*} Optional section. May be omitted without loss of continuity.
Risk Neutral Valuation | 8.7

The way to achieve the desired expected value is to set

$$p^{*} \equiv P^{*} (S_{1} = H) = \frac{B - L}{H - L}$$

$$1 - p^{*} = P^{*} (S_{1} = L) = \frac{H - B}{H - L}$$
(8.9)

The no arbitrage conditions guarantee that the probabilities exist, that is to say:

$$0 \le p^* \le 1 \text{ and } 0 \le 1 - p^* \le 1$$
 (8.10)

Conversely, if these probabilities exist, then the 'no arbitrage' conditions are satisfied.

Let us consider a numerical example to illustrate this process. Suppose² that H = 1.30, L = 0.90 and 1.06 - 0.90

B = 1.06. The risk-free rate is clearly 6 percent. The risk-neutral probabilities are given by $p^* = \frac{1.06 - 0.90}{1.30 - 0.90}$

 $= \frac{0.16}{0.40} = 0.40 \text{ and } 1 - p^* = 0.60. \text{ Under these risk-neutral probabilities, the expected stock price at time 1 is given by <math>0.40 \times 1.30 + 0.60 \times 0.90 = 1.06$. Discounting 1.06 at the risk free rate of 6 percent gives us the current stock price $S_0 = 1$.

To contrast risk-neutral valuation with the normal way of valuing a stock using risk adjusted discount rates, let p denote the actual probability $p = P(S_1 = H)$. The actual expected return on the stock is then given by $1 + r_s = pH + (1-p)L$. In the presence of risk aversion, this return r_s would be greater than the risk-free rate r. This in turn implies that p is greater than p^* . Thus what risk-neutral valuation does is to push down the probability of high stock prices and thereby, artificially reduce the expected price. It then discounts the artificially reduced expected price at the risk-free rate.

To continue the earlier numerical example, if the actual probability p = 0.55 then the expected stock price at time 1 is $0.55 \times 1.30 + 0.45 \times 0.90 = 1.12$. This implies that the expected return on the stock is 12 percent. Under risk-adjusted discount rates, we would get $S_0 = 1$ by discounting 1.12 at 12 percent; in risk-neutral valuation, we get it by discounting 1.06 at 6 percent. The key difference is that in riskneutral valuation, we reduce the probability of the good outcome (high stock price) from 55 percent to 40 percent. This reduces the expected stock price and allows the use of the risk-free rate while discounting.

We have shown that the stock can be valued by risk-neutral valuation. But we want to value all derivatives on the stock as well. Consider a derivative D. Since it is a derivative, its value at time 1 depends only on the stock price at time 1. Let the value of D at time 1 be D_1^H when the stock price is high and D_1^L when the stock price is low. These two values may be completely arbitrary. We need to show that the time 0 value of D is given by:

$$D_0 = \frac{p^* D_1^H + (1 - p^*) D_1^L}{1 + r}$$
(8.11)

To show this, we construct a replicating portfolio consisting of s stocks and b bonds, such that its time 1 values are the same as those of D:

$$sH + bB_1 = D_1^H \text{ and } sL + bB_1 = D_1^L$$
 (8.12)

^{2.} The values in this example are under our assumption that $S_0 = B_0 = 1$. The reader may find the example more intuitive if we multiply all prices by 100 so that $S_0 = B_0 = 100$ and *H*, *L* and *B* are equal to 130, 90 and 106 respectively. By doing this, the risk-free rate and the risk-neutral probabilities are unchanged.

8.8 | Derivatives and Risk Management

This is a system of two simultaneous linear equations in two unknowns (s and b) and under our assumptions, these equations always have a solution which we can find as follows.

Subtracting the second equation from the first, we find that $s = \frac{D_1^H - D_1^L}{H - L}$ and substituting this into

the first gives us $b = \frac{1}{1+r} \frac{HD_1^L - LD_1^H}{H - L}$.

Absence of arbitrage requires that at time 0 this portfolio has the same value as D.

$$D_0 = s S_0 + b B_0 = s + b = \frac{1}{1+r} \left(\frac{(1+r) - L}{H - L} D_1^H + \frac{H - (1+r)}{H - L} D_1^L \right)$$
(8.13)

We immediately recognize the risk-neutral probabilities p^* and $1 - p^*$ in this expression. Substituting these gives us the risk-neutral valuation formula:

$$D_0 = \frac{1}{1+r} \left(p^* D_1^H + (1-p^*) D_1^L \right)$$
(8.14)

Continuing our numerical example, we now consider a call option on the stock with a strike price of 1.20. When the stock price is high (1.30), this option will be exercised and will have a value $D_1^H = 1.30 - 1.20 = 0.10$. On the other hand, when the stock price is low (0.90), this option will not be exercised and will have a value $D_1^L = 0$. The replicating portfolio consisting of *s* stocks and *b* bonds is given by:

$$s = \frac{0.10 - 0}{1.30 - 0.90} = \frac{0.10}{0.40} = 0.25$$
 and

$$b = \frac{1}{1.06} \frac{1.30 \times 0 - 0.90 \times 0.10}{1.30 - 0.90} = \frac{1}{1.06} \frac{-0.90}{0.40} = -0.21226$$

The value at time 0 of this replicating portfolio is given by 0.25 - 0.21226 = 0.03774. This is the value of the derivative at time 0. Using risk-neutral valuation, the same value is obtained more easily as

$$D_0 = \frac{1}{1.06} (0.40 \times 0.10 + 0.60 \times 0) = \frac{0.04}{1.06} = 0.03774.$$

8.6^{*} NO ARBITRAGE AND RISK-NEUTRAL PROBABILITIES—MULTIPLE PERIODS

The multiple period extension to the single period model of the previous section is as follows.

The bond price at time *i* is simply $B_0(1+r)^i$. This is a straightforward extension of the single period model, where the risk-free rate is *r* per period. Clearly, the bond under consideration is a zero coupon bond.

^{*} Optional section. May be omitted without loss of continuity.

Risk Neutral Valuation | 8.9



The stock price evolves according to the following tree:

Presently in time 0, consider what could happen in time 1 when the model is the same as the one in the single period model—the stock price could move up by a factor H or down by a factor L. If time is moved to 1, the result would be one of two states, depending on whether the stock moved to the high or to the low price. From each of these states, the behaviour over the next time period is the same as in the single period model—it could move up by a factor H or down by a factor L. Thus the multiple period model is a sequence of single period models and that is what would enables the extension of single period results to the multiple periods case.

At each node in the tree, therefore, the results of the single price model can be applied to conclude that the the risk-neutral probabilities exist if and only if the 'no arbitrage' condition, L < 1 + r < H, is satisfied. The risk-neutral probability of the stock moving up is $\frac{(1+r)-L}{H-L}$ while the probability of it moving down is $\frac{H-(1+r)}{H-L}$. It is obvious that under these probabilities, the stock is priced by risk-neutral discounting as in the single period area.

neutral discounting, as in the single period case.

For example, consider the same numerical values as in the previous section:

H = 1.30, L = 0.90, r = 6 percent. The bond prices are then given by 1.00, 1.06, 1.1236, and 1.1910 at time 0, 1, 2, and 3 respectively. The behaviour of stock prices is as follows:



8.10 | Derivatives and Risk Management

To see that all derivatives are also priced by risk-neutral valuation, a replicating portfolio should be found and valued. The simple method of the previous section cannot be used because at time t = 3, the stock can take four different values $(S_0H^3, S_0H^2L, S_0HL^2, \text{ and } S_0L^3)$ and the derivative could also have four different values, denoted as D_3^{HHH} , D_3^{HLL} and D_3^{HHL} , D_3^{LLL} respectively. These four values cannot be matched using only a portfolio of stocks and bonds because four equations with only two unknowns are then obtained.

For example, if a call option maturing at time 3 with a strike price of 1.10 is considered, the value of this option at time 3 corresponding to the four different stock prices of 2.197, 1.521, 1.053, and 0.729 would be 2.197 - 1.10 = 1.097, 1.521 - 1.10 = 0.421, 0, and 0. In this case, the derivative takes only three values³: 1.097, 0.421, and 0. Using only two assets, stocks and bonds, it is not possible to match these three values.

The way to get around this problem is dynamic replication. A static portfolio of stocks and bonds does not exist, but this portfolio is allowed to change with time. However, the portfolio changes made will be self financing—whenever the portfolio is changed, the value of the new portfolio is exactly the same as the value of the old one. This means that after money has been set aside at the beginning to form the starting portfolio, there will be no further inflows or outflows until the point of maturity where the value of the portfolio at that point will be the same as the value of the derivative itself. If this can be done, it can still be argued by arbitrage, that the value of the derivative at time 0 must be the same as the value of the starting portfolio.

We will compute the replicating portfolios and the risk-neutral valuations simultaneously, starting from time t = 2 and working backwards. Suppose we are at time t = 2, and the stock price is $S_0 H^2$. We know that at time t = 3, there are now only two possible values for the stock $S_0 H^3$ and $S_0 H^2 L$ and there are only two possible values for the derivative D_3^{HHH} and D_3^{HHL} . From the single period model, we know that there is a portfolio of stocks and bonds that replicates the two possible derivative values. We also

know that the cost of this portfolio at time t = 2 is the risk neutral value $D_2^{HH} \equiv \frac{p^* D_3^{HHH} + (1 - p^*) D_3^{HHL}}{1 + r}$



Imagine that the time is t = 2 but the stock price is $S_0 HL$. From this point, we know that at time t = 3, there are two possible derivative values: D_3^{HHL} and D_3^{HLL} . We can construct another replicating portfolio,

^{3.} If the strike price had been 1.05 instead of 1.10, there would have been four values: 1.147, 0.471, 0.003, and 0.

Risk Neutral Valuation | 8.11

whose value at time t = 3 would match these derivative values. Moreover the cost of this portfolio at time

t = 2 would be equal to the risk-neutral value: $D_2^{HL} \equiv \frac{p^* D_3^{HHL} + (1 - p^*) D_3^{HLL}}{1 + r}$.

Similarly, the replicating portfolio at time t = 2 can be found when the stock price is $S_0 L^2$ and its cost D_2^{LL} is also the risk-neutral value,

$$D_2^{LL} \equiv \frac{p^* D_3^{HLL} + (1 - p^*) D_3^{LLL}}{1 + r}$$

For the numerical example (call option with strike of 1.10), the replicating portfolios and derivative values are as follows:



For example, consider the first cell at time 2. The portfolio s = 1 and b = -0.92358 has the value $1 \times 1.69 - 0.92358 \times 1.1236$ because at this point, the stock price is 1.69 and the bond price is 1.1236. At time 3, if the stock price goes up to 2.197, this portfolio will be worth $1 \times 2.197 - 0.92358 \times 1.1910 = 1.097$ which is the call option value for this stock price. If on the other hand at time 3, the stock prices goes down to 1.521, then the portfolio is worth $1 \times 1.521 - 0.92358 \times 1.1910 = 0.421$ which is the call option value for this stock price at time 2 replicates the two possible option values in the next period.

Now imagine that the time is t = 1 and the stock price is $S_0 H$. We know that at the next time period, the stock price can be either $S_0 H^2$ or $S_0 HL$ and that at those points, different replicating portfolios would have to be constructed, whose costs at time t = 2 would be D_2^{HH} and D_2^{HL} depending on whether the stock moves up or down respectively. A replicating portfolio of stocks and bonds has to be found, whose values at time t = 2 would replicate these two values. This portfolio bought at time t = 1 will be different from either of the two portfolios held at time t = 2, but it is certain that at time t = 2, the value of the portfolio that has been bought will be exactly sufficient to buy the portfolio that is to be held at that time. Moreover the cost of this portfolio would be the risk-neutral value:

$$D_1^H = \frac{p^* D_2^{HH} + (1 - p^*) D_2^{HL}}{1 + r}$$

8.12 | Derivatives and Risk Management

Similarly, at time t = 1 when the stock price is $S_0 L$, a replicating portfolio can be found and its cost would be the risk-neutral value $D_1^L = \frac{p^* D_2^{HL} + (1 - p^*) D_2^{LL}}{1 + r}$.

Finally at time t = 0, the replicating portfolio that will at time t = 1 have a value equal to D_1^H or D_1^L can be found, depending on whether the stock moves up or down respectively. The cost of this starting portfolio is again, the risk-neutral value

$$D_0 \equiv \frac{p^* D_1^H + (1 - p^*) D_1^L}{1 + r}$$



The results for the numerical example are as follows:



The dynamic replication strategy is, therefore, to start at time 0 with a portfolio [s = 0.690281 and b = -0.52953] which costs 0.16075. At the next time period, this portfolio will have to be changed to [s = 0.948839 and b = -0.84663] or to [s = 0.4413 and b = -0.31813] depending on whether the stock price has gone up to 1.30 or down to 0.9. This portfolio revision will not, however, produce cash inflow or outflow. At the stock price of 1.30 the old portfolio [s = 0.690281 and b = -0.52953] and the new portfolio [s = 0.948839 and b = -0.84663] will both have the same value 0.336063. Similarly, at the stock price of 0.90 the old portfolio [s = 0.690281 and b = -0.52953] and the new portfolio [s = 0.4413 and b = -0.52953] and the new portfolio [s = 0.4413 and b = -0.31813] will both have the same value 0.336063. Similarly, at the stock price of 0.90 the old portfolio [s = 0.690281 and b = -0.52953] and the new portfolio [s = 0.4413 and b = -0.31813] will both have the same value 0.05995.

Similarly, the portfolios will be revised in subsequent periods without any cash inflow or outflow. Since there are no intermediate cash flows and the dynamic replication strategy does replicate the cash flows of the option at its maturity at time 3, the cost of the portfolio at time 0 must be the value of the option. This value is also equal to the risk-neutral valuation of the option.

Thus in the multiple period model, we have a dynamic replicating portfolio instead of a static replicating portfolio. But the principle of risk-neutral valuation is not affected. The value of any derivative is still given by discounting the risk-neutral expectation at the risk-free rate.

8.7^{*} RISK-NEUTRAL VALUATION IN CONTINUOUS TIME

The continuous time model is essentially a multiple period model in which the number of time periods is very large and each period is very short. As expected, the multiple period results extend to this case as well. The value of any derivative is given by discounting the risk-neutral expectation at the risk free rate.

Moreover, some things do become simpler in continuous time. The distribution of the stock price is log normal. The shift to risk-neutral probabilities consists simply of changing the expected return of the stock to the risk-free rate, leaving the volatility unchanged. These are the simplifications that will allow the derivation of the Black- Scholes formula in Chapter 9.

Chapter Summary

The preferred way of adjusting for risk while valuing derivatives is to shift from actual probabilities to risk-neutral probabilities. Once the shift to risk-neutral probabilities has been made, one can ignore risk completely and simply discount expected values at the risk-free rate. The entire task of risk adjustment is subsumed in the process of distorting the probabilities. This technique is known as risk-neutral valuation.

Risk-neutral valuation can be justified using expected utility according to which, risk aversion arises from the principle of diminishing marginal utility. In bad states of the world, consumption is low and the marginal utility of consumption is high, while the reverse is the case in good states of the world. Risk-neutral valuation adjusts for this by making the risk-neutral probabilities of bad states (low returns) higher than the actual probabilities and the risk-neutral probabilities.

Risk-neutral probabilities are intimately related to the absence of arbitrage. A market is free of arbitrage opportunities if and only if, there is a set of risk-neutral probabilities under which all assets can be valued using risk-neutral valuation. For most of the derivatives considered in this book, the risk-neutral probabilities are uniquely determined by the requirement that the underlying asset earns the risk-free rate of return under the risk-neutral probabilities. In other words, while the expected return is changed, the variance and other characteristics of the distribution are unchanged while going to risk-neutral probabilities.

Suggestions for Further Reading

This chapter covers the foundation of modern derivative pricing theory. Some of the seminal papers in this field are: **Breeden, D (1980)** "Consumption Risk in Futures Markets", *Journal of Finance*, 35, 503–20.

- **Cox, JC, Ross SA and Rubinstein, M** "Option Pricing: A Simplified Approach", *Journal of Financial Economics*, 7, 229–263.
- Cox, JC, Ingersoll, JE and Ross, SA (1985) "An Intertemporal General Equilibrium Model of Asset Prices", *Econometrica*, 53, 363–84.

^{*} Optional section. May be omitted without loss of continuity.

8.14 | Derivatives and Risk Management

- Harrison, Michael J. and Kreps, David M. (1979) "Martingales and Arbitrage in Multiperiod Securities Markets", Journal of Economic Theory, 20, 1979, 381–408.
- Harrison, Michael J. and Pliska, Stanley R. (1981) "Martingales and Stochastic Integrals in the Theory of Continuous Trading", *Stochastic Process and their Applications*, 11, 1981, 215–260.

A good introduction to some of these ideas in discrete time is:

LeRoy, Stephen F. and Werner, Jan (2000), Principles of Financial Economics, Cambridge University Press.

Problems and Questions

- 1. A manager who is thinking of changing his job thinks that there is 80% chance of getting a bonus of Rs 1 million after 9 month and estimates the certainty equivalent of the bonus to be Rs 0.6 million. The manager would get the bonus only if he stays on the job until the bonus is announced. If the risk free rate of return is 8% p.a. continuously compounded, how much would the new employer have to pay him now to compensate him for the loss of the bonus?
- 2. In the above example, what is the risk adjusted discount rate that is implicitly being used by the manager?
- 3. The CBOT trades binary option on the key interest rate (the Federal Funds taraget) set by the US Federal Reserve. The binary call option pays \$1,000 if the rate announced by the Federal Reserve at a specific meeting is more than the strike. Consider an option maturing at the next Federal Reserve meeting with a strike 25 basis points above the current rate. If the risk neutral probability of an interest rate hike of 25 basis points or more is 0.23, what is the value of the option just before the meeting?
- 4. In the above problem, if the option is trading at \$300 what is the risk neutral probability of an interest rate hike of 25 basis points or more?
- 5. A company has issued a bond with a maturity of one year which pays interest at the rate of 10.5% on maturity. The risk neutral probability that the company would default on the payment is 5%. In case of default, all creditors (including the bond holders) would receive 40% of their claims through the liquidation of the assets of the company. What is the fair value of the bond if the risk free interest rate is 8% continuously compounded?
- 6. A venture capitalist is contemplating investing in a risk start up. She thinks that there is risk neutral probability of 70% that the company would fail and she would lose her entire investment. If the company does succeed, she thinks the company would be worth Rs 100 million at the end of five years. The risk free rate is 5.5% p.a. semi-annually compounded. How much should the venture capitalist pay for a 35% stake in the start up company?
- 7. In the above example, suppose that the venture capitalist wants to buy a put option that will allow her to recover her original investment without interest at the end of five years if the company fails. What is the fair price for her put option?
- 8. Consider the following assessment of a company's prospects at the end of a year.

Scenario	Risk neutral probability	Share price
Demand for the product is very high and there	0.05	200
is very little competition.		
Demand for the product is very high but there	0.30	150
is a lot of competition		
Demand for the product is low but there is	0.20	100
very little competition.		
Demand for the product is low and there is a	0.45	45
lot of competition.		

If the risk free rate is 8% p.a. annually compounded, what is the fair value of the share today?

Chapter **Nine**

The Binomial Option Pricing Model

This chapter develops the binomial option pricing model which is a very versatile numerical method for valuing American and European options. A binomial tree assumes that during a short interval of time, the stock can take only two values — the up move and the down move. However, over a longer period of time after several such moves, the stock can take a large number of distinct values that approximates the observed behaviour of stock prices. This chapter shows how to construct a binomial tree that matches the estimated volatility of the stock and the known risk-free rate. It also shows how the binomial tree can be used to value various kinds of options.

9.1 SINGLE STEP BINOMIAL

Suppose that the stock price is currently 120 and that there are only two possibilities of how it will go next year—it can either go up by 25 percent to 150 or down by 20 percent to 96 (see Figure 9.1). Assume that the risk-free rate is 5 percent. Suppose also, that the probability of the up move to 150 is 70 percent and that the probability of the down move to 96 is 30 percent.



Figure 9.1 Stock price tree: Actual Probabilities

The expected stock price next year is therefore, $150 \times 0.70 + 96 \times 30 = 133.80$. The expected return on the stock is $\frac{133.80}{120} - 1 = 11.5$ percent.

In other words, the current stock price of 120 is obtained by discounting the expected stock price of the coming year of 133.80 by the risk — adjusted discount rate of 11.5 percent.

The principle of risk-neutral valuation described in the previous chapter says that there is a different way to do this. We can find a set of risk-neutral probabilities such that the stock earns only the risk-free rate of 5 percent. In this case, the risk-neutral probability p^* of the up move is $\frac{5}{9}$ and the risk-neutral

probability
$$1 - p^*$$
 of the down move is $\frac{4}{9}$ as shown in Figure 9.2. This is because $\frac{5}{9} \times 25\% + \frac{4}{9} \times -20\% = 5\%$

9.2 | Derivatives and Risk Management



Figure 9.2 Stock price tree: Risk-neutral Probabilities

Under risk-neutral probabilities, the expected stock price in the coming year is $150 \times \frac{5}{9} + 96 \times \frac{4}{9}$ = 126. The expected rate of return on the stock is now $\frac{126}{120} - 1 = 5\%$. In other words, the current stock price of 120 is obtained by discounting the risk-neutral expected stock price next year of 126 by the risk-free rate of 5%.

We also explained in the previous chapter that the principle of risk-neutral valuation can be applied to derivatives on the stock as well. Consider for example a European call option on the stock with an exercise price of 130. The payoffs are shown in Figure 9.3:



Figure 9.3 European Call Option (Exercise Price = 130)

Under risk-neutral valuation, the expected value of the call option next year is $20 \times \frac{5}{9} + 0 \times \frac{4}{9} = 11.11$. Discounting this at the risk-free rate of 5 percent gives the immediate value of the call option as $\frac{11.11}{1.05} = 10.58$ as shown in Figure 9.4.



Figure 9.4 Call Option (Exercise Price = 130) : Risk-neutral Valuation

9.2 OPTION DELTA AND DELTA HEDGING

This tree for the call option leads to the important notion of the option delta which will be discussed at length in later chapters. The option delta is the change in the option value when the stock price changes by 1. In the above example, the stock price can be either 150 or 96 (Figure 9.1). Corresponding to these are the call option values of 20 and 0 respectively (Figure 9.4). This means that a change in the stock price of 150 - 96 = 54 corresponds to a change in the option value of 20 - 0 = 20. The option delta is therefore $\frac{20}{54} = 0.3704$.

Put differently, the call option has the same exposure to the stock price as holding 0.3704 shares. The basic idea of delta hedging is that the option can be replicated by buying 0.3704 shares. At the current price, these 0.3704 shares would cost $0.3704 \times 120 = 44.45$. The call option costs only 10.58. This means that the option is a levered position in the stock. It is equivalent to investing 10.58 of one's own funds, borrowing 44.45 - 10.58 = 33.87 and using the total funds of 44.45 to buy 0.3704 shares. This composite position is known as the replicating portfolio.

Since the borrowing of 33.87 has to be repaid with interest at 5 percent so that the total repayment is $33.87 \times 1.05 = 35.55$, the value of the replicating portfolio in the coming year is shown below in Figure 9.5.



Figure 9.5 Call Option (Exercise Price = 130) Replicating Portfolio: 0.3704 shares financed by borrowing 33.38 and investing 10.58 of own funds

It is seen that regardless of whether there is an up move or a down move in the stock price, the value of the replicating portfolio is the same as the value of the call option.

This also means that the value of the call option of 10.58 that was determined by risk-neutral valuation can also be obtained as the value of a portfolio that replicates the option. This confirms the statement made in the previous chapter that risk-neutral valuation is the same as pricing using 'no arbitrage' arguments. Since the option and the replicating portfolio have the same payoffs in future, they must have the same price now.

9.3 MULTI PERIOD BINOMIAL TREES

The assumption that the stock price can take on only two values one year from today is obviously not very realistic. To make the model more realistic, the year is divided into several sub-periods and it is assumed that within each sub-period, the stock moves only to two possible values. In the single period model, the stock price either rose by 25 percent or it fell by 20 percent. These two numbers have an interesting relationship. If the stock rises by 25 percent and then falls by 20 percent, it returns to its original value:

9.4 | Derivatives and Risk Management

 $100 \xrightarrow{+25\%} 125 \xrightarrow{-20\%} 100$. Similarly, the stock price returns to its original value if it first falls by 20 percent and then rises by 25 percent: $100 \xrightarrow{-20\%} 80 \xrightarrow{+25\%} 100$. Both of these arise from the relationship: $(1 + 0.25) \times (1 - 0.20) = 1.25 \times 0.80 = 1$

The same condition comes into play for multi period binomial models. The stock price at any point, either goes up by a factor $u = 1 + r_{up}$ where $r_{up} > 0$ and u > 1 or, it goes down by a factor $d = 1 + r_{down}$ where $r_{down} < 0$ and d < 1. Under a condition $d = \frac{1}{u}$ or $r_{down} = \frac{1}{1 + r_{up}} - 1$, the one year period is divided

into three sub-periods of four months each. One possible tree is shown in Figure 9.6:



Figure 9.6 Stock price tree with three stages

In this tree, u = 1.13750 and d = 0.87912 so that $r_{up} = 13.75\%$ and $r_{down} = -12.09\%$. These numbers are smaller in magnitude than the 25 percent and -20 percent price moves in the single period example. This is consistent with the notion that the range of variation during a four-month period is less than in a period of one year.

The up move at the end of four months takes the stock price to $120 \times 1.13750 = 136.50$. From here, the stock price can rise over the second four months by either 13.75 percent to 155.27 or fall by 12.09 percent to 120. As already discussed, this down move after an initial up move returns the stock price to its original value of 120.

Similarly an initial down move (-12.09 percent) takes the stock price to 105.49. From there an up move (+13.75 percent) returns the stock to its original value of 120. On the other hand, if the second move is also a down move, the stock price falls to 92.74.

Thus at the end of eight months, there are only three possible stock prices: 155.27, 120 and 92.74. The reason why we have only 3 and not $2 \times 2 = 4$ possibilities is that two of the paths (up move followed by down move) and (down move followed by up move) lead to the same value because *ud* is always equal to *du*.

At the end of one year, there are only four possibilities and not $3 \times 2 = 6$ possibilities. This is because the down move from 155.27 and the up move from 120 are the same (136.50). Similarly, the down move from 120 and the up move from 92.74 are the same (105.49).

In general, if there were n sub periods in a binomial tree, at the end, one would have n + 1 different stock prices possible. If a reasonably large number was chosen (say 30 or more) of sub periods, a wide range of stock prices is possible and the binomial model becomes a fairly realistic description of the actual movement of stock prices.

What is the risk-neutral probability of an up move in this model? Since the risk-free rate is 5% per annum, the rate for four months is $1.05^{1/3} - 1 = 1.64\%$. We see that if we set the probabilities of up move and down move to 0.5313 and 0.4687 respectively, then the expected return is 1.64 percent as desired:

$$0.5313 \times 13.75\% + 0.4687 \times -12.09\% = 1.64\%.$$

There is a simple formula to determine the risk-neutral probability p of an up move and the risk-neutral probability 1-p of a down move for a tree with up move u and down move d, where each sub period is of length Δt . The requirement is that:

$$p u + (1-p)d = (1+R)^{\Delta t}$$

where R is the annually compounded risk-free rate of interest. The solution is easily seen to be:

$$p = \frac{(1+R)^{\Delta t} - d}{u - d}$$

In our case, $p = \frac{1.0164 - 0.8791}{1.1375 - 0.8791} = 0.5313$ and $1 - p = 0.4687$.

We can now use the tree of Figure 9.6 to value the same European call option (exercise price = 130) that we valued using a single period model. We begin with the values of the option at the end of one year and use risk-neutral valuation to obtain the values at the end of eight months (see Figure 9.7):



Figure 9.7 European Call Option (Exercise Price = 130): Risk Neutral Valuation at time t = 8 months. Stock price tree of Figure 9.6 is used

9.6 | Derivatives and Risk Management

At the highest node at time t = 8 months, the value of the call option is seen to be 27.37. This is because from here, the up move leads to an option value of 46.62 and the down move to a value of 6.50. Weighting these by the probabilities of 0.5313 and 0.4687 and discounting for four months at the rate of 1.64 percent gives the value of 27.37. This is exactly the same as in the single period model. The difference is that the stock price at the end of 8 months is not known. There are three possibilities and the option value must be determined in each of them. This gives three possibilities of 27.37, 3.40 and 0.00.

With these values determined, the option at the end of four months can be evaluated. (Figure 9.8):



Figure 9.8European Call Option (Exercise Price = 130): Risk Neutral Valuation at
time t = 4 months Stock price tree of Figure 9.6 is used. Values at t = 8 months
are from Figure 9.7

At the upper node at time t = 4 months, the option is valued at 15.87. This is because from here, the up move leads to an option value of 27.37 and the down move to a value of 3.40. Weighting these by the probabilities of 0.5313 and 0.4687, and discounting for four months at the rate of 1.64 percent gives the value of 15.87. This is similar to what was done in the single period model except that now the future option values of 27.37 and 2.40 are themselves the result of risk-neutral valuation performed in a previous step. Thus the multi period model is simply a repeated application of the single period model.

Similarly, at the lower node at time t = 4 months, the option is valued at 1.78. This is because from here, the up move leads to an option value of 3.40 and the down move to a value of 0.00. Weighting these by the probabilities of 0.5313 and 0.4687 and discounting for four months at the rate of 1.64 percent gives the value of 1.78.



This option is now valued at time 0 (Figure 9.9):

Figure 9.9 European Call Option (Exercise Price = 130): Risk Neutral Valuation at time t = 0. Stock price tree of Figure 9.6 is used. Values at t = 8 months are from Figure 9.8

The option is currently valued at 9.12 as the up move leads to an option value of 15.87 and the down move to a value of 1.78. Weighting these by the probabilities of 0.5313 and 0.4687 and discounting for four months at the rate of 1.64 percent gives the value of 9.12.

9.4 OPTION DELTAS IN THE MULTI PERIOD MODEL

Just as in the single period model, one can compute the option delta in the multi period model. The difference is that the delta varies over time. Starting at time t = 0 when the option value is 9.12, in Figure 9.9, the option value at the end of four months is either 15.87 or 1.78 corresponding to stock prices of 136.50 and 105.49. The option delta is $\frac{15.87 - 1.78}{136.50 - 105.49} = 0.4546$. The replicating portfolio using delta hedging is then to buy 0.4546 stocks for $0.4546 \times 120 = 54.55$ using 9.12 of own funds and borrowing the balance of 45.44. The repayment after four months of this borrowing is 46.18. To see that this does replicate the option values, observe that if the stock rises to 136.50, the replicating portfolio is worth $0.4546 \times 136.50 - 46.18 = 62.05 - 46.18 = 15.87$ which is the correct value of the option. If the stock falls to 105.49, the replicating portfolio is worth $0.4546 \times 105.49 - 46.18 = 47.96 - 46.18 = 1.78$ as desired.

The difference in the multi period model is that this delta does not remain constant. Starting with a delta of 0.4546, it is seen four months later that the delta is different. From Figure 9.8, the new delta is computed as 0.6796 in the upper node and as 0.1247 in the lower node.

$$\frac{27.37 - 3.40}{155.27 - 120} = 0.6796 \text{ and } \frac{3.40 - 0}{120 - 92.74} = 0.1247$$

In the multi period model, dynamic delta hedging and dynamic replication are used. One starts with the current delta and replicating portfolio. As the stock price changes, the delta and the replicating portfolio also change. Starting with a delta of 0.4546, if the stock price rises, more stocks are bought to bring the delta up to 0.6796; if it falls, stock is sold to bring it down to 0.1247. The replicating portfolios also change accordingly. When the delta goes up to 0.6796, it costs $0.6796 \times 136.50 = 92.76$ while the option itself is worth only 15.87. So the balance of 76.89 has to be borrowed now. When the delta goes down to 0.1247, it costs $0.1247 \times 105.49 = 13.15$ while the option itself is worth only 1.78. So the balance of 11.56 has to be borrowed.

9.8 | Derivatives and Risk Management

Dynamic replication is expensive because it requires one to buy stock when the stock price is high and sell it when the stock price is low. This is one of the reasons why the option is more valuable than its intrinsic value.

9.5 MATCHING VOLATILITY AND RISK-FREE RATE

There are four things that are required to determine the binomial tree completely:

- 1 The time interval Δt between two successive nodes of the tree
- 2 The factor u by which the stock price changes in an up move
- 3 The factor d by which the stock price changes in a down move
- 4 The probability p of the up move. The probability of a down move is then given by 1-p.

The choice of Δt is governed by the trade-off between accuracy and computational time. The smaller Δt is, the more the number of nodes in the tree, the higher is the accuracy and the greater the computational time. One can also decide on the number of time intervals (nodes) to use in the tree and then determine Δt by dividing the time to maturity by the number of time intervals. It is quite common for market participants to use a tree with 500 nodes while for many less demanding purposes, a tree with 100 or less nodes may be quite adequate.

It has already been seen that the formula for p is $p = \frac{(1+R)^{\Delta t} - d}{u-d}$ where R is the annually compounded risk-free rate. In much of derivative pricing, it is more convenient to work with the continuously compounded risk-free-rate r. In terms of this, the formula for the risk-neutral probability is $p = \frac{e^{r\Delta t} - d}{u-d}$.

In practice, binomial trees are usually constructed such that the up moves and down moves are reciprocals of each other or $d = \frac{1}{2}$.

Thus there is only parameter left to determine – the up move u. It is learnt from Chapter 7 that the volatility of the stock is an important determinant of the option price, and u is chosen to match this. The optional appendix shows that $u = e^{\sigma \Delta t}$ where σ is the volatility of the stock.

9.6 MORE EXAMPLES OF BINOMIAL TREES

Suppose we want to value a European put option with the characteristics shown in Table 9.1:

Table 9.1: European put option example: Option terms

Parameter	Value
Stock price	100
Exercise price	90
Risk-free Rate (Annually Compounded)	8.00%
Volatility	50.00%
Maturity	1 year

To illustrate this, we decide to value the option using three sub periods, so that Δt is equal to $\frac{1}{3}$ years or four months. The parameters of the tree are, therefore, as shown in Table 9.2:

Parameter	Value	
Δt	$-\frac{1}{3}$	
u	$e^{0.50}\sqrt{\frac{1}{3}} = 1.3347$	
d	$\frac{1}{1.3347} = 0.7493$	
$(1+R)^{\frac{1}{3}}-1$	$1.08^{\frac{1}{3}} - 1 = 2.60\%$	
р	$\frac{1.08^{\frac{1}{3}} - 0.7493}{1.3347 - 0.7493} = 0.4727$	

 Table 9.2
 European put option example: Binomial tree parameters

The stock price tree, therefore, is as shown in Figure 9.10:



Figure 9.10 European put option example: Stock price tree with u = 1.3347

Using the risk-neutral probability of p = 0.4727, we now value the European put option with an exercise price of 90. The computations are shown in Figure 9.11.

It is seen that the value of this put option is 12.01. The option delta can also be computed at this point as $\frac{3.98 - 19.80}{133.47 - 74.93} = -0.2702$. Unlike in the case of a call option, the delta of a put option is negative.

9.10 | Derivatives and Risk Management

Thus the replicating portfolio for this option requires shorting 0.2702 shares. One needs to invest one's own funds equal to the option value of 12.01. In addition, an inflow of 27.02 received from shorting the shares is required to be invested. So an investment of a total of 27.02 + 12.01 = 39.03 at the risk-free rate is needed. Along with an interest of 2.60 percent this will repay 40.05 at the end of four months.

After an up move, the short position will be worth -36.07 so that the portfolio is worth 40.05 - 36.07 = 3.98 which is the value of the put option at this point. At this point, the option delta would be only -0.0992 so some of the shorted shares will have to be bought back to adjust the replicating portfolio accordingly. Similarly, after a down move, the delta would become -0.54 and more shares would have to be short sold to adjust the replicating portfolio.

As in the case of dynamic replication of the call option, here also one ends up buying shares at high prices and selling them at low prices. This is an important reason why the option is worth more than its intrinsic value.

9.7 BINOMIAL MODEL FOR AMERICAN OPTIONS

The binomial model can be used to value American options too. Only a small change is required in the binomial model for this purpose. The stock price tree does not change and the risk-neutral probabilities are also unaffected. What changes are the payoffs that are discounted.

In the case of a European option, it is exercised only at maturity. Prior to that, the option value consists entirely of discounted expected values, of the option's future values. In the case of an American option, there is a possibility of exercising the option at any time. This means that we must also consider whether it is better to exercise the option immediately or to hold on to it in the expectation that it will be worth more in future. Therefore at each node, we must compare the value obtained by immediate exercise with the discounted expected future value of the option.

To illustrate this process, we now value the put option of the previous section, assuming that it is American instead of European. The computations are shown in Figure 9.12.

There is no change in the nodes for time t = 1 year because at expiry date, the American and European options are identical. The difference comes prior to expiry. Consider the bottom most node at time t = 8 months. The computation for this node for the European put option (value = 31.58) has been repeated at this node, but has been struck off and replaced by the new computation 90.00–56.14 = 33.86. The value 31.58 is the discounted expected value of the option if the option is held on to. This needs to be compared with the value obtained by immediate exercise which is 33.86. Since immediate exercise gives the higher value, the option is exercised and the value of 33.86 is entered for this node. This value will be used for computing the discounted expected value at time t = 4 months.



Figure 9.11 European Put Option Example (Exercise Price = 90): Risk Neutral Valuation. The stock price tree of Figure 9.10 is used with the risk-neutral probability of p = 0.4727

9.12 | Derivatives and Risk Management



Figure 9.12 Binomial Tree for American Put Option: Exercise Price = 90. The stockprice tree of Figure 9.10 is used with the risk-neutral probability of p = 0.4727.

On the other two nodes at time t = 8 months, this comparison is not necessary because the option has zero value if exercised (the exercise price is less than the current stock price).

At time t = 4 months, in the upper node, again the option is out of the money (the exercise price is less than the current stock price). At the lower node, the comparison is necessary. The value of immediate exercise has been computed as 90.00-74.93 = 15.07 and struck off as it is lower than the value of holding on to the option (20.97).

Finally, at time t = 0, the option is out of the money and the question of immediate exercise need not be considered. The value of the option (12.61) is obtained as the discounted expected value of holding on to the option.

This value is 0.60 more than the value of 12.01 that was obtained for the European option. This difference can be explained as follows: The American option will prematurely be exercised at the lower-most node at time t = 8 months. The risk-neutral probability of reaching this node is the probability of two successive down moves which is $(1-p)^2 = (1-0.4727)^2 = 0.5273^2 = 0.2780$. When early exercise does take place, the gain is 33.86 - 31.58 = 2.28. This has to be multiplied by its risk-neutral probability of 0.2780 and discounted for eight months. The result is $\frac{2.28 \times 0.2780}{1.08} = 0.60$. This confirms that the

difference between the values of the American and European options is equal to the expected discounted gain from early exercise.

9.8^{*} BINOMIAL TREES AND THE LOG NORMAL DISTRIBUTION

In the next chapter on the Black-Scholes model, it will be argued that a plausible model of stock price behaviour is that the logarithm of the stock price follows the normal distribution. We shall also define the volatility as the standard deviation of the change in the logarithm of the stock price during a year: $\sigma = \text{STDEV} (\ln S_1 - \ln S_0).$

We now show that the binomial distribution is consistent with this model when we have a large number of nodes (Δt is small). In the binomial model, $S_{t+\Delta t}$ is equal to either $S_t u$ or to $S_t d = S_t / u$ with probability p and 1-p respectively. This means that $\ln S_{t+\Delta t}$ is equal to either $\ln S_t + \ln u$ or to $\ln S_t - \ln u$ with probability p and 1-p respectively. Therefore $\ln S_{t+\Delta t}$ is equal to $\ln S_t + x_1 + x_2 + ... + x_n$ where each x_i is equal to $\ln u$ or to $-\ln u$ with probability p and 1-p respectively.

Moreover, the x_i are independent of each other since at each stage the up move and down move have probability p and 1-p respectively regardless of whether the previous move was an up move or a down move.

If *n* is reasonably large (say 30 or more), the central limit theorem can be invoked to show that $x_1 + x_2 + ... + x_n$ follows the normal distribution approximately. Since at time *t*, $\ln S_t$ is known (is a constant) it follows that $\ln S_t + x_1 + x_2 + ... + x_n$ also follows the normal distribution approximately — the constant merely changes the mean of this distribution. Thus $\ln S_t + n\Delta t$ follows the normal distribution approximately if *n* is around 30 or more. Therefore, if Δt is small, so that there are a large number of time intervals (*n*) between S_0 and S_T it can be concluded that $\ln S_T$ is approximately normal or that S_T is approximately log normal. The binomial model is, therefore, an approximation to the Black-Scholes model that we will be discussed in the next chapter.

The choice of $u = e^{\sigma \Delta t}$ is consistent with the definition of volatility (σ) as the standard deviation of the annual change in the logarithm of the stock price $\ln S_1 - \ln S_0$. If it is assumed that there are *n* nodes between time t = 0 and time t = 1 then $\Delta t = \frac{1}{n}$ and $\ln S_1 - \ln S_0$ is equal to $x_1 + x_2 + ... + x_n$ where each

 x_i is equal to $\ln u$ or to $-\ln u$ with probability p and 1-p respectively and the x_i are independent of each other. The variance of $\ln S_1 - \ln S_0$ is therefore the sum of the variances of all the x_i or n times the

^{*} Optional section. It is recommended that this section be skipped on first reading. It srequires a knowledge of statistics including the central limit theorem.

9.14 | Derivatives and Risk Management

variance of any x_i . The mean of x_i which is denoted by m is given by $m = p \ln u + (1 - p) (- \ln u) =$ $(2p-1) \ln u$. The variance of x_i which we denote by v is then given by: $v = p(\ln u)^2 + (1-p)(\ln u)^2 - m^2 = (\ln u)^2 - m^2$. If one denotes the mean of $\ln S_1 - \ln S_0$

by *M*, then M = nm or $m = \frac{M}{n}$. The variance of $\ln S_1 - \ln S_0$ is then given by:

$$nv = n[(\ln u)^2 - m^2] = n\left[(\ln u)^2 - \frac{M^2}{n^2}\right] = n (\ln u)^2 - \frac{M^2}{n}$$

For this variance to be equal to σ^2 one must have $n (\ln u)^2 - \frac{M^2}{n} = \sigma^2$ or $(\ln u)^2 = \frac{\sigma^2}{n} + \frac{M^2}{n^2}$. If n is

relatively large then the second term $\frac{M^2}{n^2}$ becomes negligible and one has the approximate relationship:

$$(\ln u)^2 \approx \frac{\sigma^2}{n}$$
. This implies that $\ln u \approx \sqrt{\frac{\sigma^2}{n}}$ or $u \approx e^{\sqrt{\frac{\sigma^2}{n}}} = e^{\sigma} \sqrt{\frac{1}{n}} = e^{\sigma} \sqrt{\frac{\Delta t}{n}}$ where we use the fact in the last step that by construction $\Delta t = \frac{1}{2}$.

Thus the choice of u is consistent with the definition of volatility as the standard deviation of the annual change in logarithm of the stock price $\ln S_1 - \ln S_0$.

Chapter **Summary**

A binomial tree is built by dividing the time to maturity of the option into a number of sub periods. In each sub period, it is assumed that the stock price can move in only two ways — an up move by a multiplicative factor u or a down move by a multiplicative factor d. The risk-neutral probability of the up move is determined by the requirement that the stock should earn the risk-free rate.

The parameters of the binomial tree depend on the following:

- 1 the volatility of the stocks
- 2 the risk-free interest rate r (continuously compounded) or R (annually compounded)
- 3 the number of sub periods n, chosen to provide the desired accuracy. Large values of n give greater accuracy but require more computational time. In practice, n is usually between 30 and 500.

The parameters of the binomial tree are then given as follows:

Parameter	Value
Δt (length of each sub period)	$\frac{1}{n}$
<i>u</i> (up move)	$e^{\sigma}\sqrt{\Delta t}$
d (down move)	$\frac{1}{u}$
p (risk-neutral	$p = \frac{(1+R)^{\Delta t} - d}{u - d}$
probability of up move)	
	$\operatorname{OR} p = \frac{e^{r\Delta t} - d}{u - d}$

The value of a European option is determined by first determining its value at each possible stock price at maturity. The values of the option at each node is then obtained by moving back one step at a time, as the discounted expected value of the values at the two possible nodes in the next period. This is done until one comes to time t = 0. The value that one gets at this node is the value of the option.

To value an American option, one needs to make one more check at each node. The value of immediate exercise of the option must be compared with the value of holding on to the option and the higher value must be chosen.

The binomial tree also tells one about the option delta which is the change in the option value when the stock price changes by 1. It is also the number of shares that one must buy to replicate the option. Since the delta changes as one moves along the tree, the replicating portfolio also changes. Dynamic replication or dynamic delta hedging is the process of continuously adjusting the replicating portfolio as the stock price changes. In this process one ends up buying shares at high prices and selling them at low prices. This is an important reason why the option is worth more than its intrinsic value.

In the next chapter on the Black-Scholes model, it shall be argued that a plausible model of stock price behaviour is that the logarithm of the stock price follows the normal distribution. Volatility is defined as the standard deviation of change in the logarithm of the stock price during a year. The binomial distribution is consistent with this model when there are a large number of nodes.

Suggestion for Further Reading

Cox, JC, Ross SA and Rubinstein, M "Option Pricing: A Simplified Approach", *Journal of Financial Economics*, 7, 229-263.

Problems and Questions

- 1. The current share price is 85, the risk free rate is 6% continuously compounded and the volatility is 25%. Construct a binomial tree with one node per month and use it to value a two month call option with a strike of 90.
- 2. In the above example, use a tree with two nodes per month to value the same option.
- 3. The current share price is 90, the risk free rate is 15% continuously compounded and the volatility is 40%. Construct a binomial tree with one node per quarter and use it to value a one year European put option with a strike of 80.
- 4. In the above example, how would the value change if the put were American?
- 5. Compute the sensitivity of the option price to changes in the stock price in the first problem by changing the current stock price to 86, building a new tree and valuing the option. This sensitivity is known as delta and is discussed at length in Chapter 12.
- 6. Compute the sensitivity of the option price to changes in the volatility in the third problem by changing the volatility to 41%, building a new tree and valuing the option. This sensitivity is known as vega and is discussed at length in Chapter 12.
- 7 In Problem 5, compute the delta by changing the stock price to 84 instead of to 86. How much does your answer differ from the answer obtained in problem 5? This difference is a measure of the sensitivity of the delta to the stock price and is known as gamma (see Chapter 12)
- 8. The current share price is 100, the risk free rate is 5% continuously compounded and the volatility is 30%. Construct a binomial tree with two nodes per month and use it to value a one month European put option with a strike of 110.
- 9. Use the tree of problem 8 to value a European call with the same strike. Verify put call parity. This verification can be done at any node of the tree and not just for the current option value.

Chapter **Ten**

Objectives

The Black-Scholes Option Pricing Model

The Black-Scholes formula for the price of a European option on a non-dividend paying stock is perhaps the most famous formula in the whole of finance. This formula is developed in the present chapter, and its use is explained.

How the formula can be modified to deal with dividend-paying stocks, stock indices, currencies, commodities, and futures is also explained. The conditions under which the formula can be used for American options are also discussed.

10.1 THE MODEL OF STOCK PRICE BEHAVIOUR

To perform risk-neutral valuation, one needs to compute the risk-neutral expectation; and to do that one needs to know the distribution of stock prices.

The most popular distribution in statistics is the normal or bell-shaped distribution—but this is a bad model of stock prices themselves. This is because stock prices have an unlimited upside while the downside is limited to the share price dropping to zero. Over a five-year horizon, for example, it is not a rare occurence for a Rs 100 share to rise to Rs 1,000 for a return of +900% while the worst loss that can happen is that the share price drops to 0 for a loss of -100%.

At the same time, stock price returns over short periods (say a day) look very much like the normal distribution. It is as likely that a stock will rise by 5% on any given day as it is that it would fall by 5%. How is it that long period, returns behave so differently? The answer is provided by the power of compounding. Consider the +900% return over five years that we discussed earlier. If one calculates the compound annual return, it is about 58%. If instead the stock gave -58% return every year for five years, the cumulative return would be about -99% which may be unlikely but not impossible. The reason for this big difference is that when the returns are positive, each successive year's return is negative, each successive year's return is negative, each successive year's return is negative, each successive year's return is applied on a smaller and smaller base, and the cumulative return is bounded.

Taking annual returns does not completely solve the problem for a stock can rise by more than +100% in a year while it cannot fall by more than -100%. This suggests that one should look at shorter and shorter intervals and consider returns that are compounded monthly or daily, or for even shorter periods to completely remove the asymmetry between positive and negative returns. Mathematically, one should look at continuously compounded or logarithmic returns. As discussed in Chapter 3 (section 3.5), the continuously compounded or logarithmic return is defined as $\ln (1 + R)$ where R is the usual (proportional) return

defined by $\frac{P_t - P_{t-1}}{P_{t-1}}$ where P_t is the price at the end of the period, P_{t-1} is the price at the beginning of the

period (or end of the previous period) and ln denotes the natural logarithm. The logarithmic return can be

expressed more simply as
$$\ln\left(\frac{P_t}{P_{t-1}}\right)$$
 because $\ln\left(1+R\right) = \ln\left(1+\frac{P_t-P_{t-1}}{P_{t-1}}\right) = \ln\left(\frac{P_t}{P_{t-1}}\right)$

10.2 | Derivatives and Risk Management

10.2 THE LOG NORMAL DISTRIBUTION

The model of stock prices that will be adopted in this chapter is that continuously compounded or logarithmic return conforms to the bell-shaped normal distribution.

It shall be assumed that the stock pays no dividends so that the entire return from holding the stock comes through the price appreciation.

Apart from the ability to use the normal distribution, the logarithmic return has another major advantage— the return over a long period can be obtained by simply adding returns over shorter periods. For example, the cumulative return over the next five years is simply the sum of the five future annual returns:

$$\ln\left(\frac{P_5}{P_0}\right) = \ln\left(\frac{P_5}{P_4}\frac{P_4}{P_3}\frac{P_2}{P_2}\frac{P_1}{P_1}\frac{P_2}{P_0}\right) = \ln\left(\frac{P_5}{P_4}\right) + \ln\left(\frac{P_4}{P_3}\right) + \ln\left(\frac{P_3}{P_2}\right) + \ln\left(\frac{P_2}{P_1}\right) + \ln\left(\frac{P_1}{P_0}\right)$$

Since the sum of any number of normal variates is also a normal variate, the assumption can be restated as saying that the distribution of logarithmic returns over any time period (however short or long) is normal.

One can write the normal variate $\ln\left(\frac{P_T}{P_0}\right)$ as equal to $\ln\left(P_T\right) - \ln\left(P_0\right)$ and transpose the term $\ln\left(P_0\right)$ to the other side. One obtains $\ln\left(P_T\right) = \ln\left(P_0\right) + \ln\left(\frac{P_T}{P_0}\right)$.

The term $\ln (P_0)$ is a known constant since the current stock price is known, and only future stock prices are unknown. It is thus found that the logarithm of the stock price at any future date is normally distributed because it is the sum of a constant and a normal distributed return.

Since the logarithm of the future stock price is normally distributed, it is said that the stock price is log normally distributed.

10.3 THE RISK NEUTRAL LOG NORMAL DISTRIBUTION

Not only is the actual distribution of stock prices log normal, the risk-neutral distribution is also log normal. This is because, in continuous time, the only change in the risk-neutral distribution is that the expected return of the stock is changed to the risk-free rate.

This change is trickier than it sounds because of the difference between arithmetic and geometric means. Take an example. If the continuously compounded (logarithmic) return during a year is 20%, then it is known that the annual return is $e^{0.20}-1 = 22.14$ %. Now consider a situation of uncertainty where the expected logarithmic return on the stock is 20%, (it could be more or less). If the logarithmic return could be either 5% or 35% with equal probability, then the expected continuously compounded return is 20%. But the annual return could be either $e^{0.05}-1 = 5.13\%$ or $e^{0.35}-1 = 41.91$ % with equal probability and the expected annual return is 23.52% which is significantly larger than the 22.14% found in the absence of uncertainty. To understand where the difference comes from, it is assumed that in one year the return is 5.13% and in the next year, the return is 41.91%. What is the compound average annual return (which is essentially the geometric mean) agrees with the no-uncertainty annual return of 22.14%, but the arithmetic mean is different from this.

The Black-Scholes Option Pricing Model | 10.3

Risk-neutral valuation is all about the expected stock price. It requires the expected stock price to yield the risk-free rate. This means that the arithmetic mean of the annual return must be set equal to the annual risk-free rate. When one does so, the geometric mean will necessarily be less than the risk-free rate.

Continuing with the earlier example, suppose the risk-free rate is 22.14% (annual) or 20% continuously compounded. If one assumes that the logarithmic (continuously compounded) return could be either 3.88% or 33.88% with equal probability, then the expected logarithmic return is only 18.88% which is lower than the continuously compounded risk-free rate of 20%. But the annual return could be either $e^{0.0388}$ –1=3.96% or $e^{0.3388}$ –1=40.33% with equal probability and the expected annual return is 22.14% which is the annual risk-free rate.

The log normal distribution has an important property (Property I in the Appendix) that makes it easy to make this correction. If the continuously compounded risk-free rate is r and the variance of the logarithmic return of the stock is σ^2 , then the expected logarithmic return of the stock must be set equal

to $r - \frac{1}{2}\sigma^2$ to ensure that the expected annual return of the stock is equal to the annual risk-free rate. The difference between these two can also be stated in terms of the expectations of S_T and $\ln S_T$:

$$E[S_1] = e^r S_0$$
 (10.1)

but E $[\ln S_1] \neq \ln S_0 + r$. Instead,

E
$$[\ln S_1] = \ln S_0 + r - \frac{1}{2} \sigma^2$$
 (10.2)

Equation (10.1) expresses the requirement of risk-neutral valuation that the stock should earn the risk-free rate. The second equation (10.2) shows the adjustment to the logarithmic return required to give effect to this requirement.

For example, if S_0 is 100, r is 10% and σ is 15%, then the expected logarithmic return on the stock is not 10% but 8.875% because $0.1 - \frac{1}{2} \ 0.15^2 = 0.1 - 0.01125 = 0.08875$. We have $E[S_1] = e^{0.1} \times 100 = 110.5171$, $\ln S_0 = 4.60517$ but $E[\ln S_1] \neq 4.60517 + 0.1 = 4.70517$. Instead $E[\ln S_1] = 4.60517 + 0.1 - \frac{1}{2} \ 0.15^2 = 4.60517 + 0.1 - 0.01125 = 4.69392$.

Similar expressions hold for time periods other than a year. The cumulative logarithmic return over T

years is equal to $rT - \frac{1}{2}\sigma^2 T$ and the expected value of $\ln S_T$ is $\ln S_0 + rT - \frac{1}{2}\sigma^2 T$.

The variance of $\ln S_T$ is much easier to compute—it is simply $\sigma^2 T$. Knowing the mean and the variance completely defines the risk-neutral distribution of the logarithm of the future stock price: $\ln S_T$ is normally distributed with a mean of $\ln S_0 + rT - \frac{1}{2}\sigma^2 T$ and a variance of $\sigma^2 T$. The stock price itself is log normally distributed. By the properties of the log normal distribution, the expectation of S_T is $e^{rT}S_0$ indicating that the stock earns the risk-free rate over this period. To summarize:

$$E[S_t] = e^{rT} S_0$$

$$E[\ln S_t] = \ln S_0 + rT - \frac{1}{2} \sigma^2 T$$

$$Var [\ln S_t] = \sigma^2 T$$
(10.3)

10.4 | Derivatives and Risk Management

10.4 THE BLACK-SCHOLES FORMULA

We are now ready to derive the Black-Scholes formula for the value of a European call option on a stock.

A European call option is the right without an obligation to buy the stock at an exercise price X at a time T. Since the option is European, it can not be exercised before maturity. At time T, the holder of the option will check whether it is worthwhile to exercise the option. If the stock price at maturity S_T is greater than X, it is advantageous to exercise the option. The gain from exercising the option is $S_T - X$. If the stock price at maturity S_T is less than X, it does not make sense to exercise the option. The value of the option is, therefore, zero.

Under risk-neutral valuation, the value of this option at time t = 0, is given by computing the riskneutral expectation of the option value and discounting it at the risk-free rate. We must first find the probability that S_T is greater than X, and then find the conditional expected value of $S_T - X$ given that S_T is greater than X.

For the log normal distribution, both these quantities can be computed.

$$P(S_T > X) = P(\ln S_T > \ln X) = P\left(\frac{\ln S_T - M}{\sqrt{V}} > \frac{\ln X - M}{\sqrt{V}}\right)$$
$$= P\left(z > \frac{\ln X - M}{\sqrt{V}}\right) = P\left(z < -\frac{\ln X - M}{\sqrt{V}}\right) = N\left(\frac{M - \ln X}{\sqrt{V}}\right)$$
(10.4)

where *M* is the mean and *V* the variance of the normal variate $\ln S_T$; *z* is therefore a standard normal variate; and *N*(·) denotes the cumulative normal distribution. In the previous section, *M* and *V* have already been determined for the risk-neutral distribution of $\ln S_T$:

$$M = \ln S_0 + rT - \frac{1}{2}\sigma^2 T \text{ and } V = \sigma^2 T$$
 (10.5)

Therefore

$$P(S_T > X) = N\left(\frac{\ln S_0 + rT - \frac{1}{2}\sigma^2 T - \ln X}{\sigma\sqrt{V}}\right) = N(d_2)$$

$$d_2 = \frac{\ln \frac{S_0}{X} + rT - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{V}}$$
(10.6)

where

It is necessary to compute the risk-neutral expectation of $S_T - X$ given that S_T is greater than X. This can be broken into two parts:

$$E[S_T - X | S_T > X] = E[S_T | S_T > X] - E[X | S_T > X] = E[S_T | S_T > X] - X.$$
(10.7)

The following property of a log normal distribution is now used (Property II in the Appendix): If ln Y is normally distributed with mean M and variance V, then the conditional expectation E(Y|Y>X)

is given by $E(Y) \frac{N(d_1)}{N(d_2)}$ where $d_2 = \frac{M - \ln X}{\sqrt{V}}$ and $d_1 = d_2 + \sqrt{V}$.

The Black-Scholes Option Pricing Model | 10.5

Plugging in the values of M and V, it is obtained that

$$d_{1} = \frac{\ln \frac{S_{0}}{X} + rT + \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}} \text{ and } d_{2} = \frac{\ln \frac{S_{0}}{X} + rT - \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}}$$
(10.8)

Under the risk-neutral distribution, $E(S_T) = e^{rT} S_0$ and the conditional expectation can be computed as:

$$E[S_T - X | S_T > X] = S_0 e^{rT} \frac{N(d_1)}{N(d_2)} - X$$
(10.9)

To get the risk-neutral value of the option, all that is required is to multiply the above expression by the risk-neutral probability that S_T is greater than X and discount it by the risk-free rate. This gives one the value C of a call option:

$$C = e^{-rT} N(d_2) \left(S_0 e^{rT} \frac{N(d_1)}{N(d_2)} - X \right) = S_0 N(d_1) - X e^{-rT} N(d_2)$$
(10.10)

Equations (10.10) and (10.8) gives one the famous Black-Scholes formula for the value of a European call option on a non-dividend paying stock:

$$C = S_0 N(d_1) - X e^{-rT} N(d_2)$$

$$d_1 = \frac{\ln \frac{S_0}{X} + rT + \frac{1}{2} \sigma^2 T}{\sigma \sqrt{T}}$$

$$d_2 = \frac{\ln \frac{S_0}{X} + rT - \frac{1}{2} \sigma^2 T}{\sigma \sqrt{T}} = d_1 - \sigma \sqrt{T}$$
(10.11)

where S_0 is the current stock price, X is the exercise price of the call option, T is the maturity of the option, σ is the volatility of the stock (standard deviation of the logarithmic return) and r is the continuously compounded risk-free rate.

For example, if the current stock price is 100, the volatility is 23% and the risk-free rate is 8%, then the value of a three-month call option with an exercise price of 105 is 3.34 as shown below :

$$d_{1} = \frac{\ln \frac{100}{105} + \frac{0.08}{4} + \frac{1}{2} \cdot 0.23^{2} \cdot \frac{1}{4}}{0.23 \sqrt{1/4}} = \frac{-0.048790 + 0.02 + 0.0066125}{0.115} = -0.192849$$
$$d_{2} = d_{1} - 0.23 \sqrt{\frac{1}{4}} = -0.192849 - 0.115 = -0.307849$$
$$N(d_{1}) = 0.423539 \text{ and } N(d_{2}) = 0.379099 \text{ while } e^{-0.08/4} = 0.980199$$
$$C = 100 \times 0.423539 - 105 \times 0.980199 \times 0.379099 = 3.34$$

Using put-call parity, one can then compute the value of a European put option as follows:

$$P = X e^{-rT} N (-d_2) - S_0 N (-d_1)$$

10.6 | Derivatives and Risk Management

$$d_{1} = \frac{\ln \frac{S_{0}}{X} + rT + \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}}$$
$$d_{2} = \frac{\ln \frac{S_{0}}{X} + rT - \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}} = d_{1} - \sigma\sqrt{T}$$
(10.12)

Continuing the earlier example, if the current stock price is 100, the volatility is 23% and the risk-free rate is 8%, then the value of a three-month put option with an exercise price of 105 is 6.26 as shown below:

$$d_1 = -0.192849$$
 and $d_2 = -0.307849$ as for the call option
 $N(-d_1) = 0.576461$ and $N(-d_2) = 0.620901$ while $e^{-0.08/4} = 0.980199$
 $P = 105 \times 0.980199 \times 0.620901 - 100 \times 0.576461 = 6.26$

10.5 BLACK-SCHOLES MODEL WITH DIVIDENDS

Dividends are normally paid once a year or once a quarter, but many assets other than stocks can be regarded as paying a dividend continuously.

We now consider a stock or other asset that pays dividends continuously to provide a constant dividend yield q. Over a short time period Δt , the stock pays a dividend of $qS_t \Delta t$. In this case, the risk-neutral valuation changes because while the stock must still yield a total return equal to the risk-free rate r, only a portion r - q of this comes through price appreciation, while the remaining q comes through dividends.

The expectations of the stock price and log stock price then become E $(S_T) = e^{(r-q)T}S_0$ and E $[\ln S_T]$

= ln S_0 + (r - q) T - $\frac{1}{2}\sigma^2$ T. These can be rewritten a little differently as:

$$\mathcal{E}\left(S_{T}\right) = e^{rT}\left(S_{0}e^{-qT}\right) \tag{10.13}$$

and

$$E\left[\ln S_T\right] = (\ln S_0 - qT) + rT - \frac{1}{2}\sigma^2 T = \ln\left(S_0 e^{-qT}\right) + rT - \frac{1}{2}\sigma^2 T$$
(10.14)

The only change that takes place is that S_0 is replaced by $S_0 e^{-qT}$. Instead of performing all the computations of the previous section once again, one can simply change S_0 to $S_0 e^{-qT}$ everywhere in the final formula.

There is a simple economic interpretation of this change. Imagine the stock being broken up into two assets—the first asset earns all the dividends but none of the capital appreciation and the second asset earns all the capital appreciation but none of the dividends. We consider the call option to be an option on this second asset, that by definition pays no dividend. Since the price of this asset at maturity is the same as the price of the stock, this approach gives the correct answer and the only change is to replace the price (S_0) of the stock by the price (denoted S_0) of the asset that earns only the capital appreciation and not the dividends.

Under risk-neutral valuation, \hat{S}_0 is simply the expected future stock price discounted at the risk free rate. Because of the relationship $E(S_T) = e^{rT} (S_0 e^{-qT})$ that we derived in equation (10.13), the equality obtained is:

$$\dot{S}_0 = S_0 e^{-qT} \tag{10.15}$$

The Black-Scholes Option Pricing Model | 10.7

Changing S_0 to S_0e^{-qT} everywhere in the Black-Scholes formulas—equations (10.11) and (10.12)—gives us the generalized Black-Scholes formula for an asset that pays a dividend continuously at a constant rate q:

$$C = S_0 e^{-qT} N(d_1) - X e^{-rT} N(d_2)$$

$$P = X e^{-rT} N(-d_2) - S_0 e^{-qT} N(-d_1)$$

$$d_1 = \frac{\ln \frac{S_0}{X} + (r-q)T + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln \frac{S_0}{X} + (r-q)T - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$
(10.16)

Continuing the example of the previous section, if the current stock price is 100, the volatility is 23%, the risk-free rate is 8% and the dividend yield is 3%, then the value of a three-month call option with an exercise price of 105 is 3.03 while the put with the same exercise price is worth 6.70 as shown below :

$$d_{1} = \frac{\ln \frac{100}{105} + \frac{0.08 - 0.03}{4} + \frac{1}{2} 0.23^{2} \frac{1}{4}}{0.23\sqrt{1/4}} = -0.258067$$

$$d_{2} = d_{1} - 0.23\sqrt{\frac{1}{4}} = -0.258067 - 0.115000 = -0.373067$$

$$N(d_{1}) = 0.398178 N(d_{2}) = 0.354549 N(-d_{1}) = 0.601822 N(-d_{2}) = 0.645451$$

$$e^{-0.08/4} = 0.980199 e^{-0.03/4} = 0.992528$$

$$C = 100 \times 0.992528 \times 0.398178 - 105 \times 0.980199 \times 0.354549 = 3.03$$

$$P = 105 \times 0.980199 \times 0.645451 - 100 \times 0.992528 \times 0.601822 = 6.70$$

The Black-Scholes model with continuous dividend yields is very useful for options on stock indices, currencies, commodities, and futures which can all be regarded as paying a continuous dividend. These applications are described later in this chapter.

It might appear at first glance that the model would provide only a crude approximation for stocks which pay lumpy dividends annually or quarterly. Surprisingly, the model can still be applied if we interpret q with some care to ensure that equation (10.15) is satisfied. In other words, if we interpret q as the average dividend yield during the life of the option so that $\dot{S}_0 = S_0 e^{-qT}$, then the model given by Equation (10.16) is valid.

Another common situation is that only one or two dividends are expected to be paid during the life of the option and the timing and quantum of these dividends are known. This would be the case for stocks that pay a regular quarterly dividend. In this case, we could simply subtract the present value of these known dividends and use

$$\dot{S}_0 = S_0 - \sum_{i=1}^k e^{-rt_i} D_i$$
(10.17)

where there are k dividends of amount D_i at time t_i . Since the dividends are known with certainty, discounting is done at the risk-free rate. We can replace S_0 by this value of S_0 in equations (10.11) and

10.8 | Derivatives and Risk Management

(10.12). Alternatively, we can find the value of q such that $S_0 e^{-qT}$ is equal to the right hand side of equation (10.17) and then use equation (10.16) with this value of q.

Continuing the example of the previous section, assume that the current stock price is 100, the volatility is 23%, the risk-free rate is 8% and a dividend of Rs 2 will be received in a month's time. We use equation (10.17) to compute \dot{S}_0 and use this as the value of *S* in equations (10.11) and (10.12) to find the option values. The value of a three-month call option with an exercise price of 105 is 2.22 while the put option with the same exercise price is worth 8.12 as shown below :

$$\begin{split} \dot{S}_0 &= 100 - e^{-0.08/12} \ x \ 3 = 97.0199 \\ d_1 &= \frac{\ln \frac{97.0199}{105} + \frac{0.08}{4} + \frac{1}{2} \ 0.23^2 \frac{1}{4}}{0.23 \sqrt{1/4}} = -0.455925 \\ d_2 &= d_1 - 0.23 \sqrt{\frac{1}{4}} = -0.455925 - 0.115000 = -0.570925 \\ N(d_1) &= 0.324222 \ N(d_2) = 0.284025 \ N(-d_1) = 0.675778 \ N(-d_2) = 0.715975 \\ e^{-0.08/4} &= 0.980199 \\ C &= 100 \times 0.970199 \times 0.324222 - 105 \times 0.980199 \times 0.284025 = 2.22 \\ P &= 105 \times 0.980199 \times 0.715975 - 100 \times 0.970199 \times 0.675778 = 8.12 \end{split}$$

The alternative method is to observe that $\dot{S}_0 = 97.0199$ is the same as assuming a continuous dividend yield of 12.1015% because $100 \times e^{-0.121015/4} = 97.0199$. We can easily verify that using this value of q in equation (10.16) also produces the same put and call values of 2.22 and 8.12:

$$d_{1} = \frac{\ln \frac{100}{105} + \frac{0.08 - 0.121015}{4} + \frac{1}{2} 0.23^{2} \frac{1}{4}}{0.23 \sqrt{1/4}} = -0.455925$$

$$d_{2} = d_{1} - 0.23 \sqrt{\frac{1}{4}} = -0.455925 - 0.115000 = -0.570925$$

$$N(d_{1}) = 0.324222 N(d_{2}) = 0.284025 N(-d_{1}) = 0.675778 N(-d_{2}) = 0.715975$$

$$e^{-0.08/4} = 0.980199 e^{-0.121015/4} = 0.970199$$

$$C = 100 \times 0.970199 \times 0.324222 - 105 \times 0.980199 \times 0.284025 = 2.22$$

$$P = 105 \times 0.980199 \times 0.715975 - 100 \times 0.970199 \times 0.675778 = 8.12$$

10.6 OPTIONS ON STOCK INDICES

Consider a stock index like the S&P 500 index in the United States. As the name suggests, this index consists of 500 stocks. Many of these stocks pay dividends quarterly, but each stock may pay the dividend on a different date. This means that the dividends on this index are paid in small bits all through the year. It is quite reasonable to assume that the dividends are paid continuously throughout the year. To value an option on the S&P 500, therefore, we compute the weighted average dividend yield of all stocks in the index and set q to this average dividend yield. Equation (10.16) then gives the value of options on a stock index.

For example, if the Nifty index is quoting at 1500, the volatility is 15%, the risk-free rate is 7%, the dividend yield on the index is 1%, then the price of a one-month call option with a strike of 1550 is 10.36 while the price of a one-month put at the same strike is 52.60:

$$d_{1} = -.6201 \ d_{2} = -.6634$$

$$N(d_{1}) = .267587 \ N(d_{2}) = .253528 \ N(-d_{1}) = .732413 \ N(-d_{2}) = .746472$$

$$e^{-0.07/12} = 0.994184 \ e^{-0.01/12} = 0.999167$$

$$C = 1500 \times 0.999167 \times 0.267587 - 1550 \times 0.994184 \times 0.253528 = 10.36$$

$$P = 1550 \times 0.994184 \times 0.746472 - 1500 \times 0.999167 \times 0.732413 = 52.60$$

As already discussed, it is not necessary that the dividend yield be constant all through the year. What is required is that q must be set equal to the average dividend yield during the life of the option. Thus, any seasonal variation in q can easily be handled.

This refinement is particularly important in India where companies tend to pay annual dividends and most companies use the fiscal year as their accounting year. This produces a strong seasonality in the dividend yield with significant dividend payments during June, July and August and very low dividend payments in other months of the year. While valuing index options in India, it is necessary to use a seasonally adjusted q.

Moreover, the commonly used indices in India have only a small number of stocks—the Sensex has only 30 stocks and the Nifty has only 50 stocks. It is possible that in some months of the year, only one or two constituents of the index are expected to pay dividends and the timing and quantum of these dividends is known. In this situation, it may be appropriate to use the model with known dividends as in equation (10.17).

However in this case, it is necessary to account for the weightage of each stock in the index. For example, suppose the Nifty is trading at 1500 and stock X with a market price of Rs 5,000 per share pays a dividend of Rs 200 per share. Suppose also that stock X has 5% weight in the Nifty so that out of the total index value of 1500, stock X accounts for 75 ($1500 \times 0.05 = 75$). If we regard Nifty as a portfolio of stocks, it contains 0.015 shares of stock X ($0.015 \times 5000 = 75$). The Nifty portfolio would thus receive a dividend of $0.015 \times Rs200 = Rs 3$. This whole computation can be summarized as:

$$200 \ \frac{1500 \times 0.05}{5000} = 3$$

Thus we would modify equation (10.17) as follows:

$$\dot{S}_0 = S_0 - \sum_{i=1}^k e^{-rt_i} D_i \frac{S_0 w_i}{Y_i}$$
(10.18)

where S_0 is the value of the index, company *i* pays a dividend of D_i per share at time t_i , company *i* has a weight w_i in the index and the market price of one share of company *i* is Y_i .

For example, assume that the Nifty index is quoting at 2000, the volatility is 14%, the risk free rate is 7% and that the only dividend during the next two months is a dividend of Rs 1.50 per share on a stock that is quoting at Rs 450 and has a weightage of 6% in the index. The equation (10.18) is used to compute \dot{S}_0 and this is used as the value of S in Equations (10.11) and (10.12) to find the option values. The price of a two-month call option with a strike of 2025 is 18.62 while the price of a one-month put at the same strike is 92.13:

10.10 | Derivatives and Risk Management

$$\begin{split} \dot{S}_0 &= 2000 - e^{-0.08/12} \ge 5 \frac{2000 \ge 0.06}{450} = 1998.6755 \\ d_1 &= \frac{\ln \frac{1998.6755}{2100} + \frac{0.08}{6} + \frac{1}{2} \ 0.14^2 \ \frac{1}{6}}{0.14 \ \sqrt{1/6}} = -0.603379 \\ d_2 &= d_1 - 0.14 \ \sqrt{\frac{1}{6}} = -0.603379 - 0.057155 = -0.660533 \\ N(d_1) &= 0.273128 \ N(d_2) = 0.254456 \ N(-d_1) = 0.726872 \ N(-d_2) = 0.745544 \\ e^{-0.08/6} &= 0.986755 \\ C &= 1998.6755 \ge 0.273128 - 2100 \ge 0.986755 \ge 0.254456 = 18.62 \\ P &= 2100 \ge 0.986755 \ge 0.745544 - 1998.6755 \ge 0.726872 = 92.13 \end{split}$$

Another way of doing this is again to convert the known dividend into an equivalent index dividend yield. An approximate way of computing the index dividend yield is as follows. The dividend yield on the stock that is paying a dividend of 5 on a price of 450 during the next two months is $\frac{5}{450} \frac{1}{2/12} = 6.67\%$. But this stock has only a 6% weight in the index and so the index dividend yield is only 6% ×

6.67% or 0.40%.

This computation is only an approximation because it ignores the timing of the payment of the dividend. The exact equivalent index dividend yield is 0.3975% which may be verified as $\hat{S}_0 = 2000 x e^{-0.003975/6} =$ 1998.6755. If we use this dividend yield as the value of q and use equation (10.16), we obtain the same prices of the options:

$$d_{1} = \frac{\ln \frac{2000}{2100} + \frac{0.08 - 0.003975}{6} + \frac{1}{2} 0.14^{2} \frac{1}{6}}{0.14\sqrt{1/6}} = -0.603379}$$

$$d_{2} = d_{1} - 0.14 \sqrt{\frac{1}{6}} = -0.603379 - 0.057155 = -0.660533}$$

$$N(d_{1}) = 0.273128 N(d_{2}) = 0.254456 N(-d_{1}) = 0.726872 N(-d_{2}) = 0.745544$$

$$e^{-0.08/6} = 0.986755 e^{-0.003975/6} = 0.999338$$

$$C = 2000 \times 0.999338 \times 0.273128 - 2100 \times 0.986755 \times 0.254456 = 18.62$$

$$P = 2100 \times 0.986755 \times 0.745544 - 2000 \times 0.999338 \times 0.726872 = 92.13$$

10.7 OPTIONS ON CURRENCIES

When we consider an option on a foreign currency, the notion of a continuous dividend yield applies perfectly because the foreign currency earns interest at the foreign risk-free rate. Thus we can use equation (10.16) with q set equal to the foreign interest rate which we denote r_f . The option value is then given by:

$$C = S_0 e^{-r_f T} N(d_1) - X e^{-r T} N(d_2)$$

$$P = X e^{-rT} N(-d_2) - S_0 e^{-r_f T} N(-d_1)$$

The Black-Scholes Option Pricing Model | 10.11

$$d_{1} = \frac{\ln \frac{S_{0}}{X} + (r - r_{f})T + \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}}$$

$$d_{2} = \frac{\ln \frac{S_{0}}{X} + (r - r_{f})T - \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}} = d_{1} - \sigma\sqrt{T}$$
(10.19)

For example consider a three month option on the rupee-dollar rate at an exercise price of 44 when the spot rate is 45, the Indian interest rate is 5%, the US interest rate is 2% and the volatility of the exchange rate is 11%. The price of a call is 1.77 while the price of a put is 0.45 as shown below:

$$d_{1} = \frac{\ln \frac{45}{44} + \frac{0.05 - 0.02}{4} + \frac{1}{2} 0.11^{2} \frac{1}{4}}{0.11 \sqrt{1/4}} = 0.572461$$

$$d_{2} = d_{1} - 0.11 \sqrt{\frac{1}{4}} = 0.572461 - 0.055000 = 0.517461$$

$$N(d_{1}) = 0.716495 N(d_{2}) = 0.697583 N(-d_{1}) = 0.283505 N(-d_{2}) = 0.302417$$

$$e^{-0.05/4} = 0.987578 \quad e^{-0.02/4} = 0.995012$$

$$C = 45 \times 0.995012 \times 0.716495 - 44 \times 0.987578 \times 0.697583 = 1.77$$

$$P = 44 \times 0.987578 \times 0.302417 - 45 \times 0.995012 \times 0.283505 = 0.45$$

10.8 OPTIONS ON FUTURES

When the Black-Scholes valuation model was derived with dividends, the risk-free return earned on the stock was divided into two components—dividends at the rate q and capital appreciation at the rate r - q. When one considers applying this model to options on futures, it is most convenient to ask what the rate of growth (capital appreciation) of the futures price is? The answer is that the rate of growth is zero. This is evident from the relation:

$$F_0 = E^*[S_T]$$

which does not depend on time T at all. The futures price depends only on the date of maturity of the contract and not on how far away one is from that maturity date. The expected rate of growth of the futures price is therefore zero. Under risk-neutral valuation, we can think of the futures contract, therefore, as paying a continuous dividend at the risk-free rate. We can thus set q = r in equation (10.16) to value options on futures:

$$C = e^{-rT} (F_0 N (d_1) - X N (d_2))$$

$$P = e^{-rT} (X N (-d_2) - F_0 N (-d_1))$$

$$d_1 = \frac{\ln \frac{F_0}{X} + \frac{1}{2} \sigma^2 T}{\sigma \sqrt{T}}$$

10.12 | Derivatives and Risk Management

$$d_{2} = \frac{\ln \frac{F_{0}}{X} - \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}} = d_{1} - \sigma\sqrt{T}$$
(10.20)

It may appear odd to describe futures as paying a dividend at the rate r when it does not produce any cash flow at all. There are many ways to convince oneself that this assumption gives the correct option valuation. One approach is to compare buying the stock to buying the future. Buying the stock involves incurring a cost of carry at the rate r which is saved when buying futures. This saving can be interpreted as the dividend at the rate r. Another approach is the one taken in the preceding paragraphs: if the growth rate r - q of the asset is zero, then r must be equal to q. A third approach is to deduce things from first principles by going through the derivation of the Black-Scholes formula again after accounting for the fact that instead of $E[S_T] = e^{rT}S_0$ as in equation (10.3), one now has $E[F_T] = F_0$. Whichever approach may be taken, one would arrive at the equation (10.20) for the price of options on futures.

As the futures and the spot price are identical at maturity, this formula can actually be used to value options on any asset in terms of the futures price instead of the spot price. For example, an option on currencies that we valued in equation (10.19) in terms of the spot exchange rate S_0 can also be valued using equation (10.20) if F_0 is the futures or forward¹ exchange rate.

For example, consider a three month option on the rupee-dollar rate at an exercise price of 44 when the forward rate is 44.75, the Indian interest rate is 5%, the US interest rate is 2% and the volatility of the exchange rate is 11%. The price of a call is 1.75 while the price of a put is 0.45 as shown below. These values are very close to the values that we are computed earlier using the spot exchange rate because the forward rate of 44.75 is very close to the theoretical value of 44.78, computed using the cost of carry model.

$$d_{1} = \frac{\ln \frac{44.7500}{44} + \frac{1}{2} 0.11^{2} \frac{1}{4}}{0.11 \sqrt{1/4}} = 0.562078$$

$$d_{2} = d_{1} - 0.11 \sqrt{\frac{1}{4}} = 0.562078 - 0.055000 = 0.507078$$

$$N(d_{1}) = 0.712969 N(d_{2}) = 0.693950 N(-d_{1}) = 0.287031 N(-d_{2}) = 0.306050$$

$$e^{-0.05/4} = 0.987578$$

$$C = 0.987578 (44.7500 \times 0.712969 - 44 \times 0.693950) = 1.75$$

$$P = 0.987578 (44 \times 0.306050 - 44.7500 \times 0.287031) = 0.45$$

Similarly, options on stock indices can be valued using equation (10.20), if we use the index futures price for F_0 . The advantage this has is that we do not need to estimate the dividend yield—this estimate is impounded in the index futures price.

In the Indian context, where index futures are mispriced relative to the cash index because of short sale restrictions and other barriers to index arbitrage, it is often preferable that index options are treated as options on index futures. The implicit assumption is that the short sale restrictions affect futures and options in the same way so that while both futures and options are mispriced relative to cash, they are correctly priced relative to each other.

¹ In most of the models discussed in this book, the futures and forward prices are the same as discussed in Chapter 3.
The Black-Scholes Option Pricing Model | 10.13

For example, assume that the Nifty index is quoting at 2000, the volatility is 14%, the risk-free rate is 7% and two month Nifty futures are at 1995. The price of a two month Nifty call with an exercise price of 1900 is 127.95 while the price of a put at the same exercise price is 7.78.

$$d_{1} = \frac{\ln \frac{44.7500}{44} + \frac{1}{2} 0.11^{2} \frac{1}{4}}{0.11 \sqrt{1/4}} = 0.562078$$

$$d_{2} = d_{1} - 0.11 \sqrt{\frac{1}{4}} = 0.562078 - 0.055000 = 0.507078$$

$$N(d_{1}) = 0.712969 N(d_{2}) = 0.693950 N(-d_{1}) = 0.287031 N(-d_{2}) = 0.306050$$

$$e^{-0.05/4} = 0.987578$$

$$C = 0.987578 (44.7500 \times 0.712969 - 44 \times 0.693950) = 1.75$$

$$P = 0.987578 (44 \times 0.306050 - 44.7500 \times 0.287031) = 0.45$$

Note that unless there are large dividends expected over the next two months, the Nifty futures are mispriced since they are trading below the spot Nifty. Pricing the options relative to the futures is likely to be more accurate than pricing them relative to the spot Nifty.

10.9 OPTIONS ON COMMODITIES

As understood from the discussion of commodity futures in Chapter 3, commodities can be regarded as paying a continuous dividend equal to their convenience yields. We can then set q equal to the estimated convenience yield and use equation (10.16). However, the best estimate of the convenience yield is probably contained in the futures price. In this case, it is easiest to value commodity options using equation (10.20) with the futures price.

10.10 AMERICAN OPTIONS

The Black-Scholes models that have been discussed so far are strictly speaking, valid only for European options. To derive these models we discounted the risk neutral expectation of the payoff at the maturity of the option. This valuation assumes that the option can be exercised only at maturity which is the case for European options.

However, the Black-Scholes formula is frequently used to value American options too. There are several reasons why this is more reasonable than it looks.

First, an American call option on a non dividend paying stock will never be exercised before maturity and is therefore no different from a European option. The reason is that an option is always worth more when alive, than when it is dead. This was discussed in Chapter 6 when it was demonstrated that the value of a call option is always more than the excess of the current stock price over the discounted exercise price $(S_0 - X e^{-rt})$. The reason the option is worth more is the insurance value of the option. Immediate exercise gives the holder an excess of the current stock price over the exercise price. By exercising the option, the holder suffers two losses:

1. The holder loses the insurance value of the option. $C - (S_0 - X e^{-rt})$.

10.14 | Derivatives and Risk Management

2. He has to pay the exercise price immediately, instead of paying it only at the maturity of the option. The loss here is the interest loss—the difference between the exercise price and the discounted exercise price. $X - X e^{-rt}$.

Consider the first example in this chapter, where a three-month European call option was valued with an exercise price of 105 when the current stock price was 100, the volatility 23% and the risk free rate, 8%. What would happen if this option were really American? There is no question of an early exercise at this point because the option is out of the money. But suppose that one month has passed and the stock price has risen to 108. Should the holder exercise the option? Immediate exercise fetches the holder 108 -105 = 3 while the Black-Scholes formula explains that the European option is worth 6.54 which is much more than 3.34:

$$d_{1} = \frac{\ln \frac{108}{105} + \frac{0.08}{6} + \frac{1}{2} \ 0.23^{2} \frac{1}{6}}{0.23 \sqrt{1/6}} = 0.488967$$
$$d_{2} = d_{1} - 0.23 \sqrt{\frac{1}{6}} = 0.488967 - 0.093897 = 0.395069$$
$$N(d_{1}) = 0.687567 \ N(d_{2}) = 0.653604 \text{ while } e^{-0.08/6} = 0.986755$$
$$C = 108 \times 0.687567 - 105 \times 0.986755 \times 0.653604 = 6.54$$

We can break up this difference (6.54 - 3.00 - 3.54) into its two components:

- 1. The insurance value of the option: $C (S_0 Xe^{-rt}) = 6.54 4.39 = 2.15$
- 2. The interest loss on the exercise price: $X Xe^{-rt} = 105.00 103.61 = 1.39$

Even if the holder of an option wants to book his profits and exit, he should sell the option instead of exercising it.

This means that the use of the Black-Scholes formula is perfectly valid even for American call options, provided the stocks do not pay a dividend. If a stock pays a dividend, this can no longer be asserted because we can only say that $C > S_0 - X e^{-rt}$ and not $C > S_0 - X e^{-rt}$. However if the dividends are modest, S_0 is not much lower than S_0 and C is unlikely to be lower than $S_0 - X e^{-rt}$. Even if C is slightly lower than $S_0 - X e^{-rt}$, it may still not be desirable to exercise the option immediately because of the interest loss $X - X e^{-rt}$. If options that are close to the money $S_0 \approx X$ are considered, immediate exercise makes sense only when the dividend yield (the equivalent q that we have computed so often in this chapter) is significantly more than the interest rate.

Let us change the above example to one where a dividend of Rs 2 per share is due in a month's time. The value of the option of Rs 3.59 is still more than the value of Rs 3 that can be had by immediate exercise:

$$\dot{S}_0 = 108 - e^{-0.08/12} \times 5 = 103.0332$$
$$d_1 = \frac{\ln \frac{103.0332}{105} + \frac{0.08}{6} + \frac{1}{2} 0.23^2 \frac{1}{6}}{0.23\sqrt{1/6}} = -0.012430$$

The Black-Scholes Option Pricing Model | 10.15

$$d_2 = d_1 - 0.23 \sqrt{\frac{1}{6}} = -0.012430 - 0.093897 = -0.106327$$

N(d_1) = 0.495041 N(d_2) = 0.457661 while e^{-0.08/6} = 0.986755
C = 103.0332 × 0.495041 - 105 × 0.986755 × 0.457661 = 3.59

To understand how this comes about, one observes that the dividend is large enough to swamp the insurance value of the option $C = 3.59 < S_0 - X e^{-rt} = 108 - 103.61 = 4.39$ and $C - (S_0 - X e^{-rt}) = 3.59 - 4.39 = -0.80$. However, the interest loss on the exercise price $X - X e^{-rt} = 105.00 - 103.61 = 1.39$ more than makes up for the -0.80. The option is worth more than immediate exercise by -0.80 + 1.39 = 0.59.

In most realistic situations, American call options would never be exercised prior to their maturity. They are therefore worth exactly the same as European options. Hence the Black-Scholes formula is a very good approximation for the value of an American call. The only exception is when large dividends are due, during the life of the option.

The situation regarding American puts is somewhat different. Here, early exercise is possible even if the stock pays no dividends because there are two forces acting in opposite directions. When the holder exercises early:

- 1. The holder loses the insurance value of the option. $P (X e^{-rt} S_0)$.
- 2. He receives the exercise price immediately instead of getting it only at the maturity of the option. There is an interest gain here—the difference between the exercise price and the discounted exercise price $X X e^{-rt}$.

For deep in the money puts, the insurance value is negligible while the interest gains from early exercise is substantial. Early exercise of deep in the money American puts is quite common. For example, if the current stock price is 90, the volatility is 23% and the risk- free rate is 8%, then the value of a two-month European put option with an exercise price of 105 is 13.87, which is less than the value of 15 that can be obtained by immediate exercise:

$$d_{1} = \frac{\ln \frac{90}{105} + \frac{0.08}{6} + \frac{1}{2} \cdot \frac{0.23^{2}}{6}}{0.23 \sqrt{1/6}} = -1.452750$$

$$d_{2} = d_{1} - 0.23 \sqrt{\frac{1}{6}} = -1.452750 - 0.093897 = -1.546647$$

$$N(-d_{1}) = 0.926853 \quad N(-d_{2}) = 0.939026 \text{ while } e^{-0.08/6} = 0.986755$$

$$P = 105 \times 0.986755 \times 0.939026 - 90 \times 0.926853 = 13.87$$

In this case, the option is so deep in the money that the insurance value of the option is very small: $P - (Xe^{-rt} - S_0) = 13.87 - (103.61 - 90) = 0.26$. As compared to this, the interest gain from early exercise is quite large: $X - Xe^{-rt} = 105 - 103.61 = 1.39$. The net effect is that immediate exercise is worth 1.23 more than a European option.

For near money options, however, the loss of insurance value is likely to exceed the interest gain and immediate exercise is not worthwhile. In the earlier example, if the current stock price were 100 instead of 90, then it is better to hold on to the option. Immediate exercise fetches only 105 - 100 = 5 while the European option is worth 5.89:

10.16 | Derivatives and Risk Management

$$d_{1} = \frac{\ln \frac{100}{105} + \frac{0.08}{6} + \frac{1}{2} 0.23^{2} \frac{1}{6}}{0.23 \sqrt{1/6}} = -0.330665$$
$$d_{2} = d_{1} - 0.23 \sqrt{\frac{1}{6}} = -0.330665 - 0.093897 = -0.424562$$
$$N(-d_{1}) = 0.629551 N(-d_{2}) = 0.664422 \text{ while } e^{-0.08/6} = 0.986755$$
$$P = 105 \times 0.986755 \times 0.664422 - 100 \times 0.629551 = 5.89$$

Since the option is closer to the money, the insurance value of the option is larger than before: $P - (X e^{-rt} - S_0) = 5.89 - (103.61 - 100) = 2.28$ and this is well above the interest gain of 1.39 which is the same as before. There is now a net advantage of 0.89 to holding on to the option.

Of course, even if immediate exercise is not worthwhile, the option need not be held on to maturity. Sometime during the life of the option, if it moves significantly into the money, it may be optimal to exercise it. Thus, the American option is probably worth more than 5.89 because one may exercise sometime before maturity.

Another factor is dividends, whose effect on put options is the opposite of the effect on call options. Dividends make it less profitable to exercise American put options early. This is because the payment of the dividend would reduce the stock price and make the put more valuable.

The net effect of all this is that for American put options which are not too deep in the money, the Black-Scholes formula is a good approximation. Many traders therefore value exchange traded American puts using the Black-Scholes formula with one simple adjustment: if the Black-Scholes value is less than intrinsic value (X - S), they replace the Black-Scholes value by the intrinsic value. In our earlier example, where the current stock price is 90, the volatility is 23%, and the risk-free rate is 8%, and the value of a two-month European put option with an exercise price of 105 is 13.87, this approach would value an American option at the intrinsic value of 15 that can be obtained by immediate exercise. This is the same as assuming that an American option will either be exercised immediately or at maturity, and never in between.

A variety of methods are available to value American options more accurately. One approach is the binomial model discussed in the previous chapter. This model values American options as easily as European options. In terms of computation time, however, the binomial model takes much more time than Black-Scholes. Moreover, unlike Black-Scholes, the binomial model is difficult to accommodate in a pocket calculator.

A binomial model with 500 nodes produces a value of 6.17 for the American option with an exercise price of 105 when the current stock price is 100, the volatility is 23% and the risk- free rate is 8%. This is more than the Black-Scholes value of 5.89 for the European option and is also more than the intrinsic value of 5.00; this means that while the option should not be exercised immediately, there is some chance that it would be exercised at some point prior to maturity. The same binomial model produces a value of 15.00 for the American option when the current stock price is 90 instead of 100. This is the same as the intrinsic value because immediate exercise is indeed optimal for this deep in the money option.

Other analytical approximations to the valuation of American options have been developed, but they will not be covered in this book. From a practical point of view, it is believed that the choice is mostly between the inaccurate Black-Scholes and the time consuming binomial model.

The Black-Scholes Option Pricing Model | 10.17

10.11* APPENDIX ON LOG NORMAL DISTRIBUTION

This appendix is provided for completeness and does not contain anything new. It only provides a proof of the following two properties of the log normal distribution used in the derivation of the Black-Scholes formula.

Property I

If lnY is normally distributed with mean M and variance V, then the expectation of Y is given by:

$$E(Y) = e^{M + V/2}$$
(10.21)

Alternatively if E(Y) and V are known, then M can be computed as

$$M = \ln E [Y] - V/2 \tag{10.22}$$

Property II:

If ln Y is normally distributed with mean M and variance V, then the conditional expectation E(Y|Y>X)is given by:

$$E(Y | Y > X) = E(Y) N(d_1)/N(d_2)$$
(10.23)

where

$$d_2 = \frac{M - \ln X}{\sqrt{V}}$$
 and $d_1 = d_2 + \sqrt{V} = \frac{M - \ln X + V}{\sqrt{V}}$

Both of the above are proved by proving the following property.

Property III:

If ln Y is normally distributed with mean M and variance V, and Z is defined as equal to Y when Y > X and equal to zero otherwise, then the expectation of Z is given by

$$E(Z) = e^{M + V/2} N(d_1)$$
 where $d_1 = \frac{M - \ln X + V}{\sqrt{V}}$ (10.24)

Property I follows from Property III by setting X equal to 0 so that $d_1 = d_2 = \infty$ and $N(d_1) = 1$. Property II follows from Property III by noting that

$$E(Y \mid Y > X) = E(Z)/P(Y > X)$$

and using the fact derived in the text that

$$P(Y > X) = P(\ln Y > \ln X) = P\left(\frac{\ln Y - M}{\sqrt{V}} > \frac{\ln X - M}{\sqrt{V}}\right)$$
$$= P\left(z > \frac{\ln X - M}{\sqrt{V}}\right) = P\left(z < -\frac{\ln X - M}{\sqrt{V}}\right) = N\left(\frac{M - \ln X}{\sqrt{V}}\right) = N(d_2) \quad (10.25)$$

where

 $d_2 = \frac{M - \ln X}{\sqrt{V}}$ Proceeding to therefore, prove property III.

Optional section. It is recommended that this section be skipped on first reading as it requires the knowledge of calculus and integration.

10.18 | Derivatives and Risk Management

Making the change of variable $Q = \frac{\ln Y - M}{\sqrt{V}}$ we get

- (a) $Z = e^{Q\sqrt{V+M}}$ if $Q > \frac{\ln X M}{\sqrt{V}}$ and zero otherwise
- (b) Q is a standard normal variate with density

$$n(Q) = \frac{1}{\sqrt{2\pi}} \mathrm{e}^{-\frac{1}{2}Q^2}$$

The equation obtained is:

$$E[Z] = \int_{\frac{\ln X - M}{\sqrt{V}}}^{\infty} e^{Q\sqrt{V} + M} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}Q^2} dQ$$
(10.26)
$$E[Z] = \int_{\frac{\ln X - M}{\sqrt{V}}}^{\infty} \frac{1}{\sqrt{2\pi}} e^{Q\sqrt{V} + M - \frac{1}{2}Q^2} dQ$$

The exponent $Q\sqrt{V} + M - \frac{1}{2}Q^2$ can be rewritten as

$$-\frac{1}{2}(Q^2 - 2 Q\sqrt{V} - 2M) = -\frac{1}{2}(Q^2 - 2 Q\sqrt{V} + V - V - 2M)$$
$$= -\frac{1}{2}(Q - \sqrt{V})^2 + \frac{1}{2}V + M$$

Using another change of variable $W = (Q - \sqrt{V})$ we get

$$E[Z] = e^{\frac{1}{2}V} e^{M} \int_{\frac{\ln X - M - V}{\sqrt{V}}}^{\infty} n(W) dW = e^{M + \frac{1}{2}V} N(d_{1})$$
(10.27)

where

$$d_1 = -\frac{\ln X - M - V}{\sqrt{V}} = \frac{M - \ln X + V}{\sqrt{V}}$$

which is the same as Equation (10.24) or property III.

Chapter Summary

The Black-Scholes model assumes that the logarithm of the future stock price is normally distributed. The stock price itself is log normally distributed. The risk-neutral distribution is also log normal, but the expected return of the stock is changed to the risk-free rate. If the continuously compounded risk free rate is r and the variance of the

logarithmic return of the stock is σ^2 then the expected logarithmic return of the stock must be set equal to $r - \frac{1}{2}\sigma^2$

to ensure that the expected annual return of the stock is equal to the annual risk-free rate.

The Black-Scholes formulas for European call and options on a non-dividend paying stock are:

The Black-Scholes Option Pricing Model | 10.19

$$C = S_0 N (d_1) - X e^{-rT} N (d_2)$$

$$P = X e^{-rT} N (-d_2) - S_0 N (-d_1)$$

$$d_1 = \frac{\ln \frac{S_0}{X} + rT + \frac{1}{2} \sigma^2 T}{\sigma \sqrt{T}}$$

$$d_2 = \frac{\ln \frac{S_0}{X} + rT - \frac{1}{2} \sigma^2 T}{\sigma \sqrt{T}} = d_1 - \sigma \sqrt{T}$$

where C is the value of the call, P is the value of the put, S_0 is the current stock price, X is the exercise price of the call option, T is the maturity of the option, σ is the volatility of the stock and r is the continuously compounded risk-free rate.

The generalized Black-Scholes formula for an asset that pays a dividend continuously at a constant rate q is:

$$C = S_0 e^{-qT} N(d_1) - X e^{-rT} N(d_2)$$

$$P = X e^{-rT} N(-d_2) - S_0 e^{-qT} N(-d_1)$$

$$d_1 = \frac{\ln \frac{S_0}{X} + (r-q)T + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln \frac{S_0}{X} + (r-q)T - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

If the timing t_i and quantum D_i of the k dividends to be paid during the life of the option are known, we can simply subtract the present value of these known dividends and use $\dot{S}_0 = S_0 - \sum_{i=1}^{k} e^{-rt_i} D_i$ in lieu of S_0 in the Black-

Scholes formula.

The Black-Scholes formula for an option on a foreign currency is given by the following formula where r_f is the foreign interest rate:

$$C = S_0 e^{-r/T} N(d_1) - X e^{-r/T} N(d_2)$$

$$P = X e^{-r/T} N(-d_2) - S_0 e^{-r/T} N(-d_1)$$

$$d_1 = \frac{\ln \frac{S_0}{X} + (r - r_f)T + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln \frac{S_0}{X} + (r - r_f)T - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

The Black-Scholes formula for an option on a futures contract is:

$$C = e^{-rT} (F_0 N (d_1) - X N (d_2))$$

$$P = e^{-rT} (X N (-d_2) - F_0 N (-d_1))$$

10.20 | Derivatives and Risk Management

$$d_1 = \frac{\ln \frac{F_0}{X} + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}$$
$$d_2 = \frac{\ln \frac{F_0}{X} - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

This formula can also be used to value an option on an asset using its futures price instead of its spot price. Options on commodities can be valued by regarding the commodity as paying a continuous dividend equal to its

convenience yield. Alternatively, they can be valued using the futures price and the formula for options on futures. Turning to American options, if there are no dividends, American call options would never be exercised prior to their maturity. They are, therefore, worth exactly the same as European options. Even if there are dividends, in most realistic situations, the Black-Scholes formula is a very good approximation to the value of an American call. The only exception is when large dividends are due during the life of the option.

American put options can often be exercised before maturity. However, for American put options which are not too deep in the money, the Black-Scholes formula is a good approximation. Many traders, therefore, value exchange traded American puts using the Black-Scholes formula with one simple adjustment: if the Black-Scholes value is less than intrinsic value (X - S), they replace the Black-Scholes value by the intrinsic value.

Though there are some analytical approximations to obtain a more accurate valuation of the American put option, the simplest approach is to use the binomial model discussed in the previous chapter.

Suggestions for Further Reading

The Black-Scholes model was first derived in

- Black, Fischer and Scholes, Myron (1973) "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy*, 81, May-June 1973, 229–246.
- Merton, RC (1973) "Theory of Rational Option Pricing", *Bell Journal of Economics and Management Science*, 4, 141–83.

The extension of the model to futures contracts was done in

Black, F (1976) "Pricing of Commodity Contracts", Journal of Financial Economics, 3, 167–79.

The extension to foreign currency options was presented in

Garman, Mark B. and Kohlhagen, Steven W. (1983), "Foreign currency option values", *Journal of International* Money and Finance, 2(3), 231–237.

The following papers discuss the valuation of American puts:

- Barone Adesi, G and Whaley, RE (1987) "Efficient Analytic Approximation of American Option Values", *Journal* of Finance, 42, 301–20.
- Brennan, M and Schwartz, ES (1977) "Valuation of American Put Options", Journal of Finance, 32, 449-62.
- Geske, R and Johnson, HE (1984) "The American Put Option Valued Analytically", *Journal of Finance*, 39, 1511–1524.
- Johnson, HE (1983) "An Analytic Approximation for the American Put Price", *Journal of Financial and Quantitative* Analysis, 18, 141–48.

The Black-Scholes Option Pricing Model | 10.21

Problems and Questions

- 1. If the current stock price is 150, the volatility is 21% and the risk free rate is 7% continuously compounded, what is the value of a three-month call option with an exercise price of 160? What is the value of a three-month put option with the same exercise price?
- 2. If the current stock price is 130, the volatility is 30%, the dividend yield is 2% and the risk free rate is 5% continuously compounded, what is the value of a three-month call option with an exercise price of 140?
- 3. If the current stock price is 80, the volatility is 25% and the risk free rate is 5.5% continuously compounded, what is the value of a three-month put option with an exercise price of 60 if the stock is expected to pay a dividend of Rs 5 at the end of one month?
- 4. If the Nifty index is quoting at 4500, the volatility is 14%, the risk free rate is 6.25% continuously compounded, the dividend yield on the index is 1%, then what is the price of a one-month call option with a strike of 4650?
- 5. What is the price of a three month call option on the rupee-dollar rate at an exercise price of 41 when the spot rate is 40, the Indian interest rate is 6.5%, the US interest rate is 5% and the volatility of the exchange rate is 11%? Assume that both interest rates are on simple interest basis.
- 6. In problem 4, if one-month Nifty futures are trading at 4520 what is the price of the option if you use the Black formula for options on futures? Why is the answer different?
- 7. If one-month gold futures are trading at \$625 what is the value of a one-month put option with a strike of \$615 assuming a volatility of 10% and a risk free interest rate of 4.75% continuously compounded?
- 8. Suppose that the Nifty index is quoting at 4000, the volatility is 16%, the risk free rate is 6.5% continuously compounded and that the only dividend during the next quarter is a dividend of Rs 150 per share (payable after two months) on a stock that is quoting at Rs 4,500 and has a weightage of 2% in the index. What is the price of a three-month put option with a strike of 4000?

Chapter **ELEVEN**

Uses of Options: Simple Option Strategies

The objective of this chapter is to present some simple option strategies that involve only one or two options. The strategies in this chapter are simple in the additional sense that they are methods of taking a bullish or bearish view on an asset, or hedging an existing exposure. All the strategies discussed here can therefore be regarded as alternatives to buying or selling the underlying asset. There is a separate chapter on more complex strategies that involve several options or are designed to take views on the volatility of the asset, rather than its price trend.

11.1 BUYING OPTIONS INSTEAD OF TRADING FORWARDS OR SPOT

This section considers a number of situations where an investor or hedger might wish to *buy* an option instead of buying or selling the underlying itself in the spot, forward, or futures markets. In this section, the possibility of *selling* options will not be considered.

11.1.1 Buying a Call instead of Buying the Share

Consider an investor who wants to buy the share of a company because he thinks its price will go up from 100 to 125. One alternative to buying the share is to buy a call option on the share with a strike price of say 110. One can visualize three scenarios at the maturity of the option:

- 1 The share does rise to 125, he exercises the call at 110 and makes a profit of 15 less the option premium.
- 2 The share price rises, but to only 110, the option expires worthless and the investor does not profit from the rise in prices. He also loses the option premium.
- 3 The share price falls to 80 contrary to his expectation. He loses the option premium, but is protected from the loss of 20 in the share price.

There are both advantages and disadvantages in buying a call instead of buying a share. The biggest advantage is that the call allows the investor to profit from a rise in share prices without the risk of loss



Figure 11.1 Buying the share exposes the investor to both the upside and the downside

11.2 | Derivatives and Risk Management



Figure 11.2 Buying an at-the-money call retains the entire upside and eliminates the downside at the cost of a steep option premium

if the share falls for an unexpected reason. In many situations, this is a big advantage. The use of an option allows investors to take views on stocks more freely because they know that their downside is limited. Some investors, therefore buy the stock if they are completely confident about their analysis, and buy an option if they are less confident.

There are two disadvantages: (a) there is a cost (premium) for buying the option, and (b) the investors pr ofits only if the share price rises above the strike price. It is possible to reduce either of these disadvantages at the cost of increasing the other. For example, if an investor bought an option at a strike of 100 instead of at 110, he is guaranteed to profit from any rise in prices however small, but the option premium will now be much higher. For example, if the option maturity is three months, the risk free rate is 5 percent and the volatility is 25 percent, then the 110-Call costs only Rs 1.98 while the 100-Call would cost Rs 5.60.



Figure 11.3 Buying an out of the money call, eliminates the downside and retains a part of the upside at the cost of a modest option premium.

A comparison of all three alternatives is presented in Figures 11.1, 11.2 and 11.3.

11.1.2 Buying a Put Instead of Shorting the Share

The opposite of the above situation arises when the investor does not hold the share but has a bearish view on it. He could sell futures on the stock (or borrow the stock and short sell it if that is permitted).

This strategy, however, exposes him to serious losses (Figure 11.4) if his bearish view turns out to be wrong. The purchase of a put option allows him to benefit from his bearish view if it turns out to be right, while limiting his losses if the share rises against his expectations. The situation of the bearish investor buying a put is the exact mirror image of the bullish investor buying a call (see Figure 11.5). The advantages and disadvantages of using options instead of short selling are also very similar.



Figure 11.4 Short selling the stock or selling a future on the stock has an unlimited downside (if the stock rises) and a large upside (if the stock falls)



Figure 11.5 Buying an out of the money put retains a part of the upside (from a falling stock price) and eliminates the downside at the cost of a modest option premium

11.1.3 Hedging with a Put instead of Selling the Share

We now consider a slightly different situation. An investor is bullish on a share that he holds, but thinks of selling, because he's worried about the likelihood of the share price falling. An attractive alternative is to buy a put option (compare Figure 11.6 and Figure 11.7). If the share is trading at 100 and the investor thinks that the price could fall to 80 in the unlikely event of the company's profits coming out well below market expectations, the investor could in that case, buy a put option at 90. This limits his loss to Rs 10 in the unlikely event of a profit decline. On the other hand, if the profits are in line with market expectations, the share price will probably rise in accordance with his overall

11.4 | Derivatives and Risk Management

bullish view and he would benefit from the rise. Just as in the earlier case, one can consider different scenarios (for example, if the share price falls to 80, falls to 90 and then rises to 120) and compare the performance of the put option as against an outright sale of the share. This is left as an exercise for the reader to undertake.



Figure 11.6 Holding a share unhedged, exposes the investor to both the upside and the downside



Figure 11.7 Hedging a long stock position with an out-of-the money put eliminates most of the downside while retaining the upside at the cost of a modest option premium

11.1.4 Foreign Exchange Hedging with a Call instead of a Forward

Let us now consider a foreign exchange example. An Indian company has imported some raw materials and has to make a payment of \$100,000 in 30 days. The company is concerned that the dollar might appreciate against the Indian rupee, and is considering taking forward cover at the 30 day forward rate of Rs 43.90/\$. This would allow the company to lock in an exchange rate on its future payment obligation. On the other hand, there is a possibility that the dollar might actually depreciate to, say, Rs 43.20/\$ and the company may end up paying a lot more than what it would have paid, if it had not hedged (see Figures 11.8 and 11.9).

Uses of Options: Simple Option Strategies | 11.5



Figure 11.8 An unhedged payable is essentially a short position in the foreign currency. It has both upside potential and downside risk

A similar situation arose with Lufthansa several years ago when it hedged a significant part of a large dollar payment that was due for the purchase of Boeing aircraft. The dollar actually depreciated sharply and Lufthansa ended up paying far more than what it would have paid had it not hedged. The Chairman of Lufthansa almost lost his job as critics both within Lufthansa and without, accused him of incompetence.



Figure 11.9 Hedging the payable using a forward contract locks in the rate, but eliminates the upside potential completely

So how can a company protect itself from this risk? One solution is to use a call option. The Indian company referred to above, could buy a call option to buy \$100,000 at a strike price of Rs 43.90/\$. The



Figure 11.10 Hedging a payable with an at-the-money call eliminates the downside and retains the downside at the cost of a substantial option premium

11.6 | Derivatives and Risk Management

option ensures that the company will never have to pay more than Rs 43.90/\$, but the company is also free to buy the \$100,000 in the spot market on the due date if that is cheaper. It thus gets the best of both worlds (Figure 11.10).

Alas, there is no free lunch in this world and the option would cost the company a lot of money. If the option expires without worth, somebody can always accuse the management of wasting money, buying a worthless option. It is easy to criticize any decision in hindsight.

11.1.5 The Option Premium

In all the above examples, one of the major stumbling blocks in the use of options is the fact that the options cost money. The more useful they are, the greater the option premium. A lot of the option strategies that shall be discussed now, are attempts to deal with this difficulty. The key element in all of them is the *selling* of options to earn some option premium. In many cases, the premium for buying an option that we need is earned by selling an other option. Once again, one must remember that there is no free lunch. The option that is sold can impose losses on the seller if the holder exercises them. In many of these strategies it is hoped that the sold option will either not be exercised, or that the loss incurred by exercise will be tolerable. This assumption must always be carefully examined.

11.2 RANGE FORWARDS

Consider again, the example of the Indian company that has to make a payment of \$100,000 in 30 days, and is thinking of taking forward cover at the 30 day forward rate of Rs 43.90/\$ or buying a call option at this strike price. It the company buys a call option at Rs 43.90/\$ and also sells a put option at Rs 43.20/\$ the call option means that the company is fully protected against any dollar appreciation as its cost can be no worse than Rs 43.90/\$. However, because of the sold put option, it enjoys only a part of the upside of a dollar depreciation. Specifically, if the dollar depreciates to Rs 43.20/\$, then the put option forces the company to sell its \$100,000 at Rs 43.20/\$ when the spot rate is even better than the paid amount.

Effectively, if the dollar costs less than Rs 43.20, the company pays Rs 43.20 (it delivers into the sold put); if the dollar costs more than Rs 43.90, the company pays Rs 43.90 (it exercises the call); and if the dollar costs between Rs 43.20 and Rs 43.90, the company pays the actual rate (both the bought and sold options expire without value). This is often called a range forward because the company does not lock in a single forward rate, but locks in only a range of Rs 43.20 to Rs 43.90.

What the company gets by sacrificing part of the upside (dollar cheaper than Rs 43.20) is that it earns a part of the premium spent on buying the call at Rs 43.90. It gets only a part of the premium because the call at Rs 43.90 is at the money and therefore very expensive while the put at Rs 43.20 is out of the money and therefore relatively cheap. Thus this range forward still costs money though less than just buying a call at Rs 43.90.

The company can change its strategy to make the range forward zero cost (Figures 11.11, 11.12 and 11.13). For this to happen, both options (the option bought and the option sold) have to be roughly equally out of the money so that they are worth the same. An example might be to buy a call option at Rs 44.30/\$ and sell a put at Rs 43.50/\$. This means that the company locks in a rate in the range Rs 43.50/\$ to Rs 44.30/\$. Such a range forward might have near-zero cost. (Since the dealer with whom we trade options will charge a bid-ask spread and other commission, the option sold will actually have to be more valuable than the option purchased, for the range forward to be truly zero cost. This might imply a range of say Rs 43.50/\$ to Rs 44.35/\$.

Uses of Options: Simple Option Strategies | 11.7



Figure 11.11 The unhedged payable to be hedged using a range forward



Figure 11.12 The zero cost range forward used to hedged the payable



Figure 11.13 When the payable is hedged using a range forward, both the downside and the upside are limited. But within this range the payable is effectively unhedged

11.3 BULL AND BEAR SPREADS

Spreads also involve buying and selling options but the bought and sold options are of the same type (either both are calls or both are puts).

11.8 | Derivatives and Risk Management

Early in this chapter, we talked of buying a call instead of buying a share and saw that the major disadvantage was the cost of the option. The idea of a bull spread is to recover part of the cost by selling a call option at a higher price. The investor in that example was expecting the share price to go up from 100 to 150 and was considering buying a call option at 110. To turn this into a bull spread, he would sell a call option at 125. The sold option limits his upside, however high the share price rises. His gains are limited to 15 (less the cost of the spread) because he will have to sell the share at 125. In return, he gets back a part of the premium paid for the purchased option. The sold option at 125 is further out of the money and therefore, cheaper than the purchased option at 110. This bull spread does cost money to set up, but it promises a chance of a large profit at fairly low cost if the bullish view turns out to be right (Figure 11.14).



Figure 11.14 Bull Spread

A bear spread (Figure 11.15) on the other hand is set when the investor has a bearish view on the stock. In this case, the sold option is at a lower strike than the bought option. For example, selling a call at 80 and buying a call at 120 when the stock is trading at 100. If the bearish view turns out to be wrong,the investor could lose a maximum of 40 (when the stock rises to or above 120 and she has to sell it at 80). If the bearish view turns out right, the options expire worthless. How does she make money then? The answer is that the bear spread set up using call options earns premium (the in the money call at 80 is worth far more than the out of the money call at 120). The attentive reader would have noted that



Figure 11.15 Bear Spread

Uses of Options: Simple Option Strategies | 11.9

the diagram for the bear spread is similar to the diagram for the import payable hedged with a range forward. The reason for this is put call parity. The sold put in the range forward when combined with the pre-existing short position in the asset is equivalent to a sold call by put-call parity. Therefore, the diagrams are quite similar.

A bear spread is perhaps better set up using put options (Figure 11.16). An investor could, for example, buy a put option at 100 and sell a put option at 80. If the bearish view turns out right and the stock falls to 80 or lower, the spread makes a profit of Rs 20 (less the cost of the spread). If the bearish view turns out wrong and the stock rises, both options expire worthless and she loses the premium paid for the spread. In this case, the spread has a cost because the sold option at 80 is out of the money and is less valuable than the bought option which is at the money.

Similarly, a bull spread can also be set up with put options. In this case, premium is earned while setting up the spread, while on expiry either the options expire worthless or the investor loses money.



Figure 11.16 Bear Spread with Puts

11.4 COVERED OPTION WRITING

Covered call writing refers to selling call options on stocks that the investor already holds. If the option does get exercised, the investor simply delivers the stocks that he holds and collects the strike price. There are a variety of situations where covered call writing makes sense:

- Covered call writing can be regarded as an automatic system of profit booking. Suppose an investor buys a stock at 100 with a price target of 130. He would therefore, expect to book his profits and exit at 130. Selling a covered call with a strike price of 130 appears to do exactly this while also earning a premium in the process. Once again, we must remember that there is no free lunch in finance. The risk that exists here is that the stock rises to 130 and falls to 120 during the life of the option. The option might not get exercised (it certainly will not be exercised if it is European) and the investor would have lost the opportunity to book profits when the stock reached 130.
- Covered call writing might also make sense when the investor thinks that the stock has limited upside potential. Suppose an investor believes that the stock that he has bought at 100 is likely to rise gradually, but not go above 130 in the next three months. He could then sell a call at 130 and pocket the premium. If the stock does rise beyond 130, he loses the upside, but he regards this possibility as somewhat remote and is willing take to the risk.
- Covered call writing is also a way of converting future capital gains into current income for accounting purposes. Some institutional investors may be able to pay dividends only out of the

11.10 | Derivatives and Risk Management

current income and not out of unrealized capital gains. They may then be forced to sell their holdings to book the profits for accounting purposes. Covered call writing is an alternative that produces current income without having to sell the stocks. (Of course, if the call is exercised a sale does take place). Another aspect of covered call writing is that it moves income from the future to the present. Expected future capital appreciation is converted into current income.

Chapter Summary

It has been concluded in this chapter, that it is sometimes advantageous to take a bullish view on a stock by buying a call option on the stock, rather than buying the stock itself. Similarly, it is sometimes better to hedge a stock with a put rather than to sell the stock. Similar arguments apply to other assets and not just to stocks. It sometimes makes sense to hedge an import payable with a currency option.

In some cases, the upfront cost of the option can be reduced by selling another option. In the foreign exchange context, selling a put while buying a call leads to a range forward. The resulting hedged position is structurally the same as a bear spread which is achieved by selling a call at a low strike and buying a call at a higher strike, or by buying a put at a high price and selling a put at a low price. A bearish spread makes money if the underlying declines in value. If the sold option is at a lower price than the purchased option, then the position is a bull spread which makes money if the underlying rises in value.

Covered call writing is a way of earning premium income at low risk because the seller owns the underlying asset.

In all these cases, it is important to recognize that there is no free lunch. Every purchased option implies the payment of an option premium, while every sold option carries with it the risks of a loss (or a reduction in profit) if the option is exercised.

Suggestions for Further Reading

Baird, Allen. Jan (1993) "Position Risk Profiles", Chapter 4 in Option Market Making—Trading and Risk Analysis for the Financial and Commodity Option Markets, New York, Wiley.

Problems and Questions

- 1. The rupee-dollar exchange rate is Rs 42/\$, the Indian interest rate is 7.5%, the US interest rate is 4.5% and the volatility of the exchange rate is 12%. Compare the performance of the following strategies for an exporter over a one month horizon:
 - (a) no hedging
 - (b) hedging with futures
 - (c) hedging with at-the-money dollar puts
 - In your comparison, consider the following exchange rates at maturity: Rs 41/\$, Rs 42/\$ and Rs 43/\$.
- 2. In the above problem, if the company wants downside protection at Rs 41/\$ and is not willing to pay a premium, design a range forward strategy.
- 3. In the above problem, if the banks sells options at 10% above the price at which it buys them (5% above and below fair value respectively), how will your range forward strategy have to be modified?
- 4. In the above problem, if the company is willing to sell calls for twice the quantity for which it buys downside protection, at what strike can the calls be sold?

Uses of Options: Simple Option Strategies | 11.11

- 5. A mutual fund manager thinks that valuations in the Indian stock market are full if not stretched and does not see any potential for large gains over the next two months. He decides to sell Nifty calls at a strike of 4000 when Nifty is trading at 3900, the risk free rate is 5.5% continuously compounded and the volatility is 15%. Draw a profit diagram for the strategy. At what index level does he break even?
- 6. In problem 5, if the fund manager wants to use the premium earned to buy downside protection for his portfolio, what strategy would you suggest?
- 7. In problem 5, if the fund manager is worried about the index rising contrary to his expectations, how can he limit losses by modifying his strategy?
- 8. Would you recommend that the fund manager in problem 6 use a ratio spread?

Chapter **Twelve**

The Greeks of the Black-Scholes Model

The option price in the Black-Scholes model is a function of the asset price, the exercise price, the volatility, the time to maturity, and the interest rate. Of these, only the exercise price remains fixed throughout the life of the option. All the other parameters can change over time. The Greeks describe how the option price changes when these parameters change. In the chapter we will look at the definition of the various Greeks, understand their behaviour and explain how the Greeks are used for hedging options and for risk management.

12.1 DELTA

We have already discussed the notion of option delta in Chapter 9. The option delta denoted by the Greek letter Δ is the change in the option value when the stock price changes by a small amount say Re 1.

Mathematically, the definition is $\Delta = \frac{\partial f}{\partial S}$ where *f* is the option price and *S* is the price of the underlying asset.

Buying the option, therefore, gives the same exposure to the stock price as buying Δ shares. For the same reason, the option can be hedged against stock price risk by selling Δ shares. If the stock price falls by Re 1, the option value falls by Rs Δ while the short position of Δ shares gains Rs Δ leaving the hedged position with no profit or loss. The same is true when the share price falls by Re 1.

The delta of a European call option on a non-dividend paying stock is given by

$$\Delta = N(d_1) \tag{12.1}$$

For put options on a non-dividend paying stock, the formula is

$$\Delta = -N(-d_1) = N(d_1) - 1 \tag{12.2}$$

The delta of an option is close to zero when the option is deep out of the money. When the option is deep in the money, the delta is close to 1 for a call and close to -1 for a put. When the option is at the money, the delta is close to 0.5 for a call and close to -0.5 for a put. In terms of stock prices (for a fixed strike), the position is as follows:

- 1. When the stock price is very low, the call is deep out of the money and has a delta close to zero. At this point, the put is deep in the money and has a delta close to -1.
- 2. When the stock price is close to the strike, the call delta is around 0.5 and the put delta is around -0.5.
- 3. When the stock price is very high, the call is deep in the money and has a delta close to 1. At this point, the put is deep out of the money and has a delta close to 0. In all cases, the delta of the put can be obtained by subtracting 1 from the delta of the call. This last fact is evident by comparing the formula for the put delta with that of the call delta. It can also be obtained simply from put call parity which says that a call plus a bond is the same as a put plus the stock. Hence, the delta of the call plus the delta of the bond (equal to 0) must be equal to the delta of the stock (equal to 1) plus the delta of the put.

12.2 | Derivatives and Risk Management

These facts can be seen in Figure 12.1 which shows the behaviour of the delta as the stock price changes. The call deltas are shown on the left axis and the put deltas on the right axis, but the graph of the delta is identical for both options. This reflects the fact that the put delta is simply 1 less than the call delta.



Figure 12.1 Delta of call and put option as a function of the asset price. The call delta rises from neatly zero for a deep out of the money call to nearly one for a deep in the money call. The put delta moves from nearly zero for a deep out of the money put to nearly minus one for a deep in the money put

Another interesting question to ask is what happens to the delta as the maturity of the option increases but the asset price is kept fixed. Figure 12.2 depicts the behaviour of call deltas as a function of maturity. This behaviour might appear perplexing, but has a very simple explanation. As the time to maturity increases, the discounted strike price $X e^{-rt}$ decreases. Regardless of whether X was more or less than S, the discounted strike $X e^{-rt}$ becomes less than S for sufficiently large t. In discounted terms, the call option



Figure 12.2 Call delta as a function of maturity. The call delta rises towards one as the matutity rises. This is because with a positive risk-free rate, the strike becomes much less than the forward price (risk-neutral expected price at maturity) when the maturity of the option is sufficiently long

The Greeks of the Black-Scholes Model | 12.3

moves deep into the money. As a result, the option delta moves towards one regardless of the strike price. Another way to look at it is to compare the strike not to the spot price of the asset, but to the forward price $S e^{rt}$ of the asset. Again regardless of whether X was more or less than S, the strike X becomes less than the forward price $S e^{rt}$ for sufficiently large t and this causes the delta to move towards one.

The most interesting behaviour is that of the in-the-money call whose delta falls and then rises. Close to maturity, a call that is even slightly in the money has a delta of almost unity. But a short period before maturity, it is not certain that the call will expire in the money. There is always a chance that the stock will move out of the money. This causes the delta to drop below one. Over much longer periods of time, the random fluctuations in the stock tend to cancel out (standard deviation rises only with the square root of time) and the mean growth rate (equal to the risk-free rate) dominates the picture. This forces delta back towards unity.

The same argument can be applied to put options (Figure 12.3). Over long maturities, the put moves deep out of the money in terms of either the discounted exercise price or the forward price. As a result, its delta moves towards zero. Again, the out-of-the money option whose delta is zero at expiry, shows non-monotonic behaviour. Its delta initially turns negative on the possibility that random fluctuations in the asset price could move it into the money prior to expiry. In the long run, however, these random fluctuations are swamped by the growth of the stock price at the risk free-rate.

One might ask what would happen to the behaviour of delta over time, if the distortion caused by the risk-free rate is eliminated. One way to offset this is to have the stock pay a dividend that reduces the growth rate of the stock price. Another way is to consider an option on a future and keep the futures price fixed instead of keeping the stock price fixed as the option maturity changes. For this purpose, we need to consider the Black-Scholes model for options on an asset that pays a continuous dividend yield.



Figure 12.3 Put delta as a function of matutity. The put delta tends towards zero as the maturity rises. This is because with a postive risk free-rate, the strike becomes much less than the forward price (risk-neutral expected price at maturity) when the maturity of the option is sufficiently long.

For call options on an asset that pays a continuous dividend at the rate q, the delta is given by:

$$\Delta = N(d_1) \mathrm{e}^{-qt} \tag{12.3}$$

12.4 | Derivatives and Risk Management

For put options on an asset that pays a continuous dividend at the rate q, the delta is given by:

$$\Delta = -N(-d_1)e^{-qt} = (N(d_1) - 1)e^{-qt}$$
(12.4)

For an option on a futures, q = r and the deltas are $N(d_1)e^{-rt}$ and $(N(d_1) - 1)e^{-rt}$ respectively for call and put options. Figure 12.4 shows the behaviour of the call delta as a function of maturity when the call is on a future instead of on the spot.

The call deltas now go to zero as the maturity increases. Again, the explanation is simple. The payoff of the option arises at maturity and under risk-neutral valuation, the change in the option price is the expected change in the payoff at maturity discounted at the risk-free rate. If the option is deep in the money, a Re 1 change in the asset price causes almost a Re 1 change in the expected payoff at maturity. However, when this is discounted to today at the risk-free rate, the delta drops below 1. If maturity is very far away, the delta would drop to zero.



Figure 12.4 Delta of call option on futures instead of spot. The delta goes to zero at long maturities because the payoff at expiry has to be discounted to today at the risk-free rate.

Exactly the same thing happens to a put option (Figure 12.5). The delta drops to zero because even if the payoff at maturity changes by -1, this has to be discounted to today and this causes the delta to drop towards zero. That both call and put deltas go to zero might appear to be at odds with put-call parity, but it is not. When a stock pays a dividend at the rate q, put- call parity says that a call plus a bond is the same as a put plus e^{-qt} stocks. As the stock pays dividends, the dividends have to be reinvested in the stock so that at maturity, the holding rises to 1 stock. Hence the delta of the put is equal to the delta of the call minus e^{-qt} . This is what we find in the equation for the put delta as well. As *t* becomes very large, the term e^{-qt} becomes very small and the call and put deltas come close to each other.

Our intuitive expectation that an at-the-money option should have a delta close to 0.5 for a call and -0.5 for a put regardless of maturity, does not hold true whether we define at-the-money relative to the spot price or the forward price¹. The delta of at-the-money options is close to 0.5 or -0.5 only for options expiring in a few months.

¹ The inquisitive reader might wonder whether the at-the-money call and put deltas would remain at 0.5 and -0.5 for all maturities, if the interest rate were set equal to zero so that all the discounting problems go away. Unfortunately, even this is not true, but it would take us too far afield to explain this phenomenon which arises out of the extreme skewness of the log normal distribution. This phenomenon is of no practical relevance in any case.

The Greeks of the Black-Scholes Model $\mid~12.5$



Måturty

Figure 12.5 Delta of put option on futures instead of spot. The delta goes to zero at long maturities because the payoff at expiry has to be discounted to today at the risk-free rate

From a hedging point of view, delta is often regarded as the least problematic of all the option Greeks. This is because the delta can be hedged using the underlying asset or futures on that asset. All the other Greeks can be hedged only using other options. From a practical point of view, this difference is very important because options are much less liquid than the futures or the underlying. Bid-ask spreads are much wider. Long option positions also require the payment of an options premium which is not required for a futures. For this reason, in many option trading and arbitrage strategies, it is assumed that the delta risk can be completely eliminated, and attention is focused on the remaining option Greeks.

12.2 GAMMA

The essence of an option is its non-linearity which manifests itself in the change in the delta from 0 for an out-of-the-money option to ± 1 for a deep-in-the-money option. The best measure of this non-linearity is the rate of change of the delta as the price of the underlying asset changes. Mathematically, this

measure known as gamma (Γ) can be defined as $\Gamma = \frac{\partial \Delta}{\partial S} = \frac{\partial^2 f}{\partial S^2}$ where f is the option price and S is the

price of the underlying asset.

For a European call or put option on a stock that does not pay dividends, the gamma is given by

$$\Gamma = \frac{N'(d_I)}{S\sigma\sqrt{t}} \tag{12.5}$$

where $N'(x) = (1/\sqrt{2\pi})e^{-\frac{1}{2}x^2}$ is the normal density function. For a European call or put option on a stock that pays dividends at the rate q, the gamma is given by

$$\Gamma = \frac{e^{-qt} N'(d_1)}{S\sigma\sqrt{t}}$$
(12.6)

12.6 | Derivatives and Risk Management

It has already been seen that the call and put deltas differ only by a constant (namely unity). It stands to reason, therefore, that the rate of change of the delta is the same for both call and puts. That is why the call and put gammas are the same. The gamma of an option is always positive. Negative gammas arise when options are sold.

Figure 12.6 shows the graph of gamma as the stock price changes. The graph looks very much like the bell-shaped curve of the normal density function with the peak arising when the asset price is close to the strike price, with the gamma falling off towards zero when the option goes deep in or out of the money. The similarity is obvious because the normal density is the numerator in the formula for gamma. There is however a difference—unlike the bell-shaped curve, there is an asymmetry in the graph of the gamma. This is caused by the presence of the asset price in the denominator of the formula for gamma. Because of this, the gamma is lower for high stock prices than for low stock prices when both are equally away from the money. In-the-money calls (and out-of-the-money puts) have higher gammas than out-of-the-money calls (and in-the-money puts) when the asset price is equally far from the strike price in percentage terms.



Figure 12.6 Gamma of call and put options are highest when the underlying price is close to the strike price. The gamma drops off to zero when the option is deep in or out of the money

We now turn to the behaviour of the gamma as a function of the time to maturity (Figure 12.7). The key observation here is that the gamma of an at-the-money option explodes as the option approaches its maturity date. This is because close to expiry, a slightly out-of-the-money option has a delta close to zero, while a slightly in-the-money option has a delta close to ± 1 . A relatively small change in the asset price can, therefore, cause a large change in the option delta. This is akin to saying that the option gamma is very high. The non-linearity of the option payoff is most pronounced near maturity. Far away from maturity, the non-linearity is smoothened away by the uncertainty of the asset price over the remaining life of the option.

If an option is in-the-money or out-of-the-money, then as the option approaches maturity, the option delta moves towards ± 1 or 0. Small changes in the asset price do not cause any change in the delta from these levels. Hence, the gamma decays to zero as the option approaches expiry.

The Greeks of the Black-Scholes Model | 12.7



Gàmmà of Càll or Put Opt ons Volat hty = 25%, Rf = 5%

Figure 12.7 The at-the-money option has a very high gamma close to expiry. Its gamma falls off rapidly to zero as the maturity increases. Out of-the-money options have near zero gammas close to expiry and slightly higher gamma at longer maturities

Since gamma is a measure of the non-linearity of the option, it is in some ways the most important of all the Greeks:

- Consider a person who has sold an option and is hedging the short position using delta hedging. As we have already seen in Chapter 9, the cost of this delta hedging arises because of the need to buy high and sell low. When the stock price rises, the delta of a call option rises, forcing the option seller to buy more stocks at the higher price to maintain delta neutrality. Similarly, when the stock price falls, the delta of a call option falls, forcing the option seller to sell stocks at the lower price to maintain delta neutrality.
- The more the delta changes with changing stock prices, the greater the costs of dynamic hedging. This is another way of saying that options with large gamma are more difficult to replicate and are more risky for the seller who intends to do dynamic hedging. Moreover, a large and sudden move in the stock price can force a very costly rebalancing of the delta hedge. We have used the example of a call option, but the situation is the same for put options. The delta of a put is negative and the option seller sells stocks. A rise in the stock price increases the delta (it becomes less negative) and the put seller has to buy back some of the stocks shorted as part of the original delta hedge. Again, the option seller ends up buying high and selling low.
- A person who has bought an option is in the opposite position. Suppose she decides to eliminate the delta risk by selling stocks equal to the delta of the call option. Fluctuations in the stock price work to her advantage. When the stock price rises and the delta of the call option rises, she sells more stocks at the higher price to maintain delta neutrality. Similarly, when the stock price and the delta fall, she buys stocks at the lower price to maintain delta neutrality. She thus buys low and sells high, making a profit from the fluctuations in the stock price. Again, the gamma determines how much the delta changes and, therefore, how much profit the option buyer makes.

Gamma thus measures the extent to which the option holder benefits (and the option seller suffers) from stock price fluctuations. The relation can actually be quantified²: the realized gain from a delta

² See also Eq. (12.14) and footnote 3.

12.8 | Derivatives and Risk Management

hedged long option over a short period is equal to half the gamma times the square of the asset price times the realized volatility of the underlying asset during that short period. If the option is held to maturity (dynamically delta hedged), then the the realized gain is equal to half the average gamma (average over the life of the option) times the square of the asset price times the realized volatility of the underlying asset.

Of course, there is no free lunch: the option premium is equal to half the expected average gamma (average over the life of the option) times the square of the asset price times the expected (implied) volatility of the underlying asset. Therefore, the delta hedged option buyer makes money only if the realized volatility is more than the implied volatility. In fact, the gain is equal to half the average gamma times the square of the asset price times the difference between the realized and implied volatilities. If the realized volatility is less than the implied, then it is the delta hedged option seller who makes money.

This makes it clear that holding a delta hedged option (long or short) can be regarded as a bet on how volatile the asset price will be during the life of the option. Viewed this way, options with large gammas are desirable; these options produce large gains if the realized volatility behaves as expected.

12.3 THETA

Theta is the change in the value of the option with the passage of time. Mathematically, the definition is

 $\Theta = -\frac{\partial f}{\partial t}$ where f is the option price and t is the time to maturity of the option. The negative sign is

required because as we move forward in calendar time, we are moving closer to maturity and the time to maturity (t) decreases. It must be emphasized that in the definition of theta we are assuming that there is a passage of time, but nothing else changes. The asset price in particular does not change. In other words, theta actually tells us what happens to the option price when one day elapses without any change in the asset price. This remark applies to all the other Greeks as well; each Greek tells us what happens if one parameter changes while everything else remains unchanged. In the language of mathematics, these are all 'partial' derivatives.

For a European call option on a stock that does not pay dividends, the theta is given by

$$\Theta = -\frac{SN'(d_1)\sigma}{2\sqrt{t}} - Xe^{-rt}N(d_2)$$
(12.7)

where $N'(x) = (1/\sqrt{2\pi})e^{-\frac{1}{2}x^2}$ is the normal density function. For a European put option on a stock that does not pay dividends, the theta is given by

$$\Theta = -\frac{SN'(d_1)\sigma}{2\sqrt{t}} + Xe^{-rt}N(-d_2)$$
(12.8)

Figure 12.8 shows the graph of theta as a function of the asset price. This looks like the inverse of the graphs for gamma; in other words, the theta behaves very much like the negative of the gamma. We see that the theta is usually negative, reflecting the fact that options are usually worth more when they have a longer period left to maturity. We have already seen in Chapter 12 that it is never optimal to exercise an American call option on a non-dividend paying stock prior to maturity.

Therefore, the theta of the European call option is always negative. In the case of a put option, it may be optimal to exercise prior to maturity, if the option is deep in the money. This is seen in the slightly

The Greeks of the Black-Scholes Model | 12.9

positive theta for such options. However, if the American put options are either out of the money or only moderately in the money, it is not optimal to exercise them prior to maturity. This is why the theta of the corresponding European put options is also negative.



Figure 12.8 Theta of at-the-money calls and puts are highly negative, implying large time decay of option values. The thetas fall away but remain negative zero if the options are deep out of the money. Deep in the money calls also have nearly zero theta. A deep in the money put can have a positive theta because premature exercise is optimal. The value of a deep in the money European put rises with every passing day as the exercise moves one day closer

If the underlying asset pays dividends at the rate q, the theta of the call option is given by

$$\Theta = qSN(d_1)e^{-qt} - \frac{Se^{-qt}N'(d_1)\sigma}{2\sqrt{t}} - rXe^{-rt}N(d_2)$$
(12.9)

while the theta of the put option is given by

$$\Theta = -qSN(-d_1)e^{-qt} - \frac{Se^{-qt}N'(d_1)\sigma}{2\sqrt{t}} + rXe^{-rt}N(-d_2)$$
(12.10)

The presence of the dividend yield introduces a positive term in the theta of the call option. As we saw in Chapter 10, if the asset pays a sufficiently large dividend, it may be optimal to exercise an American call option before maturity. The theta of the corresponding European call option can therefore be positive. On the other hand, the presence of dividends make early exercise less likely for American puts and this is reflected in the additional negative term in the theta of the European put.

We now turn to the behaviour of the theta as a function of time to maturity. Figure 12.9 and Figure 12.10 show this for call and put options respectively. Once again, these look like the inverse of the graphs for gamma; in other words, the theta behaves very much like the negative of the gamma. The theta of the at-the-money option explodes to a large negative value as it approaches maturity. The large positive gamma of such options means that they make a lot of money for their holders if the stock fluctuates in either direction. At the same time, the large negative theta means that the option loses a lot of value if the stock does not fluctuate. As usual, there is no free lunch.

12.10 | Derivatives and Risk Management



Time to maturity

Figure 12.9 Theta of an at-the-money call option is highly negative close to maturity. The theta moves closer to zero at long maturity. Out of the money options have zero theta close to maturity, but the thetas are negative at longer maturities.



Time to maturity

Figure 12.10 Theta of an at-the-money put option is highly negative close to maturity. The theta moves closer to zero at long maturity. In the money puts have near zero theta close to maturity, but the thetas are negative at longer maturities. Deep out of the money puts have positive thetas close to maturity, but the thetas fall away to zero at longer maturities.

The Greeks of the Black-Scholes Model | 12.11

12.4 THE BLACK-SCHOLES EQUATION

The three Greeks that we have discussed – delta, gamma and theta – are all related. Let us first consider the situation where the delta is zero (for example, the option has been delta hedged to eliminate the delta risk. The relationship in this case is:

$$\Theta + \frac{1}{2}\sigma^2 S^2 \Gamma = rf \qquad (12.11)$$

To understand this equation let us consider a short time interval h and see what happens to the option over this period. For this purpose, we multiply both sides of the above equation by h to get:

$$\Im h + \frac{1}{2}\sigma^2 h S^2 \Gamma = rhf \qquad (12.12)$$

The right hand side is the expected return from holding the option. In the risk-neutral world, the option must earn the risk-free rate of return and this what we see on the right hand side *rhf*. Over the period *h*, notice what happens if the stock price does not change? In this case, the change in value is due only to passage of time and this change is given by Θh which is the first term on the left hand side. Now consider the effect due to fluctuations in the stock price. During the discussion of gamma, we saw that the gain due to this is half the gamma times the square of the asset price times the realized volatility of the underlying asset during that short period. The market's expectation about what the realized volatility will be, is given by the implied volatility σ^2 . Since σ^2 is the annualized volatility, we know that the volatility over a short period *h* is expected to be $\sigma^2 h$. Multiplying this by half the gamma and then by the

square of the asset price, we get the second term on the left hand side $\frac{1}{2}\sigma^2 h S^2 \Gamma$.

Adding these two terms together (the change due to efflux of time without change in the stock price and the change due to fluctuation in the stock price) gives us the total expected return from holding the option. This must under risk-neutral valuation be equal to the risk- free return over the period h, namely the right hand side *rhf*.

We now turn to the general case where the delta is non zero. Now there will be an additional component equal to the delta times the expected change in the stock price. Under risk- neutral valuation, the stock also earns the risk-free return and therefore the expected change in the stock price over the small interval h is rhS. Multiplying this by the delta of the option gives us the change in the option price due to the expected change in the asset price: rhS Δ . Adding this term to the left hand side of our earlier equation gives us the following equation:

$$\Theta h + rhs \Delta + \frac{1}{2} \sigma^2 h S^2 \Gamma = rhf \qquad (12.13)$$

We can divide this whole equation throughout by h to get the Black-Scholes differential equation³:

³ Mathematically-oriented readers can think of the second and third terms on the left hand side $\left(r S \Delta + \frac{1}{2} \sigma^2 S^2 \Gamma\right)$ as a second order Taylor series expansion of the option price as a function of *S*. This also explains why the realized gains of a Δ hedged option is given by $\frac{1}{2} \sigma^2 S^2 \Gamma$ as mentioned earlier in this chapter. Similarly, the first term (Θ) on the left hand side of Eq. 12.14 is a first order Taylor series expansion of the option price as a function of calendar time (-t). The other determinants of the option price (the strike price, the interest rate, and the volatility) are assumed to remain unchanged. It is beyond the scope of this book to explain why a second order expansion with respect to the asset price is required, while a first order expansion with respect to time is sufficient.

12.12 | Derivatives and Risk Management

$$\Theta + r S\Delta + \frac{1}{2} \sigma^2 S^2 \Gamma = rf \qquad (12.14)$$

When Black and Scholes derived their option pricing model over 30 years ago, the concept of riskneutral valuation had not been developed. They had to derive their model the hard way by first obtaining the above differential equation and then trying to solve it. This is not a very easy equation to solve, but Black and Scholes realized that this was essentially the same as the heat equation in physics for which the solution is known. Applying this solution and imposing the boundary condition that at expiry, the option value is equal to the intrinsic value, they were able to get the Black Scholes formula.

In our discussion of theta, we noted that the graph of theta looks very much like the graph of the negative of gamma. The Black-Scholes differential equation tells us why this is so. If we consider the simplified case where the interest rate is zero, then the Black-Scholes differential equation reduces to

 $\Theta = -\frac{1}{2}\sigma^2 S^2 \Gamma$. This exemplifies the 'no free lunch' condition in its simplest form: if the option holder

makes a lot of money from fluctuations in the asset price (large positive gamma), he takes a large loss if the asset price remains unchanged (large negative gamma). If the interest rate is non-zero, the relationship is more complex, but over short time intervals the effect of the interest rate is quite small and the qualitative picture (the shape of the graphs) is still quite similar.

12.5 VEGA

Our discussion of gamma focussed on the realized volatility over a holding period. This is adequate if the option is held to maturity. But if an option is sold prior to maturity, the gamma only captures the effect of the realized volatility till the time of sale. If for example, the realized volatility has been high till the time of sale, there would be a realized gain during this period. But it is possible, that this higher realized volatility might lead the market to believe that the realized volatility will be higher in future also. In other words, the implied volatility used by the market to price options might itself go up.

The term vega is used to denote the change in option value if the implied volatility changes. Vega is not a letter of the Greek alphabet unlike delta, gamma, and theta. It is however referred to as a Greek of the option. It is denoted by the capital letter V written in a calligraphic font as V. Mathematically, vega is

defined as $\mathcal{V} = \frac{\partial f}{\partial \sigma}$ where f is the option price and σ is the volatility.

For a European call or put option on a non dividend paying stock, the vega is given by

$$\mathcal{V} = SN'(d_1)\sqrt{t} \tag{12.15}$$

If the stock pays dividends at the rate q, the vega is given by:

$$\mathcal{V} = S \, \mathrm{e}^{-qt} N'(d_1) \, t \tag{12.16}$$

Comparison with equations for gamma show that they both involve the normal density $N'(d_1)$, but that this density is multiplied or divided differently by terms involving the asset price and the implied volatility. If different options on the same underlying (same asset price and same implied volatility) are considered, the gammas and vegas of all these options would be proportional to each other.

For this reason, the graph of the vega as a function of the asset price looks very similar to the graph of the gamma as may be seen from Figure 12.11.

The Greeks of the Black-Scholes Model | 12.13



Figure 12.11 Vega of call and put options are highest when the underlying price is close to the strike price. The vega drops off to zero when the option is deep in or out of the money

When viewed as a function of time to maturity, the graph of the vega is very different from that of gamma as may be seen from Figure 12.12. The vega of any option goes to zero as the option reaches maturity. This is because the change in the implied volatility impacts the future fluctuation of the asset price. As the option approaches maturity, there is practically no future left before expiry and the change in the implied volatility makes no difference at all. By the same logic, an option that has a long time to maturity is strongly impacted by changes in the vega because there is a very long future over which the changed volatility can act.



Figure 12.12 All options have near zero vega close to maturity because the future volatility has little impact when the future is very short. The vega rises with maturity over most practically relevant ranges of maturity. However, at very long maturities, the vega starts falling because the effect of the interest rate is to make all calls deep in the money, and all puts deep out of the money and in both cases, the vega falls towards zero

12.14 | Derivatives and Risk Management

Why then does the vega start declining when the maturity is very long? This happens for the same reason that the delta of long maturity calls goes to 1 for very long maturity. The effect of the interest rate, at sufficiently long maturities causes the calls to become deep in the money. At this stage, the volatility has little impact because the probability that the option would move out of the money is quite small. Similarly, very long maturity puts are deep out of the money (delta approaching 0) and the volatility is immaterial because of the low probability that any fluctuations could drive the option into the money.

12.6 RHO

Rho is the change in the option value when the risk-free interest rate changes.

Mathematically, it is defined by $\rho = \frac{\partial f}{\partial r}$ where f is the option price and r is the risk free rate. For a call option, the rho is given by:

$$\rho = tX \,\mathrm{e}^{-rt} N\left(d_2\right) \tag{12.17}$$

while for a put option, the rho is:

$$\rho = -tX \,\mathrm{e}^{-rt} N \,(-d_2) \tag{12.18}$$

These equations tells us that the rho is simply the change in the discounted exercise price multiplied by the risk-neutral probability that the option would be exercised. In these formulas, we have not distinguished between assets that pay dividends and those that do not. The same formula can be applied to both cases. However, it must be noted that if the asset pays dividend, then the definition of d_2 itself is different.

When the asset is a foreign currency, the 'dividend yield' is the foreign interest rate. Just as we compute a rho with respect to the domestic interest rate, we can compute a rho with respect to the foreign interest rate.

These rhos are equal to $\rho_f = -tSN(d_1)e^{-qt}$ for calls and $\rho_f = tSN(-d_1)e^{-qt}$ for puts.

Another important situation is that of options on futures. In this case q=r and a change in r causes an equal change in q. The impact on the option price is the combined effect of the change in r and q. The formula in this case is simplicity itself: $\rho = -tf$ where f is the option price. This result might appear strange but it becomes obvious when one recalls that the formula for the option on futures can be obtained by first taking the Black-Scholes formula with the interest rate set to zero and then discounting the resulting option value at the risk- free rate. This being so, changing the risk-free rate only changes the discounting of the option value and the result is evident.

12.7 NUMERICAL EXAMPLE OF OPTION GREEKS

We illustrate the computation of the option Greeks by considering four different examples. We consider two assets—a stock that does not pay dividends and a foreign currency. On each of these two assets, we consider calls and puts. This gives us four options.

The specific options to be considered are the following, where the interest rates are continuously compounded.
Option Greeks Examples: Option Data						
Underlying Asset	Stock (no d	lividends)	Curi	rency		
Option Type	Call	Put	Call	Put		
Asset Price	450	750	44	44		
Exercise Price	450	725	43.5	44		
Risk free rate	5%	5%	5%	5%		
Volatility	35%	25%	6%	6%		
Maturity (months)	3	1	6	6		
Maturity (years)	0.25	0.08	0.50	0.50		
Dividend Yield	0	0	3%	3%		

The Greeks of the Black-Scholes Model | 12.15

A few terms like $N(d_1)$ and $\sigma\sqrt{t}$ appear in the formulas of several of the Greeks. We now compute these intermediate quantities so that the Greeks themselves can be computed quite easily.

Option Greeks Examples: Intermediate Quantities						
Underlying Asset	Stock (no	Stock (no dividends)				
Option Type	Call	Put	Call	Put		
d_1	0.1589	0.5636	0.5263	0.2569		
d_2	-0.0161	0.4914	0.4839	0.2145		
$N(d_1)$	0.5631	0.7135	0.7007	0.6014		
$N(d_2)$	0.4936	0.6884	0.6858	0.5849		
$N(-d_1)$	0.4369	0.2865	0.2993	0.3986		
$N(-d_2)$	0.5064	0.3116	0.3142	0.4151		
$N'(d_1)$	0.3939	0.3404	0.3473	0.3860		
$\sigma\sqrt{t}$	0.1750	0.0722	0.0424	0.0424		
$X e^{-rt}$	444.4100	721.9855	42.4260	42.9136		

Given the above intermediate quantities, it is a simple matter to compute the various Greeks by simple multiplication or division.

Option Greeks Examples: The Greeks						
Underlying Asset	Stock (n	o dividends)	C	Currency		
Option Type	Call	Put	Call	Put		
Delta (Δ)	0.5631	-0.2865	-0.2949	0.5924		
Gamma (Γ)	0.0050	0.0063	0.1833	0.2037		
Theta (Θ)	-73.0127	-99.2883	-0.3614	-1.1829		
Theta(Θ) per day	-0.2000	-0.2720	-0.0010	-0.0032		
Vega (V)	88.6355	73.6905	10.6460	11.8304		
Rho (ρ)	54.8389	-18.7458	-6.6660	12.5505		
Foreign Rho (ρf)			6.4875	-13.0333		

The only new feature is the row for theta per day. In all the Black-Scholes formulas, we usually measure time in years. This is inconvenient when it comes to theta because we are typically interested in

12.16 | Derivatives and Risk Management

the decay of option value when one day elapses. Therefore, practitioners typically divide the theta computed from the formula by 365 to arrive at the theta per day.

Chapter Summary

Delta (Δ) is the sensitivity of the option price to changes in the price of the underlying. It is the key parameter that is used in dynamic hedging, but is otherwise not regarded as a critical risk factor because it is easy to eliminate by positions in the spot or futures market for the underlying. Gamma (Γ) is the sensitivity of the delta to changes in the asset price. Vega (V) is the sensitivity of the option price with respect to changes in its implied volatility. Gamma and vega are always positive for long option positions and negative for short option positions. They are in many ways, the most important measures characterizing an option position.

Theta (Θ) measures the rate of time decay of an option. It is the rate at which the option price changes through the passage of time if nothing else changes. Theta is usually negative for long option positions and positive for short option positions. Rho (ρ) is the sensitivity of the option price with respect to the interest rate and is in many ways the least interesting of all the option Greeks.

The formulas for the various Greeks in the Black-Scholes model for non-dividend paying assets are as follows:

Call on Non-Dividend Paying Asset	Put on Non-Dividend Paying Asset
$\Delta = N(d_1)$	$\Delta = -N(-d_1) = (N(d_1) - 1)$
$\Gamma = \frac{N'(d_1)}{S\sigma\sqrt{t}}$	$\Gamma = \frac{N'(d_1)}{S\sigma\sqrt{t}}$
$\Theta = -\frac{SN'(d_1)\sigma}{2\sqrt{t}} - rXe^{-rt} N(d_2)$	$\Theta = -\frac{SN'(d_1)\sigma}{2\sqrt{t}} + rXe^{-rt} N(-d_2)$
$\mathcal{V} = SN'(d_1)\sqrt{t}$	$\mathcal{V} = SN'(d_1)\sqrt{t}$
$\rho = tXe^{-rt}N(d_2)$	$\rho = -tXe^{-rt}N(-d_2)$

These formulas need to be modified in case of assets paying continuous dividends at the rate q. Moreover, the definition of rho needs some care. For currency options, q is the foreign interest rate and there is a rho denoted ρ_f with respect to this interest rate in addition to the usual ρ with respect to the domestic interest rate. For options on futures, r and q are one and the same and there is a separate formula for rho in this case. The formulas for the various Greeks in the Black-Scholes model for assets paying continuous dividends at the rate q are as follows:

Call on Asset Paying Dividend q	Put on Asset Paying Dividend q			
$\Delta = N(d_1) \mathrm{e}^{-qt}$	$\Delta = -N(-d_1)e^{-qt} = (N(d_1) - 1)e^{-qt}$			
$\Gamma = \frac{e^{-qt} N'(d_1)}{S \sigma \sqrt{t}}$	$\Gamma = \frac{e^{-qt} N'(d_1)}{S \sigma \sqrt{t}}$			
$\Theta = qSN(d_1)e^{-qt} - \frac{Se^{-qt}N'(d_1)\sigma}{2\sqrt{t}} - rXe^{-qt}N(d_2)$	$\Theta = -qSN(-d_1)e^{-qt} - \frac{Se^{-qt}N'(d_1)\sigma}{2\sqrt{t}}$			
-	$+ rXe^{-rt}N(-d_2)$			
$\mathcal{V} = S \mathrm{e}^{-qt} N'(d_1) \sqrt{t}$	$\mathcal{V} = S \mathrm{e}^{-qt} N'(d_1) \sqrt{t}$			
$\rho = tX e^{-rt} N(d_2)$	$\rho = -tX e^{-rt} N(-d_2)$			
$\rho_f = -tSN(d_1)e^{-qt}$	$\rho_f = t SN \left(-d_1\right) \mathrm{e}^{-qt}$			
Put or Call on Futures $\rho = -tf$				

The Greeks of the Black-Scholes Model | 12.17

Suggestions for Further Reading

Baird, Allen. Jan (1993) "Option Risks", Chapter 3 in *Option Market Making –Trading and Risk Analysis for the Financial and Commodity Option Markets*, New York, Wiley.

Problems and Questions

- 1. If the current stock price is 150, the volatility is 21% and the risk free rate is 7% continuously compounded, what is the delta, gamma and vega of a three-month call option with an exercise price of 160?
- 2. In problem 1, what is the delta of a three-month put option with the same exercise price?
- 3. If the current stock price is 130, the volatility is 30%, the dividend yield is 2% and the risk free rate is 5% continuously compounded, what is the delta, gamma and vega of a three-month call option with an exercise price of 140?
- 4. In the previous problem, what are the theta and raho of the option?
- 5. If the Nifty index is quoting at 4500, the volatility is 14% the risk free rate is 6.25% continuously compounded, the dividend yield on the index is 1%, then what is the theta and rho of a one-month call option with a strike of 4650?
- 6. What is the rho and foreign rho of a three month call option on the rupee-dollar rate at an exercise price of 41 when the sport rate is 40, the Indian interest rate is 6.5%, the US interest rate is 5% and the volatility of the exchange rate is 11%? Assume that both interest rates are on simple interest basis.
- 7. If one-month gold futures are trading at \$625 what is the rho of a one-month put option with a strike of \$615 assuming a volatility of 10% and a risk free interest rate of 4.75% continuously compounded?
- 8. In the previous problem, what are the vega and theta of the option?

Chapter **Thirteen**

Complex Option Strategies

Objectives

Chapter 11 discussed how option strategies can be used to take a bullish or bearish view on an asset, or to hedge an existing exposure. But options are much more versatile instruments that can be used to take more nuanced views on the underlying or its volatility. This chapter looks at some of these strategies. These strategies require the use of two or more options. While the option delta provides a characterization of the bullish or bearish stance of an option strategy, it is not sufficient to describe the exposure of the position to volatility and other risks of these more complex option strategies. The gamma, vega and other Greeks of the position are therefore very important and this chapter discusses these.

13.1 STRADDLES

A straddle consists of buying a put and a call at the same exercise price. Typically, the exercise price is close to the market price at the inception of the trade. An at-the-money straddle has a very low delta because the negative delta of the put would largely offset the positive delta of the call.¹ The straddle is then practically a pure play on the volatility of the underlying.

The profit diagram of the straddle is shown in Figure 13.1.

If the price of the underlying rises substantially, the put option becomes worthless but the call option becomes highly valuable. Conversely, if the price of the underlying declines substantially, the call option becomes worthless but the put option becomes highly valuable. The straddle, therefore, makes money when the underlying makes a large move regardless of whether the move is up or down. However, if the underlying does not move much, neither the call nor the put would have much value and a major part of the option premium would be lost. That is why we say that the straddle is not a directional bet, but a bet on volatility.

For example, an investor may have reason to believe that the stock would make a large move when the company announces its earnings. If the company meets its earnings target, it might rise significantly, but if it falls badly short, the stock might tank. A straddle would be a perfect way to take advantage of such a situation. The risk of course is that the earnings are a little below target, so that the stock does not move much. In this case, the premium paid to set up the straddle would be lost without any significant gain.

It is also important to recognize that the buyer of the straddle is not betting that the realized volatility would be high is an absolute sense. The true bet is that the volatility would be more than what the market is expecting it to be. In other words, the buyer of a straddle is actually betting that the realized volatility

¹ The delta would not be zero because the long call has an unlimited upside while the long put has a limited upside (the asset price cannot go below zero). For this reason, the negative delta of the put would not completely offset the positive delta of the call and the at-the-money straddle would have a slightly positive delta. Moreover, for the delta to be close to zero, the strike must be close to the forward price of the asset and not to its spot price because it is the forward price that reflects the (risk-neutral) expected price of the asset at maturity.

13.2 | Derivatives and Risk Management



Figure 13.1 Profit Diagram of a long straddle with strike of 100

would exceed the implied volatility. For example, if the investor believes that stock A would have a realized volatility of 60% while stock B would have a realized volatility of only 30%. In the market, options on A are trading at an implied close to 60% while options on B are trading at an implied of 20%, the investor would prefer to buy a straddle on B though it is less volatile than A. This is because the high volatility of A is already reflected in the option prices and the straddle on A would cost so much money in option premium that there is little chance of a profit. Stock B on the other hand, is less volatile but the option prices are based on an even lower implied volatility and this gives rise to a profit opportunity.

Another point to remember is that during the life of the straddle, if the price of the underlying moves away from the exercise price of the options, then the straddle would acquire the character of a directional bet. For example, if in Figure 13.1, the underlying moves to 120 well before expiry, then the straddle becomes a bullish position. The delta of the call options would be significantly more than 0.5 as these options would now be deep in the money, while the put options would have deltas well below 0.5, as these options would now be deep out of the money. As a result, the straddle itself would have a large positive delta. A person who bought a straddle purely as a bet on volatility would probably close out the old straddle and create a new straddle close to the prevailing price of 120. Alternatively, she may neutralize the positive delta of the position by shorting an appropriate number of futures.

The short straddle is a position set up when the investor believes that the asset would move in a narrow range. It is a bet that the realized volatility would be lower than the implied volatility. Its profit diagram (Figure 13.2) is a mirror image of the diagram for the long straddle.

13.2 STRANGLES

A straddle can be very expensive because it involves buying two options, both of which are at the money. If the investor is truly confident of a large move in the asset price, she can reduce the option premium significantly by buying options that are slightly out of the money. When the asset is trading at 100, she can buy a call at 105 and buy a put at 95. If the asset rises above 105, the call becomes valuable. If the prices

fall below 95, the put becomes valuable. If the asset trades between 95 and 105, both options expire worthless. As compared to a straddle, the price has to move by a much larger amount for the position to make money. But otherwise, the position is similar to a straddle. The strangle is also essentially a play on the volatility of the underlying.



Figure 13.2 Profit diagram of a short straddle with a strike of 100

A strangle centred around the current asset price (or more precisely the forward price) has a delta close² to zero because the positive delta of the call (at 105) almost cancels out the negative delta of the put (at 95). The profit diagram of a long strangle is shown in Figure 13.3 while the profit diagram of a short strangle is shown in Figure 13.4.



Figure 13.3 Profit diagram of a long strangle consisting of a long call with a stirke of 120 and long put with a strike of 80

² As in the case of straddle, the delta is not exactly zero because of the unlimited upside of the call and limited upside of the put.

13.4 | Derivatives and Risk Management



Figure 13.4 Profit diagram of a short strangle consisting of a short call with a strike of 120 and short put with a strike of 80

13.3 BUTTERFLY SPREADS

The butterfly spread is a way of taking a very focused and localized view of what the asset price will be. Suppose for example that an investor believes that the fair value of the asset is 100 and would, therefore, like to make a simple bet that the asset price at expiry will be between say 90 and 110. Ideally, he would like an option strategy that makes a fixed amount of money if his view is right and loses a fixed amount of money if his view is right and loses a fixed amount of money if his view is wrong. At first sight, a short straddle at 100 also expresses the view that the asset price will trade close to 100 but it focuses on volatility rather than on the asset value. The straddle has unlimited losses if the asset trades far away from 100. The simple bet that the investor wants has limited losses.

The straddle can be modified to achieve this goal by adding a long strangle which caps the unlimited losses of the short straddle. The unlimited loss in the straddle arising from the short call at 100 can be capped by buying a call at say 120 and similarly, the loss from the short put can be capped by buying a put at say 80. The short straddle at 100 plus a long strangle (at 80 and 120) is basically a butterfly and has characteristics similar to the simple bet that our investor is looking for. Using put-call parity, it is possible to express the butterfly in a simpler fashion using three options rather than four. The two put options can be converted into calls³. to arrive at a butterfly that consists of shorting two calls at 100

³ To convert the short put at 100 into a short call at 100, one needs to add a long futures position and also short bonds with a maturity value of 100. To convert the long put at 120 into a long call at 120, we need to add a short futures position and also buy bonds with a maturity value of 120. The two futures positions cancel out leaving no net futures position to be added. The two bond positions net out to a long bond with a maturity value of 20. The initial investment, therefore, changes by the present value of 20. This difference is reflected in the fact that the butterfly constructed using a short straddle and a long strangle earns a net premium while the butterfly constructed using only calls costs a premium. The butterfly using only puts has identical initial investment and identical payoffs as the one using only calls.

Complex Option Strategies | 13.5



Figure 13.5 Payoff and profit diagram for a long butterfly consisting of two short calls at 100 and a long call each at 80 and 120. It makes money when the underlying trades in a narrow range of around 90–110 and loses money outside the range. The maximum loss is limited to the net premium paid for the position and the maximum profit is also limited to the difference between the low and middle strikes (80 and 100)

and buying calls at 80 and 120. The payoff and profit diagram would be as shown in Figure 13.5. The two calls can be equally well converted into puts to arrive at a butterfly that consists of shorting two puts at 100 and buying puts at 80 and 120. The profit and payoff diagram would be identical.

To understand the payoff diagram of the butterfly, consider the situation where the position is constructed using only call options. If the underlying trades below 80, all calls expire worthless and the payoff is zero. If the underlying trades above 120, all options expire in the money. The holder buys one unit of the asset at 80 and another at 120. The average price that she pays is 100. The two short calls at 10 require her to sell the two units at a price of 100. Thus her average payoff is again zero. So outside the range of 80-120, her payoff is zero. Between 80 and 100, only the long call at 80 is in the money and the payoff of the butterfly looks like that of a long call. At 100, she gets the peak payoff of 20 from buying at 80 and selling at 100. Between 100 and 120, she buys one unit at 80 and sells two units at 100. Her position is now that of a short call and because of this net short call, the payoff declines from 20 as the underlying rises above 100. At 120, the payoff becomes zero and above 120, when all calls are in the money, the payoff remains zero.

To understand the profit diagram, we need to convince ourselves that the butterfly costs money to set up. The simplest way to see this is to observe that the payoff is either zero or positive. It is never negative. Under risk-neutral valuation, therefore, the value of the position must always be positive. We can in fact go further and relate the butterfly price to the risk-neutral probability that the underlying would expire within a range. In our example, the payoff is zero outside the range of 80–120 and varies between 0 and 20 within this range.

Let us take the average payoff within the range of 80-120 as 10. Under risk-neutral valuation, the value of the butterfly must therefore be approximately equal to the present value of 10 times the risk-neutral probability that the underlying trades between 80 and 120. This approximation

13.6 | Derivatives and Risk Management

becomes very good as the strike prices of the butterfly come closer together. In fact, this idea can be used to recover the entire risk-neutral distribution of asset prices from traded option prices. This idea will be need in Chapter 15 when we relate the volatility smile to the risk-neutral distribution of asset prices.

When we consider the net option premium paid to set up the butterfly, the zone of positive profit of the butterfly shrinks. In Figure 13.5, the option premium paid is a little less than 10 and the zone of positive profit is about 90-110. Below about 90, the positive payoff of the long call at 90 is insufficient to recover the option premium. At 100, the profit is equal to a little over 10 (20 less the option premium). Above around 110, this profit is wiped out by the negative payoffs from the two short calls.

Several butterflies can also be combined to express much more complex views on asset prices. Consider for example, a film studio that is about to release a potential blockbuster movie. An investor believes that if the movie becomes a blockbuster, the share price of the studio company would rise to about 150, but if the film flops, it would fall to 60. He could express this view by buying a butterfly centred on 150 and another butterfly centred on 60. If he believes that there is a two-thirds chance that the film will flop, he may buy two butterflies centred on 60 and only one butterfly centred on 150. By using more and more butterflies, an investor can express any view that he may have about the risk-neutral distribution of asset prices and achieve any pattern of payoffs that he may desire.

13.4 GREEKS OF OPTION STRATEGIES

The option strategies in this chapter are complex positions that provide exposure to the volatility and to the localized regions of the risk-neutral distribution of asset prices. By their nature these tend to be dynamic positions that are modified in line with the investor's views as well as the market movements. For example, a long straddle may be established when the market is perceived to be underpricing volatility. Such an underpricing is not, however, likely to last long and if the underpricing is corrected within a few days, the investor would unwind the straddle at a profit. These positions can, therefore, often be trading positions which have a short investment horizon. While the payoff diagram provides an understanding of what happens to the position when it is held to maturity, it is an inadequate description of what could happen to a position which may be unwound or modified well before maturity.

To understand the short-term behaviour of these strategies, it is most useful to look at the option Greeks. These tell us what would happen during a short period of time as the underlying moves or its implied volatility changes, or if nothing happens (mere efflux of time). Delta and gamma tell us the impact of changes in the underlying, vega measures the impact of changes in the implied volatility and theta indicates the effect of passage of time.

We begin with the delta depicted in Figure 13.6 as a function of the underlying. All the three strategies that we have considered so far (straddle, strangle and butterfly) have near zero delta when the current asset price is close to the central strike price. As the asset price declines, both the long straddle and the long strangle move towards a delta of -1. This is because the out of the money call becomes increasingly irrelevant and the position becomes essentially an in-the-money put. This change is slower for the strangle because the put is out of the money to begin with and it takes a significant decline in asset prices to bring that into the money. If the asset price rises, the delta rises to +1 because the position is now dominated by a long call which is in-the-money. Again the strangle delta responds more slowly than the straddle delta.

Complex Option Strategies | 13.7



----- Stråddle ----- Strångle ----- Butterfly

Figure 13.6 Delta of different option strategies as a function of the underlying. The long straddle consists of a call and a put with a strike of 100. The long strangle consists of a call at 120 and a put at 80. The long butterfly consists of two short calls at 100 and a long call each at 80 and 120. The deltas of all three strategies are close to zero at the central strike price of 100. As the asset price declines, both the long straddle and the long strangle are dominated by a long put and move towards a delta of -1. Conversely, when the asset price rises, they are dominated by a long call and move towards a delta of +1. The butterfly behaves like a short straddle for modest moves in the price. But beyond a point, it reverses course and starts behaving like a long strangle and the delta reverts back to zero

The butterfly behaves differently. It has already been seen that the butterfly is akin to a short straddle protected by a long strangle. For modest changes in the asset price, therefore, its delta behaves like a short straddle. So the picture is opposite to that of the long straddle. Declines in asset prices lead to a positive and rising delta and rises in asset prices cause a negative delta. But as the asset price moves close to the 'wings' of the butterfly where the long strangle comes into play, the picture reverses. Sharp drops in the asset price cause the positive delta to decline towards zero. Similarly, sharp rises in the asset price cause the negative delta to rise towards zero. The graph of the delta of the butterfly looks like a wave that alternately rises and falls.

The gamma and vega of these option strategies are depicted in Figure 13.7 and Figure 13.8. These two graphs are very similar in shape. The gamma and vega of the straddle have a sharp peak at the centre and fall away sharply to zero. The graph for the strangle is flatter and more spread out as the gamma and vega decline more slowly. The gamma/vega of the strangle moves close to zero only when the asset price moves far away from both the strike prices. The butterfly again has a wave-like graph of the butterfly almost coincides with that of the strangle, while close to the centre it is almost a mirror image of that of the long straddle.

13.8 | Derivatives and Risk Management



Figure 13.7 Gamma of straddle, strangle and butterfly. These positions are the same as in Figure 13.6. The gamma of the straddle is highest at the centre and falls away quickly to zero. The gamma of the strangle is flatter and falls away to zero more slowly. The butterfly again has a wave like graph combining features of the short straddle and long strangle



Figure 13.8 Vega of straddle, strangle and butterfly. These positions are the same as in Figure 13.6. The shapes of these graphs is very similar to that of the gamma depicted in Figure 13.7

Gamma and vega measure the risk to the position from the volatility of the underlying. Gamma is the effect of the realized volatility, while vega is the risk due to expected future volatility as measured by the

Complex Option Strategies | 13.9

implied volatility. The long straddle and strangle are long volatility—they are essentially bets that the volatility will be high. The butterfly on the other hand is short volatility—it is a bet that the asset will not move much. High volatility, therefore, produces profits for the straddle and strangle as evidenced by their positive gamma and vega. High volatility produces losses for the butterfly when the asset price is close to the centre where it behaves like a short straddle and has a negative gamma and vega. Far from the centre, the butterfly behaves like a long strangle (positive gamma and vega) and high volatility produces profits.

Measurement of risk is not complete without asking, what happens if the asset prices do not move much. In this case, the change in the value of the position is dominated by the theta or time decay risk which is depicted in Figure 13.9. Time decay risk is typically the mirror image of volatility risk. Positions which bear a lot of losses when asset prices fluctuate (negative gamma and vega) tend to have gains when asset prices do not change much (positive theta). Conversely, positive gamma and vega is usually associated with negative theta. This is what is seen in Figure 13.9.



Figure 13.9 Theta of straddle, strangle and butterfly. These positions are the same as in Figure 13.6. The shapes of these graphs are very similar to mirror images of the gamma depicted in Figure 13.7 or the vega depicted in Figure 13.8

13.5 CALENDAR SPREADS

A calendar spread involves offsetting positions in contracts of differing maturities. For example, an investor may be short a one-month call and long a three-month call. This produces a profit diagram which resembles the profit diagram of a butterfly. In this profit diagram, the horizontal axis is the price of the underlying at the expiry of the near month contract. At this point, the far month contract is still alive and is valued using the Black-Scholes formula. It is instructive to analyse the reason for the butterfly-like behaviour.

First of all, it is easy to see that the worst possible loss on the spread is limited to the premium paid to set up the spread. (Evidently, at inception, the far month option would be more expensive than the near month option at the same strike. The spread does therefore cost money to set up.) At the expiry of the near month option, its payoff is equal to the intrinsic value. The value of the far month call at the same time is

13.10 | Derivatives and Risk Management

at least equal to the intrinsic value. Thus the pay-off at the end of the near month cannot be negative. The worst that can happen is the loss of the net premium paid. It is also clear that this worst loss arises when the underlying is either deep in the money or deep out of the money. The spread should do better when the underlying closes at or near the money.

If the underlying closes at the money, the near month option expires worthless, and the premium earned by selling it becomes part of the profit. The far month contract also loses value because the underlying is unchanged and the maturity of the option has come down by a month. This loss is however, much lower than the gain from the near month option. We saw in Chapter 12 that an at-the-money option has a very large negative theta close to expiry and, therefore, loses most of its value in this period. For example, if the underlying remains unchanged, a four-month option loses as much value in the last month as it loses in the first three months put together. Thus the calendar spread gains much more from the near month short call expiring worthless than it loses from the fall in value of the far month long call. The profit diagram (Figure 13.10) shows that the calendar spread makes a substantial profit when the underlying closes at the money.



Figure 13.10 Profit diagram of calendar spread at expiry of near month option. The calendar spread consists of a short position in a one month call option with a strike of 100 and a long position in a three month call at the same strike. The profit at near month expiry is computed by considering the Black-Scholes value of the far month option, the payoff of the near month option and the net premium paid to set up the spread. The profit diagram resembles that of the butterfly.

Now consider the impact of a fall in the underlying below the strike price. As the underlying declines below the strike price, the near month short option is completely unaffected—it simply remains worthless. The far month-long option, however, loses value as it has a positive delta. The profit of the spread thus declines as the underlying falls below the strike. The loss is not unlimited because the worst that can happen is that the far month option also becomes worthless. At this point, the net premium paid to set up the calendar spread is lost.

Complex Option Strategies | 13.11

Finally, we examine the case where the near month option expires in the money. As the underlying rises above the strike, it causes an equal loss in the short option because at expiry, an in-the-money call has a delta equal to unity. The far month-long option gains value in this situation but since its delta is less than one, the gain in the value of this option is less than the rise in the underlying price. The result is that the calendar spread suffers as the underlying rises above the strike price. However, if the underlying rises to very high levels, the delta of the far month option also approaches unity and the calendar spread is practically unaffected by further rises in the price of the underlying. The profit diagram of the calendar spread, therefore, becomes flat beyond a point.

More complex calendar spreads can also be considered where the near month and far month positions are themselves complex option strategies and not single options. An investor could be short a one-month straddle and long a three-month straddle. The one month short straddle is long theta and short vega, while the three-month long straddle is short theta and long vega. The far month option has a smaller theta and a higher vega than the near month option. This calendar spread can therefore be both positive vega and positive theta while having a small negative gamma and a near-zero delta. Apart from the small negative gamma, the spread appears to have low risk while still earning money (positive theta). In Chapter 15, however, we will see that the implied volatility of the near month and far month options need not be the same. In fact, it is possible for the near month implied volatility to rise dramatically, while the far month implied volatility rises only slightly. In this eventuality, the spread of straddles loses money though it is positive vega. A proper analysis of these complex calendar spreads requires an examination of the risk (often called time-vega) that the implied volatilities at different maturities may change in different ways. This kind of analysis is beyond the scope of this book.

Chapter Summary

Strategy	Nature of Positions
Straddle	Long a call and long a put at the same strike typically close to the current price
Strangle	Long a call and long a put at a lower strike. The two strikes are typically on either side of the current price.
Butterfly	Two short calls (typically at a strike close to the current price), one long call at a higher strike and one long call at a lower strike. By put-call parity all calls can be replaced by puts.
Calendar Spreads	A simple example of a calendar spread consists of a near-month short call and a far- month long call. More complex spreads could have a complex option strategy in the near-month and a reverse strategy in the far-month.

The table below summarizes the strategies discussed in this chapter:

Straddles and strangles are ways of taking views on the volatility of the underlying rather than the direction in which it would move. These positions have near-zero delta but positive gamma/vega and negative theta. This implies that they make money when the underlying moves a lot in either direction, but lose money when the underlying does not move much. Short straddles and strangles have the reverse behaviour. An option strategy that combines these opposing characteristics is the butterfly which behaves like a short straddle close

13.12 | Derivatives and Risk Management

to the central strike price, while behaving like a long strangle far from the centre. It has limited downside and upside, and expresses a focused view about the underlying trading within a particular range.

Calendar spreads are often analysed by looking at their profit diagram at the expiry of near-month options. The simple calendar spread consisting of a near-month short call and a far-month long call behaves like a butterfly when looked at in this way. More complex spreads can behave in more complex ways.

Suggestions for Further Reading

Baird, Allen. Jan (1993) "Position Risk Profiles", Chapter 4 in Option Market-Making and Risk Analysis for the Financial and Commodity Option Markets, New York, Wiley.

Problems and Questions

All the problems below are based on the example of Nick Leeson discussed in Chapter 1. In mid January 1995, Leeson held a short position of 70,000 options on the Nikkei index in Japan. This position consisted almost entirely of straddles with exercise prices between 18,500 and 20,000. The Nikkei was around 19,000 and the straddles were therefore either at-the-money or close to the money. The multiplier for the Nikkei contracts at the Singapore International Monetary Exchange (IMEX) was 500 yen. The exchange rate was around ± 100 /\$ so that the notional value of Leeson's position was nearly US \$ 7 billion. The volatility of Nikkei index options was around 10%. This was significantly lower than the approximately 15% level observed in the middle of 1994.

On January 17, 1995, the powerful Kobe earthquake hit Japan and within a week, the Nikkei plummeted more than 1,500 points and the volatility shot up to nearly 25%.

In answering the following problems assume for simplicity that Leeson's position consisted of 35,000 short calls and 35,000 short puts and that all the options were at-the-money and had a one-month maturity. Assume also that the implied volatility tracked the actual volatility of the index. Assume a risk free rate of 2.25% (simple interest basis).

- 1. What was the delta, gamma and vega of the straddles just before the earthquake?
- 2. Ignoring the impact due to the change in volatility, what would have been the loss from the straddles in the aftermath of the earthquake? Compare this with the delta computed in the previous problem.
- 3. Ignoring the impact due to the change in volatility, what would have been the delta of the straddles in the aftermath of the earthquake? Compare the change in delta with the gamma computed in the previous problem.
- 4. What hedging strategy should Leeson have implemented as the Nikkei began to plunge?
- 5. Leeson's actual response was to buy Nikkei futures on a large scale in an attempt to prop up the index. Evaluate this decision.
- 6. Ignoring the impact of the drop in the index what was the loss from the straddles due solely to the rise in the volatility? Compare this with the vega computed in problem 1.
- 7. The control failures at Barings were so serious that Leeson was able to hide his straddles completely from his superiors in London. Even more surprising was the fact that headquarters did not recognize that there was a serious problem when large margin calls came in from Singapore. Estimate the magnitude of the margin calls on the straddles after the earthquake (ignoring the losses on the futures positions). How do you think Leeson would have explained his cash need to headquarters?
- 8. What difference would it have made if Leeson had used strangles instead of straddles?
- 9. Do you think Leeson should have used butterflies instead of straddles or strangles?
- 10. Would your answer to problems 8 and 9 change if Leeson bought the straddles primarily to use the premium income to cover up losses elsewhere in his trading book?

Chapter Fourteen

Volatilities and Implied Volatilities

Of the five variables that enter the Black-Scholes formula for the option price (asset price, exercise price, maturity, risk-free rate and volatility), the only variable that is not directly observed is the volatility. We need to know the volatility of the asset price during the life of the option. This chapter discusses techniques for forecasting future volatility using past volatility.

Another approach to using the Black-Scholes formula is to back out the market implied volatility from traded option prices. Under the Black-Scholes assumptions, the implied volatility of one option can be used to price other options. The next chapter relaxes these assumptions and considers the situations where different implied volatilities need to be applied to different options.

14.1 HISTORICAL VOLATILITY

The simplest way to estimate volatility is to use the basic definition that volatility is the standard deviation of (logarithmic) asset returns. Given daily or weekly asset prices, it is a simple matter to compute the corresponding daily or weekly returns and compute their standard deviation. The formula for the variance σ_t^2 and the standard deviation σ_t at time *t* using *n* days of data is given by¹.

$$\sigma_t^2 = \frac{1}{n} \sum_{k=0}^n \left(r_{t-k} - E[r] \right)^2 \quad \text{or} \quad \sigma_t = \sqrt{\frac{1}{n} \sum_{k=0}^n \left(r_{t-k} - E[r] \right)}$$
(14.1)

where r_t is the logarithmic return in day or week t and E denotes the average or expected value.

A sample size (*n*) of at least 30 returns is needed to compute the volatility with any usable degree of reliability. A sample size of 100 or more is desirable to improve the accuracy of the estimate. With daily data, 100 trading days is less than six months and there is no problem in getting this amount of data. If we use weekly data, then 100 weeks is about two years of data and questions may arise about whether the volatility has changed significantly during this period. With monthly data, 100 returns would cover over eight years and it would be hard to maintain the assumption that volatilities have not changed over such a long period. It is for this reason that most estimates of volatility are usually based on daily returns and in some cases on weekly returns, but rarely on monthly returns.

The standard deviation computed in accordance with equation (14.1) is a daily volatility if computed with daily returns and a weekly volatility if computed with weekly returns. This has to be converted into an annualized volatility since all parameters in the Black-Scholes formula are typically measured in years. The maturity of the option is usually measured in years (for example, 91 days translates into 0.25 years) and the interest rate is also normally an annual rate. To be consistent with this, the volatility must also be annualized.

To annualize the volatility, the square root scaling formula used is:

^{1.} Sometimes, the simplified formula in Eq. (14.3) is used instead of this formula.

14.2 | Derivatives and Risk Management

$$\sigma_{n \ days} = \sqrt{n} \ \sigma_{1 \ day}$$

$$\sigma_{n \ weeks} = \sqrt{n} \ \sigma_{1 \ week}$$
(14.2)

For example, if there are 255 trading days in the year, then the daily volatility must be multiplied by $\sqrt{255} = 15.97$. To convert a weekly volatility to an annual volatility, the factor to be used is $\sqrt{52} = 7.21$

Table 14.1 Computation of Nifty volatility during the second quarter of 2005. Thereare 65 return observations during this quarter and the standard deviation of these 65observations is 0.9726%. The returns are computed in logarithmic form as $\ln (P_t / P_{t-1})$.The annualized volatility is computed in the text

Date	Nifty	Return	Date	Nifty	Return
31-Mar-05	1610.40		18-May-05	1579.75	-0.38%
01-Apr-05	1637.95	1.70%	19-May-05	1586.90	0.45%
04-Apr-05	1634.55	-0.21%	20-May-05	1587.75	0.05%
05-Apr-05	1624.50	-0.62%	23-May-05	1605.65	1.12%
06-Apr-05	1637.40	0.79%	24-May-05	1617.40	0.73%
07-Apr-05	1625.85	-0.71%	25-May-05	1627.65	0.63%
08-Apr-05	1609.45	-1.01%	26-May-05	1651.45	1.45%
11-Apr-05	1590.15	-1.21%	27-May-05	1652.80	0.08%
12-Apr-05	1605.20	0.94%	30-May-05	1649.65	-0.19%
13-Apr-05	1605.25	0.00%	31-May-05	1659.00	0.57%
15-Apr-05	1549.40	-3.54%	01-Jun-05	1652.60	-0.39%
18-Apr-05	1524.40	-1.63%	02-Jun-05	1634.45	-1.10%
19-Apr-05	1509.85	-0.96%	03-Jun-05	1660.55	1.58%
20-Apr-05	1527.60	1.17%	04-Jun-05	1662.10	0.09%
21-Apr-05	1543.25	1.02%	06-Jun-05	1662.45	0.02%
22-Apr-05	1558.15	0.96%	07-Jun-05	1668.60	0.37%
25-Apr-05	1561.00	0.18%	08-Jun-05	1680.70	0.72%
26-Apr-05	1551.45	-0.61%	09-Jun-05	1675.65	-0.30%
27-Apr-05	1535.65	-1.02%	10-Jun-05	1663.75	-0.71%
28-Apr-05	1538.55	0.19%	13-Jun-05	1673.80	0.60%
29-Apr-05	1508.15	-2.00%	14-Jun-05	1676.45	0.16%
02-May-05	1526.10	1.18%	15-Jun-05	1691.30	0.88%
03-May-05	1527.50	0.09%	16-Jun-05	1685.80	-0.33%
04-May-05	1544.15	1.08%	17-Jun-05	1687.90	0.12%
05-May-05	1564.95	1.34%	20-Jun-05	1705.70	1.05%
06-May-05	1578.10	0.84%	21-Jun-05	1725.35	1.15%
09-May-05	1596.25	1.14%	22-Jun-05	1739.50	0.82%
10-May-05	1589.30	-0.44%	23-Jun-05	1738.75	-0.04%
11-May-05	1584.45	-0.31%	24-Jun-05	1745.90	0.41%
12-May-05	1596.10	0.73%	27-Jun-05	1751.45	0.32%
13-May-05	1588.50	-0.48%	28-Jun-05	1728.00	-1.35%
16-May-05	1605.70	1.08%	29-Jun-05	1744.95	0.98%
17-May-05	1585.75	-1.25%	30-Jun-05	1765.95	1.20%

Volatilities and Implied Volatilities | 14.3

Table 14.1 gives data about the Nifty index values and returns during the second quarter of 2005. During this quarter, the daily standard deviation was 0.9726 percent. There are 65 trading days in this quarter or 260 trading days in the year. The annualized volatility can be computed as:

$$0.9726\% \times \sqrt{260} = 0.9726\% \times 16.1245 = 15.68\%$$

If the Black-Scholes formula is being used on 1 July, 2005 and the volatility is to be estimated using the historical volatility for the last quarter, the annualized volatility number to be used is 15.68 percent.

A critical issue to be resolved when using historical volatility is the length of the time period to be used in the computation. In the above example, one quarter of data was used, but six months or even one year of data could have been used, reducing the sampling error by using a longer time period. This however, risks contamination from data of a more distant past when the volatility was different from what it is now. There is a trade-off here between sampling error and errors due to irrelevant data. An often used rule of the thumb is to choose a sample period equal to the maturity of the option, subject to a minimum sample size of 30 to 60 observations.

Attempts to ameliorate this trade-off have led to better methods of estimating volatility that will be discussed in later sections.

14.2 EXPONENTIALLY WEIGHTED MOVING AVERAGES

In this method, the concern about contamination from irrelevant past data is addressed by applying greater weight to more recent data and progressively lower weights to more distant data. This allows the use of much larger sample sizes without worrying about contamination.

Specifically the weights are chosen to decline exponentially: today's data has weight $1 - \lambda$, yesterday's has weight $\lambda (1 - \lambda)$, day before yesterday's has weight $\lambda^2 (1 - \lambda)$, and the data of *n* days ago has weight $\lambda^n (1 - \lambda)$. Each day has weight λ times the weight of the previous day and it can be shown that the weights add up to unity. One must always choose λ between 0 and 1 – typical values range from 0.90 to 0.99. For example, if one chooses $\lambda = 0.94$, the weights for today, yesterday, and day before yesterday are 6.00%, 5.64%, and 5.32% respectively. Each day has a weight 0.94 times the weight of the previous day. The data of 10 days ago has weight 3.23%, the data of 25 days ago has weight 1.28%, and the data of 50 days ago has weight 0.27%. The first 50 days have an aggregate weight of 95% and the first 100 days have an aggregate weight of 99.8%. Though theoretically, the sample size of this method is infinite, for all practical purposes, the sample consists only of about 50-100 days.

In practical application of this method, equation (14.1) is first simplified by assuming that E[r]=0 to give the unweighted variance and standard deviation as:

$$\sigma_t^2 = \frac{1}{n} \sum_{k=0}^n r_{t-k}^2$$
(14.3)

Applying weights to this formula and extending the summation to the infinite past gives

$$\sigma_t^2 = \sum_{k=0}^n \lambda^k (1-\lambda) r_{t-k}^2$$
(14.4)

where σ_t denotes the estimated volatility at the end of day t

When working with daily returns, it is harmless to assume that the expected return E[r] is zero. This is because over such short periods, the expected return is quite small relative to the variance of the return and makes little difference to the computed estimate.

14.4 | Derivatives and Risk Management

The computation of this exponentially weighted moving average using equation (14.4) appears very complex and laborious, but the computation becomes exceedingly simple when the volatility estimate is updated from day to day. Consider the earlier example where the weights for today, yesterday, and the day before yesterday are 6.00%, 5.64%, and 5.32% respectively. If we consider the situation tomorrow, the weights of tomorrow, today, yesterday, and the day before yesterday become 6.00%, 5.64%, 5.32%, and 4.98% respectively. It is observed that the weights of today, yesterday, and the day before yesterday become 0.94 times what they were today. In fact, every day's weight is diminished by a factor of 0.94. Since all the weights (of today and past days) added up to unity today, the weights of these same days will add up to only 0.94 tomorrow. Tomorrow's weights will still add up to unity because of the additional weight of 0.06 for tomorrow's data. It can be concluded thus, that the estimated variance at the end of tomorrow would be obtained by applying a weight of 0.06 to tomorrow's squared return and a weight of 0.94 to today's *estimated variance*. Equation (14.5) shows how this works out in general.

$$\sigma_{t+1}^{2} = \sum_{k=0}^{\infty} \lambda^{k} (1-\lambda) r_{t+1-k}^{2} = (1-\lambda) r_{t+1}^{2} + \lambda \sum_{k=0}^{\infty} \lambda^{k} (1-\lambda) r_{t-k}^{2}$$
(14.5)

$$= (1-\lambda) r_{t+1}^{2} + \lambda \sigma_{t}^{2}$$

This means that every day one simply computes a weighted average of (a) the previous day's estimated variance with weight λ and (b) the squared return that day with weight $1 - \lambda$ to get the estimated variance at the end of the day. The square root is then taken to find the estimated volatility.

For example, if the estimated volatility at the end of 15 June, 2005 was 0.8095% and the logarithmic return on 16 June, 2005 was -0.33%, the volatility will be computed at the end of 16 June, 2005 as follows using a λ of 0.94:

Previous Volatility	0.8095%
Previous Variance	$0.8095 \times 0.8095 = 0.006553$
Current day return	-0.33%
Current day squared return	0.001089%
Assumed λ	0.94
End of day estimated variance	$0.94 \times 0.006553\% + 0.06 \times 0.001089\%$
	= 0.006160% + 0.000065%
	=0.006225%
End of day estimated volatility	$\sqrt{.006225\%} = 07890\%$

The end of day volatility (0.7890%) is significantly lower than the previous day's estimated volatility (0.8095%). This is because the price movement on the day (0.33%) is quite low—much lower than the previous day's estimated volatility. It is a fundamental property of the weighted average that the average always lies between the two quantities being averaged. Therefore, the new volatility must lie between 0.8095% and 0.33%. On days when the market moves little, the end of day volatility is lower than the previous day's volatility. The reverse is true when there is a large price movement (upward or downward) on a day.

While the historical volatility method required a choice of the sample period, the exponentially weighted moving average method requires a choice of λ . In some sense, the value of λ determines the effective

Volatilities and Implied Volatilities | 14.5

sample period though the theoretical sample period is always infinite. When the effective sample period is defined as the period that contributes 99% of the weight in equation (14.4), it is found that a λ of 0.94 corresponds to an effective sample period of 74 days as may be seen from Table 14.2. The effective sample period can be computed as $\ln(0.01)/\ln(\lambda)$. The parameter λ can be estimated statistically but is often chosen judgementally.

 Table 14.2
 Effective sample size for various choices of λ . The effective sample size is defined as the period which contributes 99% of the weight in the exponentially weighted moving average. The effective sample period can be computed as $ln(0.01)/ln(\lambda)$

λ	Effective sample period (99% weight)
0.90	44
0.91	49
0.92	55
0.93	63
0.94	74
0.95	90
0.96	113
0.97	151
0.98	228
0.99	458

The exponentially weighted moving average method performs very well when there is substantial volatility clustering. In many financial markets, days of high volatility tend to occur in clusters separated by periods of low volatility. When the market enters a volatility cluster, the estimated volatility rises very quickly because of the large weight on recent data compared to past data. Similarly, when the market exits a volatility cluster, the estimated volatility method, this method responds very quickly to changes in volatility. This method is therefore to be preferred when there is volatility clustering.

14.3 GARCH METHOD

The Garch method can be regarded as an extension of the exponentially weighted moving average method in that the volatility is estimated each day as a weighted average of three quantities of which two are the same as in the exponentially weighted moving average method. While the exponentially weighted moving average method weighted averages the previous day volatility and the current day squared return, the Garch method adds a third quantity—the long run average variance.

This method takes into account a common observation in financial markets that each asset has a normal level of volatility towards which the actual volatility tends to revert. When the current volatility is high, it tends to decline towards the long run mean. Similarly, when the current volatility is low, it tends to rise towards the long run mean.

The Garch method therefore estimates the volatility at the end of each day as the weighted average of (a) the long run mean variance with weight $1 - w_1 - w_2$, (b) the previous day's estimated variance with weight w_1 and (c) the squared return that day with weight w_2 . The weights w_1 and w_2 should be non-negative and less than or equal to unity.

The exponentially weighted moving average method is a special case of Garch where $w_1 + w_2 = 1$ and therefore the weight on the long run variance becomes zero.

14.6 | Derivatives and Risk Management

The Garch model requires choice of the weights w_1 and w_2 as well as the long run variance. These parameters are typically estimated by statistical methods from past data. There is therefore no need for subjective judgement in choosing any of the parameters.

Consider an example where the long run volatility is 2.213% and the weights w_1 and w_2 are 94% and 4.85%. If the previous volatility is 2.352% and the current day return is -4.336%, then the updated volatility at end of day is 2.484% as computed in Table 14.3 below.

Table 14.3 Computation of updated volatility using Garch model. The long run volatility is 2.213% and the weights w_1 and w_2 are 94% and 4.85%. If the previous volatility is 2.352% and the current day return is -4.336%, then the updated volatility at end of day is 2.484%

	Long run	Previous Day		Current day
Volatility	2.213%	2.352%	Return	-4.336%
Variance	0.04897%	0.05532%	Squared Return	0.18801%
Weight	1.15%	94.00%		4.85%
Product	0.00056%	0.05200%		0.00912%
Updated Variance	0.00056% + 0.052	00% + 0.00912% =		0.06168%
Updated volatility at e	end of day			2.484%

The updated volatility is pulled in different directions. The mean reversion of volatility is pulling the volatility down from its previous value of 2.352% to its long run value of 2.213%. However, the very high current return of 4.336% overcomes this downward pull and pushes the volatility up sharply. The larger weight on the current day return as compared to the long run variance also plays a role in this outcome. It may be observed that if we applied the exponentially weighted moving average method in this case with λ equal to w_1 (0.94), the updated volatility would be 2.516% which is even higher than the Garch estimate of 2.484%.

The Garch model introduces another new feature—it provides different estimates of volatility for different dates in the future. The computation that we have shown above is actually the forecasted volatility for the next day. The forecast for day after tomorrow will be different because of the pull towards the long run volatility of 2.213%. The weight of 1.15% on the long run variance implies that 1.15% of the gap between the long run variance and the forecasted variance will be eliminated each day. Table 14.4 shows the computation of forecasts for the first few horizon dates.

In this example, the rate of mean reversion (weight on the long run variance) is quite low—only 1.15%. This means that the convergence is rather slow. Figure 14.1 shows that it takes almost a year for the forecast volatility to move close to the long run volatility.

It can be shown that when volatility is not constant over the life of the option, the variance to be used in Black-Scholes is the average variance over the life of the option. For example, if an option is considered seven days prior to maturity, one must compute the average of the variance values listed in Table 14.4 for days 1 through 7. The average of 0.06168%, 0.06154%, 0.06139%, 0.06125%, 0.06111%, 0.06097%, and 0.06083% is 0.06125%. This average variance can be converted into an average volatility by taking the square root of 0.06125% to obtain 2.4749%. The seven day option can be valued using the Black-Scholes formula with a volatility of 2.4749%.

This complication does not arise in the exponentially weighted moving average method because there is no mean reversion in that model. Absence of mean reversion implies that the forecast volatility is the same at all forecast horizons. The volatility estimate can therefore be used directly for options of all maturities.

Volatilities and Implied Volatilities | 14.7

Table 14.4Forecasting volatility at different horizon dates in the Garch model. As wemove further into the future, the mean reversion of volatility in the Garch model movesthe forecast volatility closer to the long run value. Every day, the gap from long runvariance is reduced by 1.15% which is the weight on the long run variance

Day	Volatility	Variance	Gap from long run variance	Reduction in gap (1.15% of gap)	Reduced variance	Reduced volatility
1	2.484%	0.06168%	0.01271%	0.00015%	0.06154%	2.481%
2	2.481%	0.06154%	0.01256%	0.00014%	0.06139%	2.478%
3	2.478%	0.06139%	0.01242%	0.00014%	0.06125%	2.475%
4	2.475%	0.06125%	0.01227%	0.00014%	0.06111%	2.472%
5	2.472%	0.06111%	0.01213%	0.00014%	0.06097%	2.469%
6	2.469%	0.06097%	0.01199%	0.00014%	0.06083%	2.466%
7	2.466%	0.06083%	0.01186%	0.00014%	0.06069%	2.464%
8	2.464%	0.06069%	0.01172%	0.00013%	0.06056%	2.461%
9	2.461%	0.06056%	0.01158%	0.00013%	0.06043%	2.458%
10	2.458%	0.06043%	0.01145%	0.00013%	0.06029%	2.455%



Figure 14.1 Convergence of forecast volatility to the long run volatility in the Garch model. The mean reversion in the Garch model means that as we move further and further into the future, the forecast volatility moves closer and closer to the long run volatility. In our example, the rate of mean reversion (weight on the long run variance) is quite low—only 1.15%. This means that the convergence is rather slow

14.4 IMPLIED VOLATILITY

All the three methods discussed above start with the time series of asset prices and apply different statistical techniques to estimate volatility. All these methods therefore use past volatility to forecast future volatility using models of varying degrees of sophistication. A completely different approach is to use market price of traded options to infer the market's forecast of future volatility. In an efficient market, this market-implied forecast can be expected to be superior to forecasts based on past volatility.

14.8 | Derivatives and Risk Management

Given an option price, one way to determine the market implied volatility is to use the method of trial and error in which one tries to value the option at several different values of volatility until a volatility is found for which the Black-Scholes value matches the observed market price. For example, suppose three-month option with a strike price of 105 is trading at 2.89 when the underlying is trading at 100 and the risk-free interest rate is 5% continuously compounded. A volatility of 20% may be tried, which gives a Black-Scholes option price of 2.48. Since this is lower than the market price, one tries a higher volatility of 25% to get a Black-Scholes value of 3.44 which is higher than the market price. Now, one knows that 20% is too low and 25% is too high and so a volatility in between say 22.50% is tried, to get a value of 2.96. This is still too high and so one tries a somewhat lower volatility of 22.25% to get a value of 2.91. One is now getting close but needs to go even lower. If 22.20% is tried, get a value of 2.90, one may go down a little lower to 22.15% which gives a value of 2.89. This is the end of the quest and it is concluded that the implied volatility is 22.15%. In a computer, this kind of trial and error process can be programmed quite easily and computing the implied volatility is quite trivial.

One might wonder what is the point of finding the implied volatility of an option from a market price and then plugging the same number back into the Black-Scholes formula to find an option price. This looks like a circular process. The reason why this is not a circular process and does prove useful is that the implied volatility can be used from one option to value another option. For example, knowing the implied volatility of 22.15% from the previous example, one may use it to value an option on the same underlying with a three-month maturity and a strike of 100. The Black-Scholes formula says that this option is worth 5.04. Thus, one has successfully used one option price to get the value of another.

This method of using the implied volatility of one option to value another, works well under the Black-Scholes assumptions of constant volatility and log normal distribution of asset prices. If these assumptions are violated, one needs to be more careful in using implied volatilities to value other options. This is what will be discussed in the next chapter on volatility smiles.

Chapter Summary

There are essentially two ways to estimate the volatility of asset prices. The first approach uses the time series of asset prices and applies different statistical techniques to estimate volatility. The assumption here is that past volatility can be used to forecast future volatility. Within this approach we have seen three methods of increasing sophistication. The first uses the historical volatility of a chosen sample period as the forecast of the future volatility.

A more refined approach places greater weight on more recent data while forecasting future volatility. Under this exponentially weighted moving average method the forecasted variance can be updated using the weighted average of the previous forecast and the squared return in the current period. This method is based on the tendency of asset market volatility to occur in clusters. The even more sophisticated Garch model expresses the forecasted variance as the weighted average of the long run variance, the previous forecasted variance, and the squared return in the current period. This method takes account of the tendency of volatility to mean revert to a long run value. Under this method, there is a different volatility forecast for each forecast horizon.

A completely different approach is to use market prices of traded options to infer the market's forecast of future volatility. In an efficient market, this market implied forecast can be expected to be superior to forecasts based on past volatility. The implied volatility of one option can be used to price other options.

In much of this chapter, the discussion has stayed within the standard Black-Scholes assumptions. Situations where volatility may vary over time can be handled by using the average volatility over the life of the option, provided the distribution of the asset price remains log normal. Departures from the log normal distribution require different methods which will be considered in the next chapter on volatility smiles.

Suggestions for Further Reading

Poon, Ser-Huang and Granger, Clive W. J. (2003) "Forecasting Volatility in Financial Markets: A Review" *Journal of Economic Literature*, 41(2), 478–539.

The implications of stochastic volatility for option valuation are discussed in

Heston, Steven L and Saikat Nandi (2000) "A closed-form GARCH option valuation model", *Review of Financial Studies*, 13 (3), 585-625

Problems and Questions

Day	Return	Day	Return	Day	Return
1	2.474%	35	1.488%	69	2.056%
2	0.522%	36	4.591%	70	-7.754%
3	1.379%	37	3.897%	71	-11.647%
4	-1.379%	38	0.246%	72	11.647%
5	- 0.348%	39	8.937%	73	- 0.841%
6	0.348%	40	9.949%	74	0.316%
7	0.864%	41	1.315%	75	3.109%
8	1.537%	42	2.187%	76	0.000%
9	0.338%	43	-6.710%	77	-3.425%
10	0.841%	44	-4.187%	78	-1.062%
11	-1.010%	45	2.169%	79	-3.254%
12	4.035	46	1.067%	80	-4.973%
13	0.162%	47	0.424%	81	-2.701%
14	0.966%	48	-2.029%	82	-1.681%
15	1.116%	49	-2.1815	83	0.483%
16	0.789%	50	3.574%	84	4.477%
17	-0.631%	51	12.109%	85	-0.926%
18	2.036%	52	3.610%	86	-6.858%
19	2.147%	53	-0.273%	87	6.158%
20	3.577%	54	9.058%	88	0.584%
21	0.292%	55	4.954%	89	1.731%
22	-1.619%	56	1.103%	90	-0.689%
23	0.296%	57	-3.185%	91	-1.042%
24	0.443%	58	-5.402%	92	-0.350%
25	0.880%	59	-0.944%	93	-4.663%
26	-5.091%	60	6.911%	94	-2.982%
27	4.505%	61	-0.646%	95	-0.887%
28	1.747%	62	-2.625%	96	1.515%
29	0.719%	63	6.048%	97	0.625%
30	4.893%	64	0.234%	98	-1.506%

The first 7 problems below are based on the 100 day of return given in the table below. Assume that the initial volatility (beginning of day 1) is 1.20% per day.

14.10 | Derivatives and Risk Management

Day	Return	Day	Return	Day	Return
31	-1.374%	65	-12.542%	99	-4.658%
32	-1.253%	66	-5.830%	100	3.769%
33	2.216%	67	3.684%		
34	0.546%	68	- 8.981%		

- 1. Compute the historical volatility using these 100 days of data.
- 2. Compute the daily volatility using the exponentially weighted moving average method with $\lambda = 0.94$.
- 3. Compare the results with the estimates based on $\lambda = 0.90$. What are the advantages and disadvantages of using a lower value of lambda?
- 4. Estimate the daily volatilities using the following Garch model. Compare the results with the estimates from the EWMA models of problem 2 and problem 3. Note that the model of problem 3 has the same weight on the squared daily return as this Garch model:
 - (a) The long run mean volatility is 2.21% so that the long run mean variance is 0.04884%
 - (b) The weight on the long run mean variance is 0.01.
 - (c) the weight on the previous day's estimated variance is 0.89
 - (d) the weight on the squared daily returns is 0.10.
- 5. Can you explain why the Garch model appears to track the model of problem 2 during some periods while tacking the model of problem 3 during some other periods?
- 6. What price movement would be required on day 1 to produce an end of day volatility estimate of 2% in the models of problem 2, 3 and 4 respectively. Explain the differences between these answers.
- 7. What would be your estimate of volatility on day 3 before your know the returns on day 1, 2 or 3 using the models of problem 2, 3 and 4.

Compute the implied volatility of the following options:

- 8. The price of the underlying is 100, the strike is 100, the risk free rate is 5%, maturity is three months and call price is 15.
- 9. All the data is as in problem 8 but the option is a put.
- 10. All the numbers are as in problem 8, but the underlying asset has a dividend yield of 3%.
- 11. All the data is as in problem 10 but the option is a put.
- 12. All the numbers are as in problem 8, but the underlying is a futures contract.
- 13. All the data is as in problem 12 but the option is a put.

Compute the following option prices:

- 14. The price of a call with a strike of 90 using the data in problem 8.
- 15. The price of a call with a strike of 90 using the data in problem 9.

Chapter **FIFTEEN**

Volatility Smiles and Implied Risk Neutral Distributions

Objectives

The standard Black-Scholes formula assumes that the risk-neutral distribution of asset prices is log normal. Asset prices are close to log normality, but exhibit somewhat fatter tails than the log normal. The risk-neutral distribution can deviate from the log normal in other ways. Rather than use a different formula which is more exact, market practice is to stick to the Black-Scholes formula but use a somewhat higher or lower volatility to compensate for the departures from log normality. This technique works, but unfortunately, it requires different volatility numbers to be used for different options. The curve depicting the volatility to be used for different strike prices is known as the volatility smile. This chapter discusses how the volatility smile is estimated and used.

15.1 REVISITING THE BLACK-SCHOLES FORMULA

In Chapter 10, the derivation of the Black-Scholes formula began with the observation that the call option price is the product of (a) the risk-neutral probability that S_T is greater than X (the probability of exercise), and (b) the risk-neutral conditional expected value of $S_T - X$ given that S_T is greater than X (the conditional expected value of the option when exercised). Here S_T is the asset price at maturity and X is the exercise price of the option. The Black-Scholes formula is obtained when the properties of the log normal distribution are used to compute these two quantities.

If the true risk-neutral distribution is approximately log normal but deviates slightly from it, estimates of the risk-neutral probabilities and conditional expectations will be close to the true value but not exact. For example, if the tails are fatter than the log normal, then for an out of the money option, the true risk-neutral probability that S_T is greater than X will be substantially higher than that given by the log normal distribution. As a result the Black-Scholes formula will undervalue the out-of-the-money call option quite significantly.

The theoretically correct procedure then would be to use the true risk-neutral distribution to compute both the probability of exercise and the expected value when exercised. Multiplying these two quantities would give the true value of the call option. Market practitioners are reluctant to do this because the log normal is both very convenient and approximately correct. They prefer therefore to take the log normal approximation and then use some simple adjustments to correct for the fatter tails. For example, if the log normal distribution is used but the volatility is deliberately kept higher than it actually is, the probability of exercise of an out-of-the-money option would increase. By choosing the volatility sufficiently high, one can reproduce the true risk-neutral probability.

Figure 15.1 shows the true risk-neutral distribution with fatter tails and a narrower neck than the best fitting normal distribution (N_{best}) . It is clearly problematic to use this normal distribution (N_{best}) to value out-of-the-money options, whose values depend on the probabilities in the tails. To solve this problem, one can artificially increase the variance to fit the tails better.

15.2 | Derivatives and Risk Management



Figure 15.1 True (fat tail) risk-neutral distribution of the log asset price compared to best fitting normal distribution, N_{best} . The true distribution has fatter tails and a narrower neck



Figure 15.2 True risk-neutral distribution of the log asset price compared with a high variance normal $(N_{high var})$ which fits the tails better but gives a poor fit in the middle

Figure 15.2 shows a high variance normal distribution $(N_{high var})$ that fits the tails better, but provides a poor fit to the middle of the distribution. Using $N_{high var}$ instead of the true risk-neutral distribution will give accurate values for out of the money options but will give erroneous answers for near-money options. To value options that are near the money, one will need a distribution that fits the middle of the distribution very well.

Figure 15.3 shows a normal distribution with a low variance $(N_{low var})$ that gives a good fit to the middle of the true risk-neutral distribution of the log asset price at the cost of a very poor fit in the tails. This distribution can be used to value options that are close to the money with a good degree of accuracy.

Thus one finds that it is possible to stick to the log normal distribution provided different log normal distribution are used different stirke prices. These different distributions differ only in the choice of the variance or the volatility. Therefore, it is possible to salvage the Black-Scholes formula at the cost of using a different volatility for each option strike price.

Volatility Smiles and Implied Risk Neutral Distributions | 15.3



Figure 15.3 True risk-neutral distribution of the log asset price compared with a low variance normal $(N_{low var})$ which fits the middle better but gives a poor fit in the tails

The volatility smile in this example (where the only departure from the Black-Scholes assumptions is the existence of fatter tails) might look like Figure 15.4. In practice, many different kinds of smiles are observed because there may be many different kinds of departures from the log normal assumption.



Figure 15.4 Volatility smile used to compensate for fatter tails. The implied volatility is lowest for at the money options and is higher for options that are either in the money or out of the money

15.2 VOLATILITY SMILE AND THE RISK-NEUTRAL DISTRIBUTION

It can be shown that for any given risk-neutral distribution, there is a corresponding volatility smile such that the Black-Scholes formula used in conjunction with this smile values all options correctly. Conversely, given the volatility smile observed in the market, one can infer the risk-neutral distribution that the market is using to value options.

For example, Figure 15.5 shows a volatility smile that is downward sloping. This corresponds to a risk-neutral distribution that is skewed to the left. In other words, there is a larger probability for a major downward move in the asset price than for a major upward move. Can the market be in equilibrium in

15.4 | Derivatives and Risk Management

such a situation? Yes, it can, because the probabilities of modest changes in asset prices would be skewed in the opposite way—the probability of a modest increase in the asset price would be much higher than the probability of a modest decrease in the asset price. The current asset price would thus be equal to the discounted expected future asset price despite the skewed distribution. This situation is illustrated in Figure 15.6.



Figure 15.5 A volatility smile that is downward sloping. The implied volatility is high for options with low strike prices and low for high strike prices. Such smiles are called volatility skews because they correspond to risk-neutral distributions that are skewed to the left



Figure 15.6 Skewed Risk-neutral Distribution. There is a larger probability for a major downward move in the asset price than for a major upward move. But the probability of a modest increase in the asset price is higher than the probability of a modest decrease in the asset price

In the US equity markets after the crash of 1987, there has been a fear in the markets about sharp price falls and a desire to hedge against this risk. As a result, the risk-neutral distribution of equity price exhibits an asymmetric pattern. The risk-neutral probability of a large market decline is much higher than

Volatility Smiles and Implied Risk Neutral Distributions | 15.5

the risk-neutral probability of a large market rally. Correspondingly, the volatility smile exhibits a large skew. This illustrates the fact that there can be an asymmetry in the risk-neutral distribution even if the asymmetry in the historical or actual distribution is absent or much weaker.

15.3 OPTION COMBINATIONS AND VOLATILITY SMILES

The shape of the volatility smile can be captured by prices of various option combinations. Consider for example the strangle consisting of long positions in the out-of-the-money call and the out-of-the-money put. In the absence of a smile, this combination would trade at the same implied volatility as the at-the-money option. A fat-tailed risk-neutral distribution would produce a volatility smile in which the out of the money options would trade at higher volatility. The excess of the strangle implied volatility over the at-the-money implied volatility tells one a great deal about the shape of the volatility smile.

Similarly, volatility skews can be seen in the prices of bull and bear spreads. Consider a bull spread consisting of a long in-the-money call at low strike and short out-of-the-money call at higher strike. If both options are equally far from the money, these two options would trade at the same implied if the volatility smile does not have any skew. If the long call is replaced by a long put at the same strike, the combination of long out-of-the-money put and short out-of-the-money call would have zero price if both are equally out of the money and the volatility smile has no skew. In the presence of a skew one or the other of the two out-of-money options would trade at a higher implied volatility and therefore the combination would no longer have zero value.

In the foreign exchange market, the exchange of an out-of-the-money put for an equally out-of-themoney call is known as a risk reversal and prices of these risk reversals are closely monitored as an indicator of market expectations of asymmetric price movements. Strangles and risk reversals capture much of the information contained in a volatility smile. Finer details about the smile are reflected in more complex option combinations.

15.4 TERM STRUCTURE OF VOLATILITY

So far, the behaviour of implied volatility as a function of strike prices has been considered. Implied volatility also varies with the maturity of the option. In the previous chapter, it was seen that the mean reversion in the Garch model produces a different forecast volatility for each forecast horizon. This means that one must use different implied volatilities for options of different maturities.

Because of mean reversion, implied volatilities of long maturity options do not move too far away from the long run mean volatility. Short maturity implied volatilities on the other hand, can sometimes climb to levels 10 or even 20 times the long run mean level, and they can sometimes drop to a small fraction of the long run mean. In such situations, the term structure of volatilities becomes a critical factor in pricing options accurately.

Though the Garch model predicts a term structure of volatility, it does so using only the mean reversion of volatility. In practice, several other factors like seasonality and transient shocks affect the term structure. For example, if the underlying is an agricultural commodity, the volatility in the post harvest months can be very different from that in the pre harvest period when there is a lot of uncertainty about the size of the crop. Even in financial assets, seasonality is often present because of a regular calendar of bond issuances, monetary policy announcements, elections and fiscal budget announcements.

The 'term structure of volatility' is therefore best estimated from the market prices of traded options of different maturities but similar strikes. This leads one to the concept of a volatility surface that we discuss in the next section.

15.6 | Derivatives and Risk Management

15.5 VOLATILITY SURFACES

When implied volatility is regarded as a function of both strike and maturity, one gets a volatility surface. For example, one can plot the strike prices on the X axis, maturity on the Y axis and the implied volatility on the Z axis.

It is inconvenient for this purpose to use the raw strike prices on the axis. For example if the underlying is at 100, a one day call option with a strike of 105 is probably deep out of the money while a three-month option at the same strike is only slightly out of the money. It is useful to transform the axis so that a given point on the scale corresponds to options of different maturities that are in or out of the money to the same degree.

One way to do this is to compute the number of standard deviations separating the strike price from the forward price for the given maturity. The standard deviation for quarterly returns is much higher than the standard deviation of daily returns and therefore the 105 call with one day to expiry may be five standard deviations out of the money while the 105 call with three months to expiry may be only about half a standard deviation out of the money. Another closely related way to achieve similar comparability is to plot the delta of the option instead of the strike price on the X axis. The delta lies between 0 and 1 with the at-the-money option having a delta close to 0.5, deep out-of-the-money options having deltas well below 0.5 and deep in-the-money options having delta well below 0.5.

An example of a volatility surface using option delta on the X axis is shown in Figure 15.7.





Volatility Smiles and Implied Risk Neutral Distributions | 15.7

In practice, a volatility surface is estimated by considering the implied volatility of all liquid options in the market and fitting a smooth surface to these volatilities. Different statistical tools can be used for this purpose. To price an illiquid option, one would read its estimated implied volatility off the fitted volatility surface and use the Black-Scholes model with this volatility to determine the price of the option.

For example, suppose calls with exercise prices of 110, 120 and 130 are trading at implied volatilities of 21%, 22% and 22.5% and it is desired to price a call with an exercise of 115. One needs to interpolate between the observed volatilities to estimate the implied volatility of the 115 call. If we interpolate linearly between 110 and 120, then the implied volatility of the 115 call should be 21.50%. A more refined method is to take account of the non linearity of the smile as shown in Figure 15.8. The best fitting quadratic smile is given by $y=-0.0025\% x^2+0.6750\% x-23.0000\%$ where x is the strike and y is the implied volatility. Using this curve, the implied for a strike of 115 is seen to be 21.5625%. In reality, it may be necessary to interpolate across both strikes and maturities and it may be necessary to use more sophisticated statistical techniques to determine the best fitting volatility surface.



Figure 15.8 Interpolating a smile accounting for its curvature. Three call options with strikes of 110, 120 and 130 are observed to trade at implied volatilities are 21%, 22% and 22.5%. It is seen that the smile is not linear. A quadratic smile is fitted as shown in the graph. Using this smile, the implied of the 115 call should be 21.5625% while linear interpolation would have given an implied of 21.50%.

15.6 IMPLICATIONS OF STOCHASTIC VOLATILITY

The very notion of the term structure of volatility implies that volatility is changing over time. This randomness of volatility by itself causes a departure from the log normal assumption. For example, if the true risk-neutral distribution is either log normal with a volatility of 15% or log normal with a volatility of 20% with equal probability, the unconditional risk-neutral distribution (the distribution before we know what the volatility would be) is no longer log normal. It has been seen that any departure

15.8 | Derivatives and Risk Management

from log normality is reflected in the shape of the volatility smile. Thus the randomness of volatility induces not only a term structure of volatility, but also affects the volatility smile. In this sense therefore, the volatility surface is a single integrated entity rather than an amalgam of independent smile and term structure components.

If one knows the true risk-neutral distribution of the random volatilities, one can theoretically compute the unconditional risk-neutral distribution of asset prices and then use this to obtain the theoretical option price as the risk-neutral discounted expected value. Several such models have been developed that give formulas more complex than Black-Scholes. These models have the advantage however of not requiring the use of a volatility surface provided that the randomness of volatility is the only departure from the Black-Scholes assumptions. These models will not be discussed in this book.

15.7 IMPLICATIONS OF VOLATILITY SURFACES FOR HEDGING AND OPTION GREEKS

Option traders and market makers often use options to hedge other options. It is possible to hedge the delta of an option by taking offsetting positions in the underlying (spot or futures), but the gamma and vega of an option can be hedged only by other options because the futures and the spot have zero gamma and vega.

It is common in practice to use the Black-Scholes model to compute the gamma, vega and other Greeks even if the Black-Scholes volatilities are taken off a volatility smile or surface. Strictly speaking, this is not correct because the smile represents a departure from the Black-Scholes assumptions and this requires a correction in the Greeks as well. The volatility surface computed off market prices provides a good compensation for these departures as far as option values are concerned but not necessarily for the option Greeks. In practice, this nicety is usually ignored.

A more important problem for hedging purposes is that there is a risk that the volatility surface itself might change. Consider for example a vega neutral portfolio that consists largely of short positions in near month options and long positions in far month options. The vega neutrality obtains only under the assumption that both near month and far month volatilities move together. It is possible however for the near month volatility to increase sharply while the far month volatility moves up only marginally. This implies a steepening of the term structure of volatility. In this scenario, the losses on the short positions in near month options would far outweigh the gains in the far month long options. The vega neutral portfolio would thus experience large losses due to the changes in volatility. Though the portfolio is vega neutral, it has a large exposure to what is often called time-vega.

Similar problems can arise when a portfolio is made gamma or vega neutral by combining options of different strikes. For example, a long position in at-the-money options may be hedged with short positions in out-of-the-money options of the same maturity to make the portfolio gamma or vega neutral. If the volatility smile becomes curved, the volatility of the out-of-the-money short options can increase substantially leading to large losses that far outweigh the gains from the long options whose volatilities may have risen much less.

Chapter Summary

The risk-neutral distribution of asset prices deviates from the log normal in various ways. It may have fatter tails and may also exhibit skewness. Since the Black-Scholes formula is based on the log normal distribution, these departures from log normality require adjustments to the Black-Scholes formula. Rather than use a more

Volatility Smiles and Implied Risk Neutral Distributions | 15.9

complex and exact model, market practitioners prefer to make an ad hoc adjustment to the Black-Scholes formula by using different implied volatilities for options at different strikes. The relationship between strikes and implied volatilities is known as the volatility smile.

Moreover, the mean reversion of volatilities implies that options of different maturities must be valued using different implied volatilities. The dependence of implied volatilities on strikes and maturities is known as the volatility surface.

The volatility surface is typically estimated from liquid options traded in the market and is then used to price illiquid options. Market practitioners typically use the same volatility surface to compute the gamma, vega and other option Greeks. This is not necessarily a correct procedure but practitioners ignore this nicety.

The existence of the volatility surface creates additional risks for hedging option positions. A portfolio that is gamma and vega neutral can still be exposed to risks arising from changes in the shape of the volatility surface.

Suggestions for Further Reading

- Breeden, D.T. and Litzenberger, R.H. (1978) "Prices of State-contingent Claims Implicit in Option Prices", Journal of Business, 51(4), 621-651.
- Jackwerth, J.C. and Rubinstein, M. (1996) "Recovering probability distributions from option prices", *Journal* of Finance, 51(5), 1611-1631.
- Jackwerth, Jens Carsten (1999)"Option-implied risk-neutral distributions and implied binomial trees: A literature review", *Journal of Derivatives*, 7(2), 66-82.
- Malz A.M. (1997) "Option implied probability distributions and currency excess returns", Federal Reserve Bank of New York Staff Reports No. 32.

Additional Suggestions for Further Reading

The implications of stochastic volatility for option valuation are discussed in **Heston**, **Steven L and Saikat Nandi (2000)** "A closed-form GARCH option valuation model", *Review of Financial Studies*, 13(3), 585-625.

Problems and Questions

1. Assume that a stock which does not pay any dividends is trading at 90 and that the interest rate is 5% continuously compounded. Consider the following options all of which have a maturity of three months:

	Option 1	Option 2	Option 3
Option type	Call	Put	Call
Exercise price	81	90	100
Option price	11	3	2

Plot the volatility smile based on this data.

- 2. What departure from the log normal distribution could account for the smile in problem 1?
- 3. What implied volatility would you use for an option with a strike of 95?
- 4. Assume that a stock which has a dividend yield of 3% is trading at 90 and that the interest rate is 5% continuously compounded. Consider the following options all of which have a maturity of three months:

15.10 | Derivatives and Risk Management

	Option 1	Option 2	Option 3
Option type	Call	Put	Call
Exercise price	81	90	100
Option price	11	3	0.5

Plot the volatility smile based on this data.

- 5. What departure form the log normal distribution could accont for the smile in problem 4?
- 6. At the money options with maturities of 1 month, 3 months and 6 months are trading at implied volatilities of 30%, 28% and 24% respectively. What volatility would you use for a 2 month option?
- 7. The weather department forecasts a strong possibility of a hurricane that would knock out a maor refining capacity for the next three weeks with. What do you expect to happen to the term structure of volatility when this forecast is made public?
- 8. Consider a trader who is short one-month futures and one-month call options while being long three month call options in a possition that is approximately delta, gamma and vega neutral. What would the weather forecast in problem 7 do to this position?
Chapter SIXTEEN

Exotic Options

Objectives

European and American call and put options and various combinations of these options are known as plain vanilla options as these are simple to value and to hedge. A number of other more complex derivative products known as exotics have become popular in recent years. This chapter discusses the valuation and hedging of exotic options. These options can also be valued using risk neutral valuation, but the simple Black-Scholes formula is no longer adequate. There are special difficulties in hedging some of these options because the option delta changes abruptly when the asset price moves by a small amount near certain thresholds.

16.1 DIGITAL OR BINARY OPTIONS

Digital or binary options are in some sense simpler than even the vanilla options. A digital call at a strike of 100 pays out a fixed amount say 1 if the underlying trades above 100 at expiry and zero otherwise. Its value is equal to the discounted value of the payoff times the risk neutral probability that the underlying is above the strike at maturity. In the course of deriving the Black-Scholes formula in Chapter 10, it was found that this probability is equal to $N(d_2)$ for a non dividend paying asset. Therefore, the value of the digital call C_d is given by the following formula which is even simpler than the Black-Scholes formula.

$$C_{d} = e^{-rT} N(d_{2})$$

$$d_{2} = \frac{\ln \frac{S_{0}}{X} + rT - \frac{1}{2}\sigma^{2}T}{\sigma \sqrt{T}}$$
(16.1)

where S_0 is the current asset price, X is the exercise price of the call option, T is the maturity of the option, σ is the volatility of the stock (standard deviation of the logarithmic return) and r is the continuously compounded risk free rate.

For example, if the current stock price is 100, the volatility is 23% and the risk-free rate is 8%, then the value of a three-month digital call option with an exercise price of 105 is 3.34 as shown below:

$$d_2 = \frac{\ln \frac{100}{105} + \frac{0.08}{4} - \frac{1}{2} \cdot 0.23^2 \cdot \frac{1}{4}}{0.23\sqrt{1/4}} = \frac{-0.0487900}{0.115} = -0.307849$$

$$N(d_2) = 0.379099 \text{ while } e^{-008/4} = 0.980199$$

$$C = 0.980199 \times 0.379099 = 0.3716$$

Using put-call parity, we can then compute the value of a European digital put option as follows:

$$P_d = e^{-rT} N \left(-d_2\right)$$

The McGraw·Hill Companies

16.2 | Derivatives and Risk Management

$$d_{2} = \frac{\ln \frac{S_{0}}{X} + rT - \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}}$$
(16.2)

Continuing the earlier example, if the current stock price is 100, the volatility is 23% and the risk-free rate is 8%, then the value of a three-month digital put option with an exercise price of 105 is 0.608607 as shown below:

$$d_2 = -0.307849$$
 as for the call option
 $N(-d_2) = 0.620901$ while $e^{-008/4} = 0.980199$
 $P_d = 0.980199 \times 0.620901 = 0.608607$

In the presence of a dividend yield, the above formulas have to be modified slightly:

$$C_{d} = e^{-rT} N(d_{2})$$

$$P_{d} = e^{-rT} N(-d_{2})$$

$$d_{2} = \frac{\ln \frac{S_{0}}{X} + (r-q)T - \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}}$$
(16.3)

where S_0 is the current asset price, X is the exercise price of the call option, T is the maturity of the option, σ is the volatility of the stock (standard deviation of the logarithmic return), q is the continuous dividend yield and r is the continuously compounded risk-free rate.

The simplicity of the formula for the value of a digital call option conceals some strange behaviour close to expiry. At expiry, the call is worth nothing if the underlying is even a little bit below the strike, but is worth the full amount if the underlying is just a bit higher. In other words, the value moves discontinuously from 0 to 1 as the underlying moves past the strike. By contrast, in a plain vanilla call, the value of the option at expiry rises slowly as the option moves into the money and there is no discontinuity. Prior to expiry, the value of the digital call is the risk-neutral expected discounted value of its payoff at expiry. The averaging process involved in computing the risk-neutral expected value smooths out the discontinuity.

For example, if the asset is just above the strike, there may be, say a 51% probability that it will expire above the strike and say a 49% probability that it will expire below the strike. The risk-neutral expectation would then be $0.51 \times 1 + 0.49 \times 0 = 0.51$ and the value of the option would be the discounted value of this. If the asset is a little below the strike, the probabilities would be the other way around 49% that it will expire above the strike and say 51% probability that it will expire below the strike.

The risk-neutral expectation would then be $0.49 \times 1 + 0.51 \times 0 = 0.49$ and the value of the option would be the discounted value of this. The option value therefore changes smoothly as the asset price changes and there is no longer any discontinuity. However, close to maturity, the change in the value of the digital call is quite rapid. This situation can be seen in Figure 16.1.

The behaviour of the call can be more clearly seen in terms of the option delta. For a plain vanilla call, the delta rises gradually from 0 to 1 as the asset price rises. The delta of the digital call behaves differently and the difference is more pronounced close to maturity. The delta is close to zero except for a range of asset prices on either side of the strike. Close to expiry, this range is very narrow and within this range, the delta rises rapidly above zero, peaks, and then falls back to zero. This abrupt change in



the call delta creates a great deal of difficulty in delta hedging a digital call. When the asset price is near the strike close to expiry, the delta hedge needs to be rebalanced drastically.

Price of Underlying

Figure 16.1 Digital call value as a function of the asset price for various maturities. When we consider the option at maturity, the call value moves discontinuously from 0 to 1 as the underlying moves past the strike. Prior to maturity, this discontinuity is smoothed away by the averaging process of computing the risk neutral expectation

A small move in the asset price changes the delta drastically and the delta hedge becomes impractical under such situations. Another way of looking at the situation is in terms of the option gamma which is the rate of change of the delta. It is obvious from the sharp peak of the delta that the option gamma becomes very large near the strike when the option is close to expiry.

Since dynamic delta hedging is problematic for exotic options like the digital call, an attractive alternative is to replicate them statically with vanilla options. Consider an investor who wants to buy a digital call at a strike of 100 where the payoff is 1 if the option expires in the money. One possible way for this investor to replicate the payoffs of the digital call is a bull spread—she buys a call at 99.50 and sells a call at 100.50. Consider the range of possibilities:

- 1 The asset closes below 99.50 at expiry. In this case, the digital expires worthless and the two vanilla calls also expire worthless. Thus the bull spread has the same payoff as the digital.
- 2 The asset closes above 100.50 at expiry. The digital expires in the money and pays off 1.00. Both the vanilla calls in the bull spread also expire in the money. The investor ends up buying the asset at 99.50 and selling it at 100.50 and makes a profit of 1.00. Thus she again obtains the same payoff as the digital.

The McGraw·Hill Companies

- 16.4 | Derivatives and Risk Management
 - 3 The asset closes between 99.50 and 100 at expiry. The digital expires worthless but the spread produces a profit between 0 and 0.50.
 - 4 The asset closes between 100 and 100.50 at expiry. The digital expires in the money and pays off 1.00. The spread produces a profit between 0.50 and 1.00.



Figure 16.2 Delta of Digital Call Option. The delta is close to zero except for a range of asset prices on either side of the strike. Close to expiry, this range is very narrow and within this range, the delta rises rapidly above zero, peaks and then falls back to zero. Corresponding to this sharp peak, the gamma becomes very high near the strike

Thus we see that the spread replicates the digital call except when the asset closes in the narrow range between 99.50 and 100.50. To improve the replication the investor can consider buying two bull spreads — she buys two calls at 99.75 and sells two calls at 100.25. Now if the asset trades above 100.25 at expiry, she earns a profit of 0.50 from each of the two spreads by buying at 99.75 and selling at 100.25. The portfolio consisting of two spreads earns a profit of $0.50 \times 2 = 1.00$ which is the same as the payoff of the digital call. Therefore the portfolio replicates the digital call except in the narrower range of 99.75 to 100.25.

Theoretically, she could consider more and more refined replicating portfolios. For example, she could buy 10 calls at 99.95 and sell 10 calls at 100.05. Each of these spreads produces a profit of 0.10 when the asset closes above 100.05. The portfolio of 10 such spreads earns a profit of 1.00 which is the payoff of the digital call. In practice, she would not find liquid options at arbitrary strikes and she would be forced to live with a modest range of prices in which the portfolio of spreads does not replicate the digital call option perfectly.

Now consider an option trader who has sold a digital call to a customer and wants to hedge it. He could consider hedging his short position by buying the replicating portfolio of vanilla option spreads. He might for example sell a call at 99.50 and buy a call at 100.50. The profit on this spread would be equal to the loss on the sold option except when the asset closes in the range of 99.50 to 100.50.

The big advantage of the hedge using a spread of vanilla call options is that it is a static hedge. It remains unchanged as the asset price changes. Unlike a delta hedge which requires continuous trading to rebalance the hedge, this does not therefore produce large transaction costs from frequent trading.

16.2 BARRIER OPTIONS

Barrier options are an important and popular class of exotic options. These can be either knock-in or knock-out options. A knock-in option comes to life only when a barrier is hit. Once it has been knockedin, it is a vanilla option. A knock-out option is extinguished when a barrier is hit. If the barrier is never hit during the life of the option, it becomes a vanilla option at maturity. In both cases (knock-in and knock-out) the barrier can be above or below the current strike.

This give-rise to four possibilities — the up and in (a knock-in whose barrier is above the strike), the the up and out (a knock-out whose barrier is above the strike), the down and in (a knock-in whose barrier is below the strike), the down and out (a knock-out whose barrier is below the strike). Each of these barrier options can be a call or a put. This leads to eight possibilities in all.

There is however an important parity relation between knock in and knock out options similar to the put-call parity that links put and call options. The relationship is that a vanilla European call is the same as a knock in call combined with a knock-out call where all three calls have the same strike and both barrier options have the same barrier. For example, consider the 100 call (European), the 100 call with knock-in at 90 and the 100 call with knock-out at 90. Buying both barriers ensures that one and only one of them will be alive at maturity. If the underlying hits 90 at any time during the life of the options, the knock-out will be extinguished and the knock-in comes to life. If the underlying never hits 90 during the life of the options, the knock-out will remain alive and the knock-in never comes to life. The effect of holding both barriers therefore is the same as holding the vanilla European call.

Some of these options exhibit very strange behaviour like knocking out while in the money. For example consider an up and out call with a strike of 110 and a knock out barrier of 120 bought when the underlying is at 100. At the beginning, the option is out of the money, and as the underlying rises above 100, it moves into the money. If however, the underlying hits 120, the option is knocked out and becomes completely worthless. As the asset prices rises from 100 to 120, the option first increases in value as it moves into the money and acquires an intrinsic value. If the asset continues to rise and hits the barrier, the option loses the entire intrinsic value and becomes completely worthless. This discontinuity in the payoff of the up and out call is similar to that of the digital call at maturity that we saw earlier.

Anticipating the knock out, the up and out call would fall in value as the underlying approaches the knock out barrier. In other words, the delta which was initially positive would turn negative. Similarly, the vega would also change sign. When the option is out of the money the up and out call has positive vega as higher volatility increases the probability that the option will move into the money. However as the underlying moves close to the barrier, the vega turns negative as high volatility increases the chance that the option will be knocked out. This kind of option, therefore, behaves very differently from vanilla options whose delta and vega never change sign during the life of the option.

By the parity relationship discussed above, the up and in call also exhibits strange behaviour that mirrors the behaviour of the up and out call. For example, the 110 call that knocks in at 120 has no intrinsic value before the barrier has been hit and then suddenly acquires an intrinsic value of 10 when the barrier is touched. As the underlying rises towards 120, the up and in call would rise dramatically in

16.6 | Derivatives and Risk Management

value in anticipation of the barrier being hit. The delta and vega of the up and in call would be quite high at this stage when the up and out call has a negative delta and vega.

Many of the popular barrier options, however, do not have these strange properties. For example, the down and out call knocks out when it is out of the money and there is no discontinuity in the payoff. For example, consider a 110 call that knocks out at 90. As the underlying moves towards 90, it becomes less and less valuable as it moves out of the money. When it hits 90, the option becomes completely worthless but there is no discontinuity as the intrinsic value is zero. The delta and vega of the down and out call are always positive. The down and in call is also similarly unproblematic.

Barrier options have become quite popular because they are cheaper than vanilla options. Potential buyers of options are often deterred by the option premium. They are constantly looking at ways to reduce the premium and one alternative is to sell some options as discussed in Chapter 11.

Another alternative is to use barriers and other exotics that are less expensive. Typically the buyer of the options has a view on asset prices that makes the barrier attractive. For example, he may have an overall bullish view on the underlying asset while recognizing the possibility of a modest fall in asset prices. At the same time, he may be confident that the price will not fall too much because of a strong support level at about 10% below the current price. He may then be quite happy to buy a down and out call with a barrier 15% below the current price.

The purchase of a call option expresses his overall bullish view while protecting him against the possible downside from a fall in the asset price. Given his belief in the strong support level at 90, he does not worry too much that the asset may fall to 85 and knock out his option. At the same time, the down and out call may be significantly cheaper than the vanilla call and attenuates his reluctance to pay a high option premium.

16.2.1 Put-Call Symmetry

Under the log normal assumptions of the Black-Scholes model, it is possible to find formulas for the valuation of barrier options. There is a simple symmetry relation between puts and calls that are equally out of the money. Assume that the risk-free interest rate and the dividend yield on the asset are both zero so that the asset price has zero drift (zero expected growth rate). Now consider an out-of-the-money call with a high strike H where $H > S_0$. An equally out-of-the-money put would be one whose strike L is low

$$L < S_0$$
 and $\frac{L}{S_0} = \frac{S_0}{H}$. This relationship is the same as $\sqrt{LH} = S_0$ or $\ln S_0 = \frac{1}{2}$ ($\ln L + \ln H$). Since the

risk-neutral distribution of the asset price is log normal, one needs to ensure that the log of the stock price is midway between the logs of the two strikes to make the two options equally out of the money. For example, if the asset price is 100 and the out-of-the-money call is at 125, the out of the money put

would be at 80 because $\frac{125}{100} = \frac{100}{80}$ Intuitively, we would expect the two equally out-of-the-money options to be worth the same because of the symmetry of the log normal distribution. One's intuition is not completely wrong and in fact the risk-neutral probabilities of the two options ending up in the money

are exactly the same. However, the call is slightly more valuable than the put as shown below.

If d_1 , and d_2 denote the Black-Scholes d_1 and d_2 of the call, and d_1^* , and d_2^* denote the corresponding quantities for the put, the following relationship is arrived at:

Exotic Options | 16.7

$$-d_{2}^{*} = \frac{-\ln\frac{S_{0}}{L} + \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}} = \frac{\ln\frac{S}{H} + \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}} = d_{1}$$
(16.4)

$$-d_{1}^{*} = \frac{-\ln\frac{S}{L} - \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}} = \frac{\ln\frac{S}{H} - \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}} = d_{2}$$
(16.5)

This means that the value of the put can be written as:

$$p^{*} = LN(-d_{2}^{*}) - S_{0} N(-d_{1}^{*}) = LN(d_{1}) - S_{0} N(d_{2}) = \frac{L}{S_{0}} [S_{0} N(d_{1}) - HN(d_{2})]$$
(16.6)

Recall that the last term in square brackets $[S_0 N(d_1) - HN(d_2)]$ is the value of the call option. Thus instead of the naive intuition that $p^* = c$, we have the slightly different relationship:

$$p^* = \frac{L}{S_0} c = \frac{S_0}{H} c \text{ or } c = \frac{S_0}{L} p^* = \frac{H}{S_0} p^*$$
(16.7)

The put is less valuable than the call because though it has an equal probability of being in the money, its payoff is limited (at best the asset price can go to zero) while the potential payoff of the call is unlimited.

16.2.2 Static Hedge of the Down and In Call

Coming back to the down and in call with strike *H* and barrier *B*, we can use Eq (16.7) to find a static hedge for the down and in call. At the moment when the asset hits the barrier, the down and in call comes to life as an out of the money call. At this point the asset price is equal to the barrier *B* and we can construct an equally out of the money put with strike *L* such that $\sqrt{LH} = B$. By Eq (16.7), the value of the call that has just come to life is equal to $\frac{H}{B}$ times the value of the put with strike *L*. Consider an

investor who instead of buying the down and in call with strike *Hand* barrier *B*, buys $\frac{H}{B}$ vanilla puts

with *B* strike *L*.

Suppose this person adopts the strategy that if ever the asset hits the barrier B, he would sell the puts and buy a vanilla call with strike H. There is no cash flow when he does this because the puts and the call have the same value when the asset price is equal to B. If the barrier is not hit, he would hold the put till maturity. The payoff to this person is the same as what he would have got by buying the down and in call. If during the life of the option, the barrier is never hit, then the asset is above the barrier and therefore above L at expiry. The puts expire worthless and the down and in call is also worthless because it never came to life. If the barrier is hit at some time, the down and in call becomes a vanilla call with strike H, but our investor would also have sold the puts and bought the same vanilla call.

Thus the static hedge of the down and in call consists of the following strategy:

- 1. At inception, buy $\frac{H}{B}$ vanilla puts with a strike L such that $\sqrt{LH} = B$
- 2. If and when the barrier is hit, exchange the above puts for a vanilla call with a strike of H.

16.8 | Derivatives and Risk Management

For example, a barrier call with a strike of 100 knocked in at 90 can be statically hedged by buying $\frac{10}{9} = 1.1111$ puts with a strike of 81. If the asset price ever hits 90, these puts would be sold and the proceeds used to buy a vanilla call with a strike of 100.

16.2.3 Valuation of Down and In Call

The static hedge can also be used to value the down and in call by valuing its static hedge. For example, if the current asset price is 95, the down and in call with strike of 100 and barrier of 90 should have the same value as 1.1111 puts with a strike of 81. If the interest rate is 0, the maturity is 3 months and the volatility is 20%, the price of a 81 call using the Black-Scholes formula is 0.206458. The price of the down and in call is the same as that of 1.1111 such puts and is therefore equal to $0.206458 \times 1.1111 = 0.23$. By contrast, the value of a vanilla call with the same strike of 100 is given by the Black-Scholes formula as 1.89. The down and in call is far less expensive than the vanilla call for two reasons:

- 1 There is a high probability that the barrier will not be hit and so the down and in call would not even come into existence.
- 2 When the barrier is hit and the option comes into existence, it would be much more out of the money than it is now. The distance from the strike to the current asset price is only 100-95=5 while when the barrier is hit, the distance to the strike would be 100-90=10.

More generally a down and in call with strike H and barrier B, must be $\frac{H}{B}$ times the value of a vanilla

put with strike *L* where $\sqrt{LH} = B$ or $L = \frac{B^2}{H}$.

Under the zero drift assumption with which we began, therefore, the down and in call with a barrier B at or below the strike H has the value given by:

$$c_{di} = \frac{H}{B} p^* = \frac{H}{B} [LN(-d_2^*) - S_0 N(-d_1^*)] = BN(-d_2^*) - \frac{H}{B} S_0 N(-d_1^*) \quad (16.8)$$

It is convenient to rewrite this slightly differently as

$$c_{di} = BN(d^*) - \frac{H}{B} S_0 N(d^* - \sigma \sqrt{T})$$
(16.9)

where

$$d^* = \frac{\ln\left(\frac{B^2}{S_0H}\right)}{\sigma\sqrt{T}} + \frac{1}{2}\sigma^2 T$$
(16.10)

Observe that these formulas reduce to the **B**lack-Scholes formulas for vanilla options when the current asset price S_0 is equal to the barrier *B*.

16.2.4 Valuation of other barrier options

If we relax the zero drift assumption and reintroduce the risk-free interest r and the dividend yield q, the static hedge no longer works but there is a more complex valuation formula for the down and in call under the assumption of log normality:

Exotic Options | 16.9

$$c_{di} = \left(\frac{B}{S_0}\right)^{2\lambda} S_0 \mathrm{e}^{-qT} N(d^*) - H e^{-qT} \left(\frac{S_0}{B}\right)^{2-2\lambda} N(d^* - \sigma\sqrt{T})$$
(16.11)

where

$$d^* = \frac{\ln\left(\frac{B^2}{S_0 H}\right)}{\sigma\sqrt{T}} + \lambda\sigma^2\sqrt{T}$$
(16.12)

$$\lambda = \frac{r - q + \frac{1}{2}\sigma^2}{\sigma^2} \tag{16.13}$$

The no drift case that we discussed earlier corresponds to the special case where $\lambda = \frac{1}{2}$. Observe that these formulas also reduce to the Black-Scholes formulas for vanilla options when the current asset price S_0 is equal to the barrier *B*.

Values of down and out calls can be obtained by using the parity relationship as being equal to the value of a vanilla call less the value of a down and in call. Somewhat more complex formulas have been derived for down and in calls as well as down and out calls when the barrier is above the strike. Formulas also exist for up and out as well as up and in barrier options as well.

All of these formulas assume that the risk-neutral distribution is log normal. They therefore assume that there is no volatility smile or skew in the market. In practice, there is usually a volatility smile in the market. Practitioners are reluctant to use a formula for barrier options that does not price vanilla options correctly. As already discussed, the price of the barrier option converges to the price of vanilla options when the asset price approaches the barrier. Therefore, a valuation formula for barrier options that ignores the smile can produce values that are inconsistent with the observed prices of vanilla options.

Practitioners, therefore, often value barrier options using binomial option pricing models using the implied risk-neutral tree. In other words the asset price tree is constructed taking into account the observed volatility smile or skew. Using this tree allows vanilla options to be priced correctly. Barrier options are priced off the same tree to get a valuation of these options that does not conflict with the observed prices of vanilla options.

In some markets, concerns have been raised about the possibility of market manipulation by traders with large positions in barrier options. A trader who is short knock out options has an enormous incentive to manipulate the market price so that it touches the barrier even momentarily. Even if the market price rebounds quickly after the manipulation ceases, the options have been knocked out and the trader has made large gains that may far exceed the costs of manipulating the market. Similarly, an investor with a large position in knock in options also has an incentive to manipulate market prices so that they touch the barrier at least momentarily.

16.3 ASIAN OPTIONS

Asian options have a payoff based on the average price of the asset during the life of the option rather than just the value at maturity. For example, the payoff of an Asian call is equal to max $(\overline{S} - X, 0)$ where \overline{S} is the average price of the asset during the life of the option. The payoff of the Asian put is max $(X - \overline{S}, 0)$.

16.10 | Derivatives and Risk Management

Asian options have several uses. For example, an exporter who has a regular export business in which shipments are made every day is interested in the average price realization of her monthly ex- ports. The month end exchange rate is of little significance to her because it affects only the earnings for the shipments effected on the last day of the month. She is more worried about the average exchange rate during the month. If she wants to hedge her foreign exchange risk with an option, she would like an option whose payoff is linked to the average exchange rate rather than the month end exchange rate. An Asian put option on the foreign currency suits her requirements well.

Another example of the advantage of Asian options is the case of thin, illiquid markets. In these markets, people are often worried about the month end or year end prices being manipulated by some large traders. People might be less worried about manipulation of the average price because it is too difficult and expensive for the manipulator to manipulate prices in the same direction every day. In such markets, people may be more comfortable with Asian options.

In practice the average price used in Asian options is an arithmetic average of the asset prices. But it is instructive to consider a hypothetical case where the average is a geometric mean. The log of the geometric mean is the average of the log asset price during the life of the option. Since each of the log prices is normally distributed, their average is also normally distributed. The geometric mean therefore obeys the log normal distribution that makes it possible to use the Black-Scholes formula. The only adjustment that is required is that the volatility of the average price is lower than the volatility of the price at maturity.

This happens due to two reasons. First the prices at nearby dates are less volatile than prices at distant dates. Since the prices are being averaged all the way from the first day to the date of maturity, they are on average, at a distance of half the time to maturity. This by itself should reduce the volatility of the average to the volatility of an option with half the maturity of the original option. By the square root scaling discussed earlier in this book, this would correspond to a volatility equal to $1/\sqrt{2}$ times the volatility of the price at maturity. It can be shown that when we take this into account the volatility of the geometric mean is equal to $1/\sqrt{3}$ times the volatility of the price at maturity. Thus the value of an Asian option corresponds roughly to the value of a European option with one-third the time to maturity.

This broad correspondence is roughly correct and many practitioners hedge an Asian option by using a European option with one-third the time to maturity. However, for a proper valuation formula several adjustments need to be made. First the average used in most Asian options is an arithmetic mean and not a geometric mean. Strictly speaking the distribution of the arithmetic mean is not log normal and the use of the Black-Scholes formula is not wholly appropriate.

However, the departure of the arithmetic mean from the log normal distribution is not very large and the Black-Scholes formula can be used as a good approximation. It is however necessary to match the

mean and variance of the distribution more carefully than the simple $1/\sqrt{3}$ adjustment made above. Second, it is necessary to take account of the exact nature of the average used. For example, the option might use average of end-of-day prices or it may use an average of all prices including intra-day prices. This makes another small difference to the valuation formula. Unlike many other exotic options, the Asian option is easier to hedge than the vanilla option. As one moves closer to maturity, the value of the Asian option becomes more and more certain. The unknown evolution of the asset price in the remaining life of the option makes only a small difference to the payoff of the option because most of the prices used to compute the average price are already known. The option gamma and vega therefore come down substantially making it much easier to hedge.

16.4 CHOOSER OPTIONS

A chooser option is one in which the holder can at maturity decide whether the option is a put option or a call option. The value of the chooser at maturity is therefore equal to max (c, p) where c and p are the values of the call and put options. By put-call parity, $p = c + X e^{-rT} - S$. Therefore,

$$\max (c, p) = \max (c, c + X e^{-rT} - S) = c + \max (0, X e^{-rT} - S)$$

The chooser is therefore a combination of:

- 1. A call option with strike X
- 2. A put option with strike $X e^{-rT}$

This derivation assumes that the asset does not pay a dividend yield. In the presence of a dividend yield at the rate q, the analysis changes slightly:

$$\max(c, p) = \max(c, c + X e^{-rT} - S e^{-qT}) = c + e^{-qT} \max(0, X e^{(q-r)T} - S)$$

The chooser then becomes a combination of a call option with strike X and e^{-qT} put option with strike $Xe^{(q-r)T}$.

16.5 COMPOUND OPTIONS

A compound option is an option on an option. For example, one could consider a call option which allows the holder at maturity to buy another option. The naive approach is to consider the second option as the underlying and value the compound option by using the Black-Scholes formula. This would require estimating the volatility of the second option price using its delta. This method is incorrect because the option price is not log normally distributed. The option price is a non linear function of the underlying and this changes the distribution.

In practice, if the second option has a long maturity (much longer than the maturity of the first option), using the Black-Scholes formulas provides a tolerably good approximation. For example, an equity share of a company with some debt in its capital structure is in some models regarded as an option on the assets of the firm. The shareholders have the right to pay off the debt and take control of all the assets. This represents a call option with an exercise price equal to the face value of the debt. An option on the equity share is then an option on an option or a compound option. In practice, however, the Black-Scholes formula provides a very good approximation to the value of this option.

When this approximation is not adequate, the bivariate normal model is needed. The asset price S_1 at the maturity of the first option T_1 and the asset price S_2 at the maturity of the second option T_2 are both log normally distributed. However, these two quantities are related and therefore we cannot look at them as two separate log normally distributions. It is necessary to look at them jointly as a single distribution. The bivariate normal distribution describes the joint distribution of the logs of S_1 and S_2 . Using this joint distribution, it is possible to arrive at a valuation formula for compound options.

The McGraw·Hill Companies

16.12 | Derivatives and Risk Management

16.6 OTHER EXOTIC OPTIONS

Several other exotic options have been developed and used in the market place:

- A lookback call option allows the holder at maturity to look back at the path of asset prices during the life of the option and buy the asset at the lowest price during the period. A lookback put option allows the holder at maturity to sell the asset at the highest price observed during the life of the option. Under the log normal assumption, formulas have been derived for valuing these options.
- A shout option is in some sense like an American option. It allows the holder to 'shout' at any time during the life of the option. This is not quite the same as exercising the option. The value that would be obtained by exercising the option at the shout date (the intrinsic value at this point) is recorded. At maturity, the holder has the choice of exercising the option at maturity. He also has the option of retrospectively exercising the option at the shout date and receiving the value that he would have received at that date. In other words, the holder receives the higher of the intrinsic value at the shout date and the intrinsic value at expiry. This option can be valued using a binomial tree in a manner similar to valuing American options.
- A forward start option is one which will come into existence at some future date with a strike price dependent on the asset price at the start date. Usually, the strike price is set equal to the then prevailing asset price so that the newly created options are at the money at inception. These can be valued using a slight modification of Black-Scholes formula.

Chapter Summary

Exotic options have become popular either because they are cheaper than vanilla options or because they are more convenient or because they allow investors to take nuanced views about the evolution of asset prices. For most exotics, valuation formulas have been derived using risk neutral valuation under the assumption of a log normal distribution. In the presence of a volatility skew and smile, these options become quite difficult to value. Many exotics are quite difficult to hedge because of high gammas or discontinuities in option prices or deltas. Some exotics however are easier to hedge than vanilla options.

Suggestions for Further Reading

The following paper provides valuation formulas and pay of diagrams for various exotic options:

Faseruk, A., Deacon, C. and Strong, R. (2004) "Suggested Refinement to Courses on derivatives: Presentation of Valuation Equation, Pay Off Diagrams and Managerial Application for Second Generation Option", *Journal of Financial Management and Analysis*, 17(1), 62-76.

The difficulties of hedging exotic options statically and dynamically are described in:

Tompkins Robert G (1997) "static versus dynamic hedging of exotic options: an evaluation of hedge perfor- mance via simulation", http://sina.sharif.edu/~khoshnevisan/TOMPKINS.PDF

Problems and Questions

- 1. Assume that a stock which does not pay any dividends is trading at 95 and that the volatility is 25% continuously compounded. What is the price of a binary or digital call option with a strike price of 100 and a maturity of three month?
- 2. Determine a suitable static hedge for the digital call in problem 1. What would be the cost of this option and how does it compare with the theoretical price of the digital?
- 3. A bond promises to pay a coupon of 7% if Libor is between 5.00% and 5.50%. The coupon is zero otherwise. Can you express the coupon in terms of digital option on Libor?
- 4. Assume that a stock which does not pay any dividends is trading at 110 and that the volatility is 25% and the interest rate in zero. Construct a static hedge of a down and in call with a strike of 125 and a barrier of 100. Use this static hedge to value the down and in call.
- 5. Use Eq (16.9) to value the barrier option of problem 3 and compare the answer with that obtained by valuing the static hedge.
- 6. Assume the same data as in problem 3, except that the risk free rate is 6.25% and the dividend yield is 2.5%. What is the value of the barrier option?
- 7. A gold mining company sells call options for 100,000 ounces of gold with a strike of \$400 per ounce when gold is trading at \$300. The contract however says that if gold rises above 425, the quantity to be delivered rises to 120,000 ounces and if if gold rises above 450, the quantity to be delivered rises to 150,000 ounces. Express this option in terms of vanilla and barrier options.
- 8. An exporter receives dollar inflows uniformly throughout the month. The current hedging strategy to, hedge the exposure during the quarter is based on buying vanilla dollar put options expiring at the end of one month, two months and three months. What type of option could you recommend to exposure better? Would this have a lower or higher price than current strategy?

Chapter Seventeen

Warrants and Convertibles

Warrants are call options on a company's shares that are issued by the company itself. When the warrants are exercised, the company issues fresh stock (instead of delivering pre-existing shares) to fulfil its obligations. This chapter discusses the difference between call options and warrants in order to determine whether a different valuation procedure is required. Employee stock options are called options but they are actually warrants and so these are also discussed here. Convertible bonds can be analysed as a combination of straight bonds and warrants. This chapter discusses how this decomposition can be used to value convertibles.

17.1 WARRANTS

A warrant is essentially a call option on the stock that is issued by the company itself. The key difference from an ordinary call option is, when the warrant holder exercises the option, the company does not deliver pre-existing shares of the company, instead it issues new shares to the warrant holder. This leads to an increase in the total number of shares outstanding. This dilutes the value of the existing shareholders who end up owning a smaller fraction of the company than they did before the issue of the new shares. Warrants thus make the existing shares less valuable. Such a situation does not arise with ordinary call options where no new shares are created and only existing shares change hands.

17.2 THE WARRANT VALUATION MODEL

If one knows how to value ordinary call options on a stock, can one use this to value a warrant also? Consider two companies A and B which are identical in all respects except that Company A has issued warrants and Company B has not issued warrants. In particular, the assets of the two companies are identical. This means that the combined value of all the shares and warrants of Company A must be equal to the value of all the shares of Company B. Suppose both companies have *n* shares outstanding and Company B has *m* warrants outstanding, each of which has an exercise price of *K*. Then we have

$$nS_t + mW_t = nV_t \tag{17.1}$$

where S_t is the price per share of Company A at time *t*, W_t is the price per warrant of company A at time *t* and V_t is the price per share of company B at time *t*. The price of Company B's shares will be higher than that of Company A

$$V_t = S_t + \frac{m}{n} W_t \tag{17.2}$$

because some of the value of Company A belongs to the warrant holders and not to the shareholders. Eq (17.1) is valid at any time *t* before the warrants are exercised.

17.2 | Derivatives and Risk Management

But if the warrants are exercised at maturity (time *T*), then the assets of Company A will go up by *mK*, the amount of money paid by the warrant holders to exercise the options. The total value of Company A at this point will, therefore, exceed the value of Company B by *mK* and will be equal to $nV_T + mK$. Company A will then have n+m shares outstanding and no warrants. Therefore, the total value of Company A can also be stated as $(n + m) S_T$. This implies that

$$(n+m)S_T = nV_T + mK \tag{17.3}$$

The price of Company A's shares immediately after exercise is therefore:

$$S_T = \frac{n}{n+m} V_T + \frac{n}{n+m} K$$
(17.4)

The payoff to the warrant holder from exercising the option is $S_T - K$ which can be expressed as:

$$S_T - K = \left(\frac{n}{n+m}V_T + \frac{n}{n+m}K\right) - K = \frac{n}{n+m}(V_T - K)$$
(17.5)

Clearly, the warrant holder will exercise the warrant only if the payoff $S_T - K$ is positive. Eq (17.5) shows that $S_T - K$ is positive if and only if $V_T - K$ is positive. Thus if $V_T > K$, the warrants will be exercised and the payoff is $\frac{n}{n+m}$ ($V_T - K$). If $V_T = K$, the warrants will not be exercised and the payoff

is zero. In all cases, *n* therefore, the payoff can be expressed as $\frac{n}{n+m} \max(V_T - K, 0)$.

Clearly, max $(V_T - K, 0)$ is the payoff of a call option on the shares of Company B with the same exercise price K. Thus the payoff of the *warrants* on Company A is the same as the payoff of call

options on $\frac{n}{n+m}$ shares of Company B.

This gives one method of valuing warrants on Company A:

- 1. Consider Company B which is identical to Company A in all respects, except that it has no warrants outstanding.
- Compute the price of a call option on shares of Company B using the Black-Scholes formula or other method.
- 3. Multiply this option price by the 'dilution factor' $\frac{n}{n+m}$ which can also be written as $\frac{1}{1+m/n}$

Intuitively, this says that warrants are less valuable than ordinary call options on an otherwise identical

company. The difference depends on $\frac{m}{n}$, the ratio of new shares to the number of original shares. The higher this ratio, the greater is the adjustment to be made to the option price to arrive at the warrant price.

It is important to emphasize that in Step 2 above, the option value is the value of an option on Company B shares and not an option on Company A shares. There are two differences between these two options:

- Company B shares are more valuable than Company A shares as shown in Eq (17.2).
- The volatility of Company B shares is higher than that of Company A shares. This is because in the case of Company A, the total risk is borne partly by shareholders and partly by warrant holders.

Both these differences make options on Company B more valuable than options on Company A. It is this significantly higher option price that is adjusted downward by the dilution formula.

It is common to describe the options on Company B's shares as options on the assets of Company A. This is an accurate description if Company A (and therefore Company B) has no debt. In this case, the total value of Company B's shares is the same as the total value of its assets. Options on Company B's shares can, therefore, be regarded as options on its assets. Since the two companies are identical except for the warrants, the options can also be thought of as options on the assets of Company A. This equivalence breaks down when the company has debt. In this case, the options on Company B's shares are not at all the same as options on the assets of Company A.

It is more accurate to describe the warrant valuation model in terms of options on shares of a comparable warrant-free company (Company B). More precisely, the warrant valuation formula states that the value

of a warrant is equal to times $\frac{n}{n+m}$ the value of an option on the shares of a comparable warrantfree company.

A very common mistake in warrant valuation is to compute the value of the warrant as $\frac{n}{n+m}$ times

the value of an option on the shares of the same company (Company A). This is totally incorrect because it uses a wrong share price and a wrong volatility, and generally leads to a serious underestimation of the true warrant value.

17.2.1 Option-Like Valuation of Warrants

What would happen if warrants were valued as if they were options without applying any adjustment for dilution at all? In terms of the three steps of the Warrant Valuation Model, this would mean that neither Step 1 (shifting to a comparable warrant-free company) nor Step 3 (adjustment for dilution) is done. Since these two steps act in opposite directions, the error induced by ignoring both those steps is not as large as it might appear. Clearly, the error caused by skipping both steps is much less than the error caused by skipping only of the steps.

Thus it is far better to treat warrants as if they were options than to apply dilution adjustment to the value of options on the shares of the same company. In fact, it turns out that simply treating warrants as if they are options is a very good approximation. In most real life situations, the approximation is good enough for all practical purposes and today, the view has gained ground that the warrant valuation model is not needed in practice.

To see why this is so, consider a deep-in-the-money warrant just before maturity. At this point, the shareholders of the company know very well that the warrants are almost certain to be exercised. Anticipating this exercise, in an efficient market, the price of the share would also drop down to the 'diluted' level given by Eq (17.4). If one were to value the warrant as if it were an option, one would therefore get the correct value because the dilution is already impounded in the share price.

17.4 | Derivatives and Risk Management

The point that the share price reflects anticipated dilution well before the dilution actually occurs, is very important. In fact, even at the time of the issue of the warrants, the share price would adjust to reflect the probability that dilution would occur. Over the life of the warrant, as the probability of dilution changes, the share price would continually reflect the expected dilution.

Using risk-neutral valuation, we will now show that the option-like valuation is almost correct. In terms of the post-dilution share price S_T , the payoff from the warrant is clearly max $(S_T - K, 0)$ which is the payoff of a call option on S_T . Risk-neutral valuation implies that the warrant value is the discounted value of the risk-neutral expectation of max $(S_T - K, 0)$. In an efficient market, the current stock price S_t is the discounted risk-neutral expectation of S_T . This is a more precise way of saying that the anticipated dilution is impounded in the share price. This suggests that there is nothing wrong with the option-like valuation of the warrant.

The difficulty comes when this is interpreted to mean that the warrant can be valued using Black-Scholes without any adjustment at all. The problem is that Black-Scholes requires a log normal distribution of the stock price and the existence of the warrants changes the distribution of the stock price. Consider for example, a large outstanding issue of deep-out-of-the money warrants. If the company were to do very well, the warrants would move into the money and would be exercised leading to a large dilution. The outstanding warrants significantly reduces the potential upside of the stock since part of the upside is shared with warrant holders.

However, since the warrants are deep-out-of-the money, their current price is quite small and if the stock were to do badly, the warrant price would not drop much further. Therefore, most of the downside falls on stockholders.

The effect is to change the distribution of the stock price in an asymmetric way with the right tail (upside) being reduced more than the left tail (downside). This is only a change in the shape of the distribution since it remains true that the current stock price is the risk-neutral discounted expected value of the post dilution stock price (dilution is impounded in the stock price).

Thus the only problem with option-like valuation of warrants using the Black-Scholes formula is the validity of the log normality assumption. This also explains why the approach works so well in practice. The shape of the distribution is significantly distorted only when the options are deep-outof-the money and when the number of warrants is large as a proportion of the number of shares outstanding. In other cases, the effect is much less pronounced and the option-like valuation is quite robust.

The option-like valuation has thus become the dominant method of valuing warrants. There is however still one situation where the warrant valuation model of the previous section is needed. This is when the company is contemplating the issue of warrants and needs to compute the fair value of the warrants. If the issue of warrants has not yet been announced then even an efficient market cannot impound the expected dilution into the stock price. Thus option-like valuation loses its validity. On the other hand, the warrant valuation model has no difficulty in handling the situation.

17.3 EMPLOYEE STOCK OPTIONS

Though they are called options, employee stock options are actually warrants. They are issued to employees of the company as a component of their total compensation package and there is no explicit cash price for issuing these warrants.

For many purposes, the employee stock options can be valued as warrants. In terms of the above discussion, the option-like valuation is particularly appropriate because in practice, the warrants are issued at-the-money and almost never as deep-out-of-the money. Moreover, since the broad contours of the option plan to be implemented over the course of several years is approved well in advance, the expected dilution is impounded in the stock price well before the actual option grant. The option-like valuation of the warrant is, therefore, quite appropriate.

There are however a few special features of employee stock options that need to be kept in mind:

- Employee stock options are typically non-transferable and therefore completely illiquid.
- Employees can end up holding a very large fraction of their wealth in these options when the company does well. Since they cannot sell the options, the only way for them to diversify their portfolios is to exercise the option and sell the stock. For this reason, many employee stock options are exercised well before maturity though a similar traded stock option would be held to maturity.
- The options are typically contingent on the employee remaining with the company for a specific vesting period. Some options may also be contingent on other vesting conditions like revenue or profit targets being met.

For these reasons, the employee stock options are worth significantly less than traded options with the same maturity and strike price. A common method of valuing the options is to use a option pricing model (binomial or Black-Scholes) with an altered maturity. Instead of using the actual contractual maturity of the option, the expected life of the option based on past experience is used in the pricing model. For example, an employee stock option might allow the employee to exercise the option at any time during the next eight years. The company's historical experience may be, that on an average, the option is exercised within three years. In this case, while using an option pricing model, it may be assumed that the option maturity is three years and not eight.

Valuation of employee stock options has become important because companies are now required to perform this exercise and include the option value as a compensation cost while computing their profitability. Most companies use the Black-Scholes or similar option pricing model with an adjusted option life as discussed above.

Some companies have also attempted a more market based approach in which they offer securities that are closer to employee stock options and use an auction process to determine their market value.

17.4 CONVERTIBLE BONDS AS STRAIGHT BONDS PLUS WARRANTS

A convertible bond is a bond which gives the bond holder the right to convert the bond into shares at a pre-determined price in lieu of redemption in cash. In essence, a convertible bond is a straight bond plus some warrants.

Consider the simplest but quite unrealistic case. Assume that a company issues a three year bond with a face value of Rs 100 that pays an interest rate of 2% annually and is convertible at maturity into shares at a price of Rs 20 per share. This can be thought of as a combination of two instruments:

- A straight (non-convertible) bond with a face value of Rs 100 that pays an interest rate of 2% annually and is redeemed in cash on maturity at the end of three years; and
- Five warrants with an exercise price of Rs 20, exercisable exactly at the end of three years.

17.6 | Derivatives and Risk Management

To verify that the convertible bond is the same as this combination, consider what happens to the convertible as well to the combination of straight bond plus warrants under the two possibilities about the share price:

- 1. If the share price at maturity is Rs 20 or less, the convertible bond holder will not convert the bond and will seek redemption in cash. Under this scenario, the five warrants will also not be exercised and the combination of straight bond plus warrants simply becomes a straight bond that is redeemed in cash.
- 2. If the share price at maturity exceeds Rs 20, the convertible bond holder will convert the bond. At the conversion price of Rs 20, she will get five shares and will not get any redemption in cash. Under this scenario, the five warrants will be exercised and the exercise price will exhaust the entire redemption proceeds of the straight bond. The combination of straight bond plus warrants also ends up as five shares and no cash.

Regardless of what happens to the share price, the convertible bond is the same as a combination of straight bond plus warrants. The advantage of this decomposition is that the convertible can be valued by valuing the two components and adding up the values.

17.4.1 Using Black-Scholes to Value Convertibles

Assume that the market rate of interest for three year government bonds is 6%, the borrowing rate for the company is 8%, the volatility of the company's stock is 25% and the current market price of the company's shares is Rs 18. The combination of straight bond plus warrants can be valued as follows:

1. The straight bond can be valued by discounting at the company's cost of borrowing. This gives the value as

$$\frac{2}{1.08} + \frac{2}{1.08^2} + \frac{102}{1.08^2} = 1.85 + 1.71 + 80.97 = 84.53.$$

The straight bond is worth much less than par because the coupon rate on the bond is so low compared to the cost of borrowing of the company.

- 2. The warrants can be valued using option-like valuation. In this specific example, they can be valued as European call options on the stock with a maturity of three years, exercise price of 20, risk-free rate of 6%, volatility of 25% and market price of 18. The Black-Scholes formula gives the value of the call option as Rs 3.62. Since this is the option-like value of each warrant, the value of the five warrants embedded in the convertible is Rs 18.10.
- 3. The value of the combination of straight bond plus warrants is 84.53 + 18.10 = 102.63. The fair value of the convertible is therefore Rs 102.63. The value of the warrants more than compensates for the low coupon rate of the bond and the convertible is actually worth more than par value.

One of the nice things about the straight bond plus warrants approach is that completely different methods can be used to value the two components, In particular, the straight bond is valuing using a risk-adjusted discount rate—the company's cost of borrowing (the yield to maturity of straight debt issued by the company) exceeds the risk-free rate because it includes a compensation for the credit risk. On the other hand, the warrant is valued using risk-neutral valuation—in the Black-Scholes formula, all discounting is at the risk-free rate.

17.4.2 Binomial Option Pricing Models for Convertibles

Many convertible bonds allow the bond holder to convert at any time during the life of the bond and not just at maturity. The conversion option is then American and not European. What complicates the matter

even more is that conversion is achieved by extinguishing the bond and not paying an exercise price in cash. What the bond holder gives up is not the face value of the bond but the fair value of the straight bond on the date of conversion. Since as in the above example, convertibles usually pay a low coupon rate, the value of the straight bond is significantly less than par at issue and gradually rises towards par value at maturity date.

Thus the effective conversion price increases as the bond approaches maturity. This is one reason why a convertible may be exercised before maturity—it may be better to exercise now at today's effective conversion price than hold on to the conversion option and pay a higher effective conversion price in future. In these situations, it is necessary to use a binomial option pricing model to value the convertible.

Using binomial trees in the straight bond plus warrants approach, presents some complications. First of all as observed earlier, the straight bond part is valued using a risk-adjusted discount rate—the yield to maturity of straight debt issued by the company. The warrant is valued using risk-neutral valuation so that the associated cash flows are discounted at the risk-free rate¹. This means that within the same binomial tree, there are two sets of cash flows which must be kept separate all through the tree so that they can be discounted at two different discount rates.

The binomial tree model is best understood using an example:

- The maturity of the bond is five years
- The face value of the bond is Rs 100
- It pays an annual coupon equal to 2% of the face value
- The bond can be converted into shares at any time during the life of the bond
- The conversion price is Rs 20 per share
- The current market price of the share is Rs 18
- The share can be assumed to have an annual dividend yield of 0.5%
- The risk-free rate is 6% compounded annually
- The yield to maturity of straight debt issued by the company is 8% compounded annually
- At any time after the first two years, the company can call back the bonds for early redemption provided the share price is at least 50% above the conversion price.

The last element (usually described as a soft call) is a common feature in many convertibles and is designed to force early conversion. It is not an unconditional right to call the bond since it can be exercised only when the share price is quite high. Because the share price is much higher than the conversion price, bond holders will exercise the conversion option as soon as the company calls the bonds. Therefore, the bonds end up being converted instead of being redeemed in cash. The soft call limits the gains that bond holders can make. Instead of waiting for the share price to grow to twice or thrice the conversion price, they are forced to convert when the share price is only one-and-a-half times the conversion price. Note also that the right to call does not exist in the first two years.

Some convertibles also provide a put option to the bond holder to put the bonds to the company for premature redemption at specific points of time. The above example does not contain any put option.

The binomial tree in Figure 17.1 shows the evolution of share prices assuming that the volatility is 25%.

¹ This procedure is valid because the cash flow from the warrant part of the convertible is also a derivative on the stock price. Since risk-neutral valuation can be applied to any derivative, discounting only these cash flows at the risk-free rate is a theoretically correct procedure.

t = 0	t = 1	t = 2	t = 3	t = 4	t = 5
					62.83
				48.93	
			38.11		38.11
		29.68		29.11	
	23.11		23.11		23.11
18.00		18.00		18.00	
	14.02		14.02		14.02
		10.92		10.92	
			8.50		8.50
				6.62	
					5.16

17.8 | Derivatives and Risk Management

Figure 17.1 Binomial tree showing stock prices for convertible valuation. Initial stock price is 18, volatility is 25%, risk-free rate is 6% and dividend yield is 0.5%. In this tree, the probabilities of up moves and down moves are 0.54669 and 0.45331 respectively

To value the convertible, it is easiest to consider the value of one-fifth of the bond. This is because one-fifth of a Rs 100 bond will have a face value of Rs 20 and will therefore be convertible into exactly one share. The value of the convertible can then be obtained by multiplying the value by five. Before presenting the details of all the computations, it is useful to look at the final answer. Figure 17.2 gives the valuation tree for one-fifth of the convertible bond:

t = 0	t = 1	t = 2	t = 3	t = 4	t = 5
					63.23
					(Exercised)
				49.33	
				(Called)	
			38.51		38.51
			(Called)		(Exercised)
		30.09		30.31	
	25.39		24.97		23.51
					(Exercised)
21.03		21.38		21.09	
	18.88		19.28		20.40
		17.88		19.29	(Redeemed)
			18.26		20.40
					(Redeemed)
				19.29	20.40
					(Redeemed)

Figure 17.2 Valuation tree for one-fifth of the convertible bond which has a face value of Rs 20 and is therefore convertible into one share. The tree also shows the actions by the bond holder and the issuer (exercise, redemption and call) at various nodes of the tree. The nodes of this tree correspond to the stock prices tree shown in Figure 17.1

Consider first what happens at maturity (t = 5). If the stock price is less than Rs 20, it is best to redeem the bond in cash. He gets the face value of 20 plus the 2% coupon of 0.40 for a total of 20.40. That is what one can see in the bottom three nodes at t = 5. At higher stock prices, it is better to convert. For example, the top node at t = 5 in Figure 17.1 shows a stock price of 62.83. It is much better to get a share worth 62.83 than to get cash of 20. Adding the coupon of 0.40 to the stock value of 62.83 gives the value of the convertible as 63.23 in the top node at t = 5, in Figure 17.2.

To allow proper discounting, it is necessary to keep track of the warrant component and the straight bond component of the convertible value separately. These components are shown separately in Figures 17.3 and 17.4 respectively.

t = 0	t = 1	t = 2	t = 3	t = 4	t = 5
					62.83
				48.93	
			38.11		38.11
		28.35		29.54	
	20.23		20.33		23.11
17.3		13.11		11.92	
	8.12		6.15		0.00
		3.17		0.00	
			0.00		0.00
				0.00	
					0.00

Figure 17.3 Tree showing only the warrant part of the value of one-fifth of the convertible. Each value in Figure 17.2 is equal to the sum of the corresponding values in this tree and in Figure 17.4

t = 0	t = 1	t = 2	t = 3	t = 4	t = 5
					0.40
				0.40	
			0.40		0.40
		2.55		0.77	
	5.16		4.64		0.40
7.13		8.26		9.17	
	10.76		13.14		20.40
		14.71		19.29	
			18.26		20.40
				19.29	
					20.40

Figure 17.4 Tree showing only the straight bond part of the value of one-fifth of the convertible. Each value in Figure 17.2 is equal to the sum of the corresponding values in Figure 17.3 and this tree

17.10 | Derivatives and Risk Management

Turning now to t = 4 in Figure 17.2, consider the middle node where the value of the whole convertible is shown as 21.09. This is the sum of the values at the corresponding nodes in Figure 17.3 and Figure 17.4 (11.92 and 9.17 respectively). The value of 11.92 in Figure 17.3 is obtained by risk-neutral valuation. The values at t = 5 going forward from this node are 23.11 and 0.00 respectively. Multiplying by the probabilities of 0.54669 and 0.45331 and discounting at the risk-free rate of 6% gives $\frac{0.54669 \times 23.11 + 0.45331 \times 0.00}{1.06} = 11.92$. The value of 9.17 in Figure 17.4 is obtained by discounting at the company's borrowing cost of 8%. The values at t = 5 going forward from this node are 0.40 and 20.40 respectively. Multiplying by the probabilities of $0.54669 \times 0.40 + 0.45331 \times 20.40 = 8.77$. Adding the coupon at t = 4 of 0.40 gives 9.17.

This is what happens at each node except the nodes where the convertible is called. The company can call only when the stock price is at least Rs 30 (50% above the conversion price of 20). This is true at the top nodes of t = 3 and t = 4. For example, in the top node at t = 4, the stock price in Figure 17.1 is 48.93. When the bond is called and converted, the value is 48.93 plus the coupon of 0.40 or a total of 49.33 (see Figure 17.2). Why does the company decide to call the bond? This is because if the bond is not called, the value obtained by the procedure described in the previous paragraph turns out to be 49.47. By calling the bond, the company reduces the value to the bond holder by 0.14 which is the same as reducing the cost of servicing its debt by 0.14. That is why the company must call the bond. Similar analysis determines at which nodes the company should call the bond.

Moving back through the tree in this manner in Figure 17.2, the current (t = 0) value of the bond is seen to be 21.03. This is actually the value of one-fifth of the convertible. Therefore, the value of the full convertible bond (face value of 100) is $5 \times 21.03 = 105.15$.

This completes the valuation of the convertible bond. Of course, in practice, one would not use a tree with only one node per year to value a convertible like this. Using a computer, however, it is quite easy to construct a tree with say 100 nodes a year or 500 nodes in all, to value the convertible more accurately.

17.5^{*} CONVERTIBLE BONDS: RISK NEUTRAL VALUATION OF ALL CASH FLOWS

The straight bond plus warrants approach uses risk-neutral valuation for only the warrant part of the instrument while using risk-adjusted discount rates to value the straight bond. The alternative is to value all cash flows of the instrument using risk-neutral valuation. This requires a risk-neutral valuation of the straight bond.

There are two things that can happen to a risky straight bond. Either it gets paid in full at maturity or the company defaults and there is only a partial payment (or perhaps no payment at all). Risk-neutral valuation involves finding (a) the risk-neutral probability of default, (b) the recovery (partial payment) in case of default, (c) the expected cash flows (probability of no default times full payment plus probability of default times recovery and (d) discounting the expected cash flows at the risk-free rate.

^{*} Optional section may be omitted without loss of continuity.

Consider a one year bond with an 8% coupon and 8% yield to maturity so that the bond trades at par. Suppose that the risk-free rate is 6% and that the recovery in the event of default is 40%. Then the risk-neutral probability of default must be 3.09% because using this probability produces the correct value of the bond as shown below:

	Cash flow	Probability	Expected Value
No default	108.00	96.91%	104.66
Default	43.20	3.09%	1.33
Total			106.00
Present value a	at risk-		100.00
free rate of 6%			

If this way of valuing the straight bond is incorporated into the binomial tree, then all the cash flows of the convertible can be discounted at the risk-free rate in a consistent risk-neutral valuation. The question that arises here is what happens to the stock price when the company defaults on its debt. The simple answer is that the stock price must be zero because if a company has no money to pay its creditors, there is obviously no money left to pay the shareholders. The stock price must therefore be zero.

This suggests a simple modification to the binomial tree by adding one more node (the default node at each of the dates. In this node, the stock price is zero and the cash flow from the straight bond is equal to the recovery. It is conventional to assume that recovery is 40% of the face value plus accruec interest because this is roughly the historical experience during the last several decades in developed countries. The probability of this node is given by the risk-neutral probability of default inferred from bond prices and yields. The probabilities of other nodes must be adjusted so that the entire tree is consistent with the observed volatility of the stock and with the risk-free rate.

Another closely related method is to construct a binomial tree not for the stock price but for the value of all the assets of the company. If the value of the assets exceeds the total debt of the company, then there is no default and the value of the stock can be obtained by subtracting the debt from the value of all the assets. If the asset value is less than the total debt, then there is a default and the stock price is zero. Assuming *a pro rata* distribution to all creditors, the percentage recovery by debt holders is equal to the ratio of the asset value to the total debt. This method also allows a risk-neutral valuation of all the cash flows from the convertible.

The details of how these two methods are implemented is beyond the scope of this book.

17.6^{*} **CONVERTIBLE BOND VALUATION WITH INTEREST RATE UNCERTAINTY**

The preceding analysis has used a binomial stock price tree to determine whether any of the American options embedded in a convertible should be exercised early. In other words, it has been assumed that the decisions made by the bond holder and the company to convert, call or put the bond at differen points of time are driven only by the stock price of the company. In reality, these decisions are also influenced by the interest rate. For example, if interest rates fall sharply after the company has issued the convertible, investors may be reluctant to convert early because they are getting a high interest yield.

^{*} Optional section; may be omitted without loss of continuity.

17.12 | Derivatives and Risk Management

Similarly, the company may be more eager to call the bond and re-finance it at a lower interest rate. The binomial tree takes into account the volatility of the stock price but ignores the volatility of interest rates because it uses the same risk-free interest rate at all nodes of the tree. It also ignores any fluctuations in the company's credit spread (excess of the borrowing cost over the risk-free rate) because it assumes the borrowing cost at all nodes as well.

In practice, the error induced by considering only stock price volatility and ignoring the volatility of the risk-free and risky interest rates is modest because the volatility of stock prices is much larger than the volatility of interest rates. However, for a more sophisticated analysis, it is necessary to consider a tree of interest rates as well. Models which attempt to do this are computationally very demanding and are beyond the scope of this book.

Chapter Summary

The warrant valuation formula states that the value of a warrant is equal to $\frac{n}{n+m}$ times the value of a call option

on the shares of a comparable warrant-free company where n and m are the number of shares and warrants outstanding.

It is a serious mistake to take this as $\frac{n}{n+m}$ times the value of an option on the shares of the same company. Option-

like valuation ignores the difference between warrants and call options and takes the warrant value to be the same as the value of a call option. Option-like valuation can be used only when the warrant issue has already taken place or has been announced and the dilution is therefore impounded in the stock price. If the dilution is already in the stock price, option-like valuation is quite accurate in most cases. The only exception is in the case of a large issue of deep out-of-the-money warrants. In practice, option-like valuation has largely superseded the use of the warrant valuation formula.

Employee stock options are called options but they are actually warrants with restrictions on vesting and transferability. The effect of the transferability restriction is to induce holders to exercise the options early. It is common practice to value them by assuming a lower effective life, based on historical experience instead of using the actual contractual life.

Convertible bonds can be viewed as the sum of straight bonds and warrants. The straight bonds can be valued using risk-adjusted discount rates (the company's cost of debt or the yield to maturity of its bonds). Option-like valuation of the warrant can use the Black-Scholes model if the warrants are European. In practice, however, warrants tend to be American and the convertible also includes call and put options. This requires the use of a binomial tree. The two different valuation methods being used in this tree, require that cash flows from the straight bond and the warrant be kept separately all along the tree so that they can be discounted at the risky rate and the risk-free rate respectively. The value of the convertible is obtained by summing these two components of the instrument.

Methods for doing a risk-neutral valuation of all the cash flows of the convertible have also been developed but these are beyond the scope of this book.

Suggestions for Further Reading

The errors caused by using the equity volatility as an input in the warrant valuation formula are highlighted in:

Berger, W. and Klein, D. (1997) "Volatility of What?; Equity Warrants versus Stock Options", in Nelken, I. Volatility in the Capital Markets, state Of-The-Art Techniques for Modeling, Managing and Trading Volatility, Eric Dobby Publishing, page no. 81-94.

This issues involved Employee stock Options are discussed in:

Hull John and White, Alan (2004) "How to Value Employee Stock Options" *Financial Analysts Journal*, 60 (1), 114-119.

The warrant valuation formula was derived by:

Galai, D and Meir I. Schneller (1978) "Pricing of Warrants and the Value of the Firm," *Journal of Finance*, 33(5) 1333-42.

Option like valuation of warrants is justified by:

Darsinos, Theofanis and Stephen E. Satchell (2002) "The Implied Distribution for Stocks of Companies with Warrants and/or Executive Stock Options", Cambridge Working Paper 217, http://www.econ.cai-n.ac.uk/ dae/repec/cam/pdf/wp0217.pdf

The valuation of convertible bonds is discussed in

Tsiveriotis, Kostas and Chris Fernandes (1998) "Valuing convertible bonds with credit risk", *Journal of Fixed Income,* 8(2), 95-102

17.14 | Derivatives and Risk Management

CASES

Case 17-1 ESSAR STEEL LIMITED (1993)—VALUATION OF EURO CONVERTIBLE BONDS

In July 1993, Essar Steel Limited (then known as Essar Gujarat Limited) became the first Indian company to make an issue of convertible bonds in the international markets. A year earlier, several Indian companies led by Reliance Industries Limited had made equity issues abroad though the mechanism of Global Depository Receipts (GDRs).

The Essar issue was noteworthy for several reasons. Firstly, ever since India's credit rating dropped below investment grade to BB+ in 1991, the ability of Indian entities to borrow in international markets had been regarded as doubtful because many international investors are reluctant to pick up paper which is of speculative grade. The Essar issue indicated that there was an appettite for convertible bonds if not for straight debt. Secondly, the previous equity (GDR) issues had been priced at discounts of 15–20% below the market price of the share prevailing at the Bombay Stock Exchange. By contrast, in the Essar issue the conversion price was fixed at a premium of 5% over the Bombay price. It appeared that a skillful choice of instrument had paid off for Essar.

The existence of various options both to the bondholder and to the company means that to determine whether the company got a good deal, it is necessary to value these options properly using an appropriate option pricing model. In this context, the following questions become relevant:

Maturity	5 years
Currency	US dollars
Coupon	5.5% net of withholding tax
Conversion Price	Rs. $62.21 = $1.98 (5\% \text{ over last 5 days average price})$
Adjustment of Conversion Price	Conversion price will be adjusted to protect the bondholder in the event of issues of bonus shares, convertibles, options etc. or issues of shares at less than the market price.
Conversion Right	Bondholder can opt for conversion at any time during tenure of bond.
Call Option	Company can "call" the bonds at any time after two years if dollar price of Essar shares exceeds 130% of the conversion price. Bondholder can opt for conversion at this point, else the bonds would be redeemed at par.
Put Option	Bondholder can "put" the bonds back to the company for redemption at par after exactly three years.

The terms of the Euro Convertible Bonds (ECRs) issued by Essar were as follows:

Written by Professor Jayanth R. Varma.

Teaching material of the Indian Institute of Management, Ahmedabad, is prepared as a basis for class discussion. Cases are not designed to present illustrations of either correct or incorrect handling of administrative problems. Copyright © 1994 by the Indian Institute of Management, Ahmedabad. Revised in 1997. Reproduced with permission.

- Given that the bond is denominated in dollars and that the conversion price is also fixed in dollars, how is the exchange rate risk to be factored in? In particular, what is the relevant volatility of the share price for valuing the conversion option? In 1992, the entire Indian stock market went through a roller-coaster ride and the volatility of the Essar share prices was as high as 60%, but in the first 6 months of 1993, the volatility was a more subdued 40%. The exchange rate of the rupee against the dollar remained steady during 1993–94 due to the massive purchase of dollars by the Reserve Bank of India at a fixed reference rate.
- 2. In what way, if any, does interest rate risk enter the valuation? The yield on five year US Treasury bonds was about 5% in July 1993 and had risen to nearly 6% in March 1994. In mid 1993, the yield on five year Government of India bonds was about 13.5%, but by March 1994, that yield had dropped to about 12%. What is the relevant risk free rate of interest for valuing the conversion option?
- 3. Some newspapers estimated in mid 1993 that, if Essar were to issue straight US dollar bonds, it would have to pay an interest rate of about 350 basis points (i.e. 3.5%) above the US Treasury rate. But it was by no means clear that even at this yield, straight bonds could actually have been placed in mid 1993. The coupon rate on the ECB was of course much lower than this estimated cost of straight debt. What implication does this have for valuing the bond?
- 4. While fixing the coupon rate, an important consideration appears to have been the desire to ensure that it is not below the current dividend yield (about 5%) on the shares. What is the role of the dividend yield in pricing the ECB?

Required

At what all-in cost did Essar Gujarat get funds through the ECB?

What risks did it incur in the process?

Case 17-2 AOKAM PERDANA BERHAD

Prashant Aiyer was reading yet another gloomy report related to Malaysia in July 1997. From his Singapore office, Prashant and his team ran the ASEAN portfolio of a large London based fund. The last few weeks had been very frustrating for Prashant as a stream of bad news poured in from the whole region. When the Thai baht devalued a few weeks ago, the Malaysian economy did not, at first, appear vulnerable. But within weeks, the Malaysian ringgit too was forced to devalue by 8% to about RM 2.60 /USD. The stock market also suffered, and by July was 15% below the year's high.

The report on Prashant's desk talked about how many foreign investors were looking forward to the put options that they held on the Euro convertible bonds issued by Malaysian companies. Prashant could well understand this as he himself looked at the put option as his only salvation. He was therefore alarmed to read in the report that the put option may spell bankruptcy for financially weak companies like Aokam.

Prepared by Professors Jayanth R. Varma and V. Raghunathan.

Teaching material of the Indian Institute of management, Ahmedabad, is prepared as a basis for class discussion. Cases are not designed to present illustrations of either correct or incorrect handling of administrative problems. Copyright © 2000 by the Indian Institute of management, Ahmedabad.

Reproduced with permission

17.16 | Derivatives and Risk Management

Aokam Perdana Berhad was one of the lemons in Prashant's portfolio. In public, Prashant was quite happy to pass it off as one of the problems that he inherited from his predecessor while taking charge of the portfolio in 1996. In the privacy of his own thoughts, however, Prashant recalled how he had strongly recommended the investment to his boss back in 1994.

At that time, the convertible had looked like a wonderful opportunity to profit from the upside potential of the Malaysian timber industry while enjoying downside protection. The upside potential had vanished long ago as the Aokam stock price declined to a fraction of its 1994 levels. But till today, Prashant thought that the downside protection was intact. The put option allowed him to 'put' the bonds back to the company and demand the principal back along with a decent return. The research report in front of him was telling Prashant that this may no longer be true. Aokam may just be unable to redeem the bonds, and may end up in bankruptcy.

Prashant leaned back and wondered what had gone wrong. He knew very well that there was no use crying over spilt milk, but he was determined to learn from his mistakes. Had he missed out something while recommending Aokam in 1994? Or had a whole series of totally unforseeable disasters overtaken Aokam?

Prashant went back to his files on Aokam and re-read the 1994 prospectus with a view to analysing the whole deal afresh. The highlights of the bonds issue were as follows:

Prashant then quickly leafed through the sections of the prospectus that talked about the company's history, business and products.

Aokam Perdana Berhad, as the company is known today, was originally incorporated in June 1954 as a tin mining company with operations in Thailand. The company went public in 1966 and its shares were first listed in Kuala Lumpur Stock Exchange (KLSE) in 1973. Following the policy of the Government of Thailand to promote local equity participation, a new Thai company, Aokam Thai Limited (ATL) was formed in 1978. The company bought a 40% stake into ATL. In 1984, the company entered into joint ventures with other companies in Thailand for exploration and mining of gold. In 1987, poor tin price and a bleak outlook for the tin industry forced Aokam to sell out its holdings in

Date of issue	June 1994
Amount	US\$ 135 million
Currency	Dual currency bond with principal amount and interest payments
	denominated in dollars and redemption amount fixed in ringgits
Coupon Rate	3.5% payable annually
Maturity	10 years (June 2004)
Redemption Amount	Fixed in ringgits at RM 2575 per USD 1000 of principal amount
	(this represented the spot exchange rate at time of issue)
Conversion Right	Bondholder could convert at any time into equity shares at a
	conversion price fixed in dollar terms at USD 3.91 (RM 19.07)-a
	premium of 15.6 % over the share price at time of issue
Call option to company	premium in year 4, 3% in year 5, 2% in year 6, 1% in year seven
(for details, see Exhibit 1)	and at par thereafter) provided that the share price is at least 140%
	of the conversion price both in ringgit and in dollar terms
Put option to bondholder	Exactly at the end of five years at 125.69% of principal amount.
(for details, see Exhibit 1)	Prashant had noted in the margin that this represented a
	reasonable rate of return given that five year US Treasuries were
	then yielding about 6.8% compounded annually

ATL. In the next two years, that is, from 1987 to 1989, the company invested in other business sectors including gaming and tile manufacturing and acquired interests in other listed companies in Malaysia. In 1990, the company acquired 100% stake in Pembangunan Papan Lapis (Sabah) Sdn Bhd² (PPL), thereby entering wood based industry aggressively. By 1992, the company had shed most of their peripheral investments and wood based industry had become the core business of the group. In November 1991, the company finally changed its name to Aokam Perdana Berhad.

In 1992 APB entered into a joint venture with the Malaysian Company Idris, whereby APB undertook to provide financial and technical assistance and management and marketing services to Idris, which held licences to extract and sell logs from Sagisan concessions in Sabah. In return, Idris agreed to sell 50% of the logs to APB at a price equal to the cost of extraction plus royalties and other dues to the Government and a fixed rent of RM 17.50 per hoppus ton (each hoppus ton is equal to 50 cu. ft. of timber).

The Products

The products of the group fell into three basic types: sawn timber, moulded products and veneer wrapped products. Moulded products are essentially wood cut and moulded into different shapes. Veneer wrapped products are a variation of moulded items produced by combining different kinds of lower grade wood through finger joints, lamination and wrapping of the surface with a layer of high quality veneer. As the new plywood/veneer plant under construction became operational by 1994, the Group expected to produce plywood and veneer in addition to the three products mentioned above. Plywood is a panel product usually made of three or more layers. The group's sales by product categories from 1991 to 1994 were as follows:

Product	1991	1992	1993	1994(estimated)
Air-dried sawn timber	92%	17%	9%	5%
Kiln-dried sawn timber	2%	11%	12%	10%
Air-dried mouldings	5%	47%	15%	10%
Kiln-dried mouldings*	1%	25%	64%	75%

Sales by Product Categories

The Production Process

First, round logs are cut into sawn timber. The sawn timber is then air-dried in order to reduce the moisture content. Kiln-drying accelerates the drying process by controlling the moisture content and air flow in the kilns. The group had 50 kilns of various sizes, with a total capacity to kiln-dry about 8000 cubic metres of timber per month. The air-dried sawn timber is then cut or moulded according to profiles of end-products such as door jambs, frames, or skirtings. Finger jointing is a method used by the Group to join small pieces of wood so as to reduce wastages. The moulded wood is then wrapped with high quality veneer to remove defective appearances on the surface to give uniform colour and texture for the finished product. The chart below is a simple illustration of the integrated production process.

² Sdn Berhad is the Malaysia equivalent of Private Limited.

^{*} include veneer wrapped products.

The McGraw·Hill Companies

17.18 | Derivatives and Risk Management



Since the transportation of logs for the production facilities could be temporarily affected by continuous heavy rainfall, Aokam maintained a stockpile of three to four months of inventory.

Production Facilities

From 1990 to 1993, the Group's production facilities grew tenfold from a saw mill capable of producing basically air-dried sawn timber on 100,000 square feet of factory space in Sabah to a current fully integrated operations on over 1,000,000 square feet of factory space, all of which are held by the Group on long-term leases of 60 to 100 years. The production level during the fiscal year 1992–93 was about 85% of the optimum level and Aokam expected to reach optimum level during 1993–94. The fully integrated process had greatly enhanced the flexibility in product mix, enabling the Group to quickly switch from one timber product to another, thereby avoiding markets with softening demand and/or prices in favour of those with firmer demand and prices. The pace of expansion of the Group's processing capacity is evident from the following figures:

	Numbers in operations			
Process	1991–91	1992–93	1993–94	
Kilns	18	50	50	
Moulders	13	17	7	
Finger-jointers	3	7	7	
Laminators	0	12	12	
Veneer wrappers	2	9	10	
Decorative veneer preparation	0	5	23	
Sanders/coaters	1	4	10	

At the time of the issue, Aokam was nearing completion of a RM 270 million plywood/veneer plant, comprising eight production lines with a capacity of more than 24000 cubic metres per month of plywood and face veneer. The building of the plywood/veneer plant was part of APB's strategy to participate in the development of the wood processing industry in China.

The new plywood/veneer plant was expected to provide further synergy to the Group's existing production facilities by supplementing the existing supply of veneer required for its veneer wrapped products. The veneer used by PPL at the time was being sourced locally. It was estimated that the Group's monthly production of door jambs would double from 300,000 when the new plant came into production.

Sources of Supply

The Group did not own any timber plantations or logging rights. Their sources of timber supply were two, namely, the 30,000 acre Petramas and the 410,000 acre Sagisan concessions. The Sagisan concessions were held by Idris and the concessions were divided into four regions of different sizes, namely, Tenju, Bintang, Resolute and Sabakina. As stated earlier, Idris had agreed to sell 50% of the logs to PPL at a price equal to the cost of extraction plus royalties and other dues to the Government and a fixed rent of RM 17.50 per hoppus ton (each hoppus ton is equal to 50 cu. ft. of timber). Regarding the Petramas concessions was held by a third party. At the time the propagetus was issued, the extraction licences to the Petramas and Sabakina concessions had expired and application for their renewal had been made. The log supply contract in respect of Petramas concession was to expire in December 1994. The Group was of the opinion that based on the existing and planned production capacity, the log output from the concessions will be sufficient to meet all its current and future production requirements.

Sales & Distribution

The marketing strategy of the Group had been changing over the years as the Group shifted its mix increasingly in favour of more and more value added products. In the last two years, the Group had been establishing distribution channel for its new veneer wrapped products in North America and Europe where significant potential market existed for standardised products like flat door jambs, split door jambs and skirtings, in fact, US was a major export market for the Group.

17.20 | Derivatives and Risk Management

The Group sold directly to bulk distributors, who, in turn, sold to house developers, door manufacturers, 'do it yourself shops and other end-users. All the distributors had established marketing network in their respective territories. The Group had its own marketing team which visited end users together with distributors.

The plywood and veneer to be produced was mainly targeted for Japan and China. Other than China which was a new market, Aokam was to use its existing distribution network for sales into Japan.

The export sales of the Group by product type and the market for the three years ending June 30, 1993 and six month ending December 31, 1993 are shown below.

		1991 1992				1993			
	Sawn Timber	Moulded Products	Veneer Wrapped Products	Sawn Timber	Moulded Products	Veneer Wrapped Products	Sawn Timber	Mouded Products*	Veneer wrapped Products
Europe/North America	15%	0%	0%	1%	4%	1%	0%	0%	36%
South Korea	29%	1%	0%	11%	34%	0%	14%	17%	0%
Taiwan	0%	5%	0%	0%	23%	0%	0%	15%	0%
Japan	5%	0%	0%	3%	10%	0%	6%	7%	0%
Others	45%	0%	0%	13%	0%	0%	1%	4%	0%
Total	94%	6%	0%	28%	71%	1%	21%	43%	36%

Industry Outlook and Competition

The timber industry was Malaysia's second largest export earner after petroleum and gas. Timber accounted for RM 10.4 billion of the export earnings of Malaysia in 1992. In fact, Malaysia was one of the world's largest exporters of tropical logs, sawn timber and veneer. The timber export in Malaysia was rapidly moving towards more and more of value added products. The Government of Malaysia was also aiding the practice of sustainable forest management and undertaking of reforestation projects.

		6 months ending Dec 31 1993			
	Sawn Timber	MouldedProducts*	Veneer Wrapped Products		
Europe/North America	0%	0%	38%		
South Korea	11%	25%	1%		
Taiwan	1%	1%	0%		
Japan	7%	12%	0%		
Others	1%	3%	0%		
Total	20%	41%	39%		

The export of round logs from Peninsular Malaysia was banned in 1985. In 1993, the export of round logs from Sabah was banned. In addition, the Government of Malaysia levied a duty which decreases in relation to the higher value added downstream products. No duty is levied for export of veneer wrapped products.

The following figures show the shift of Malaysia's timber exports from low value-added to high value-added products.

Product	1989	1990	1991	1992	1993 (Jan-Aug)
Unprocessed wood*	83%	80%	75%	69%	61%
Processed Wood	17%	20%	25%	31%	39%

On account of demand-supply gap, the price of timber had been rising in 1992 and 1993. The gap had presumably arisen from:

- 1. Sabah's decision to reduce round log exports and eventual ban in 1993.
- 2. Hurricane Andrew which hit the USA in late 1992 causing considerable damage to homes mainly made of wood, and
- 3. President Clinton's plan to restrict logging in Northwest America by 50% in 1993 and 70% by 1995 to protect endangered bird species like spotted owl.

The chart below shows some of the average timber product prices from 1987 to 1993.



Average timber product prices

By the second half of 1993 the timber prices had softened, primarily on account of the austerity measures in China and bad weather in the USA which affected building activities. However, the prices were expected to recover in 1994. The long term price prospects for timber were considered fairly good, given the world-wide crunch of timber resources.

The Group exported 100% of its wood products. Until the early 1990's, Japan, South Korea and Taiwan dominated the market for downstream processed timber products. However, rising costs of production, primarily land and labour, have rendered their products less competitive. Indonesia had focusing mostly on plywood production and less on timber products.

In Malaysia, the Group enjoyed the following advantages:

- 1. Assured long term supply of logs
- 2. Many other timber plants were operating below their installed capacity because of the greater competition for logs and the resulting higher input cost.

^{*} Unprocessed wood comprises round logs and sawn timber.

17.22 | Derivatives and Risk Management

- 3. The integrated production facilities of the Group enabled it to improve upon the mix of higher value added products in comparison to timber manufacturers.
- 4. Sabah had an abundant supply of non-export grade sawn timber and odd length timber resulting from the sawing process which could be processed into valuable mouldings. Such timber could not be absorbed in Sabah on account of its small population, but could be easily absorbed for building and construction in Peninsular Malaysia.
- PPL's integrated plant was located close to the forest concessions resulting in low transportation costs.

Capital Expenditure and Finance

The downstream wood processing industry is a capital intensive industry. The following table sets out the Group's capital expenditure from 1991 to 1995:

Year	ExistingProduction (RM Million)	Plywood/Veneer (RM Million)
1991	19	
1992	65	-
1993	48	-
1994 (Projected)	15	120
1995 (Projected)	20	152

The capital expenditure had been financed from internally generated funds, domestic bank borrowings and a five year RM 50 million bond issue in 1992.

Bank borrowings in the form of term loans, trade financing, export credit financing, capital financing and revolving facilities from domestic Malaysian banks amounted to RM 59 million. These borrowings were typically secured by corporate guarantees by the company, legal charges over the Group's landed properties, fixed and floating charges over other assets and liens on deposits.

Most of the projected capital expenditure of 1994 and 1995 was to be made on the construction of the new plywood/veneer plant, the installation of machinery therein and upgrading of facilities in the existing integrated complex. There were also plans to acquire two turbines estimated to cost RM 48 million for the generation of electricity from waste wood.

Profitability

The Group had witnessed tremendous growth in the past two fiscal years. For the year ended June 30, 1993, the turnover of the Group had increased 87% and profit before tax by 75%. For the six months ended December 31, 1993, the turnover and profit before tax had registered an increase of 84% and 104% respectively over the corresponding period of the previous year. The fully integrated production facility had greatly enhanced the Group's flexibility in product mix enabling them to switch quickly from low value added products to high value added ones. Exhibit 2 provides the financial results for the Group (APB consolidated) for 1991–92 and 1992–93.

In the Group's opinion, barring unforeseen circumstances, the trends revealed in the six months ended December 31, 1993 would continue in the second half of the financial year, so that they expected the year ending June 30, 1994 profits to show a material growth over the previous year. The company also expected its new capital expenditure to result in yet higher profits. While the profitability of the
Warrants and Convertibles | 17.23

Group was sensitive to the timber prices which varied greatly in the short term, the directors of the Group believed that in the long term the timber prices would continue to increase, as they had in the past. This, combined with their assured supply of timber led them to believe that the profitability of the Group could only improve further as the timber prices rose. Prashant then turned to his analysis of the major risks involved in subscribing to the bonds.

Exchange Rate

The dual currency bond exposed the investor to exchange risk if the ringgit depreciated against the dollar. This risk would arise only of the bond were held till maturity and not if it were put or converted earlier. The ringgit had depreciated against the dollar in the 1980s, but had appreciated against the dollar in 1992 and early 1993. In an effort to dampen speculation, Bank Negara Malaysia had intervened in the foreign exchange market in late 1993 and early 1994. Exhibit 3 depicts the exchange rates, exchange controls and foreign investment regulation in Malaysia.

Timber Industry Risk

The entire profit of the Group and 99% of its turnover in the year ending 1993 came from the timber sector alone. Timber industry has certain inherent risks of its own. These risk factors include natural hazards of fire and flood, excessive governmental control and changes in environmental and forestry laws and conservation policies. However, the Group did not anticipate any shortage of timber supply in the foreseeable future due to its assumed sources of supply from the Sagisan concessions.

Malaysia Country Risk

Prashant had noted in 1994 that ever since Mahathir became Prime Minister in 1985, Malaysia had established a reputation for political stability, economic dynamism and fiscal prudence. Malaysia's sovereign rating of AA reflected this. He had also noted that since Aokam Perdana Berhad exported most of its production, its dependence on the Malaysian economy was muted.

Management

Prashant's notes of 1994 were upbeat about Aokam's management. Aokam was widely regarded as a technological pioneer not only in Malaysia but also in the Far East. It reduced wastage from 40% to less than 25% by utilising defective raw material to produce high quality veneer wrapping mouldings.

Finger jointing to increase length, laminations to extend width and thickness and veneer wrapping to remove defective appearance were crucial ingredients of this technology. Aokam had linkages with regional universities and research organisations, and many of its technical staff were trained in timber factories in Europe and US.

Pricing

Prashant had initially been put off by the fancy price-earnings multiple (of 92—93 earnings) at which Aokam was trading in Kuala Lumpur. However, eleven months of 93–94 were already over and judging from the first half results, the P/E multiple may be only in mid-teens. The P/E multiple would be even more attractive on prospective 94–95 earnings. Prashant also examined the issue of earnings dilution. Taking into account a recent 2:5 bonus issue, there would be 183 million shares. The shares to be issued if all Euro bonds converted would be only 10% of this. More importantly, the holders of the

17.24 | Derivatives and Risk Management

domestic bonds held detachable warrants or Transferable Subscription Rights (TSRs) for 28 million shares at an exercise price of RM 1.14 per share. The share price information and the KLSE Index from 1989 to 1994 has been shown in Exhibit 4.

Downside Protection

Above all, Prashant's notes of 1994 emphasised the downside protection provided by the put option.

Upside Potential

Prashant's 1994 notes mentioned the earnings growth from large capacity additions, the immense potential of the Chinese market and the impressive historical earnings growth. He saw technology and raw material supply as powerful entry barriers that lent sustainable competitive advantage to Aokam.

Before conducting a critical appraisal of his own analysis of three years ago, Prashant quickly leafed through the subsequent developments.

Shortly after the successful Euro-issue, Aokam reported its results for the year ended 30/6/94. At RM 106 million, profits were almost double that of the previous year. But trouble began next year, when profits dropped to RM 46 million. Aokam's directors attributed this to rising log prices not matched by an increase in the prices of timber products. This led to a squeeze in margins and a sharp drop in profits despite rising volumes from the new capacities coming on stream. As the same trend continued, Aokam reported a whooping loss of RM 145 million for the year ended 30/6/96. That included a write-off of RM 73 million of money advanced to log suppliers. Since these suppliers had failed to deliver the logs in the last two years, Aokam was forced to write-off these dues. "So much for assured supply of logs", muttered Prashant under his breath. Exhibit 5 shows that the financial results of APB for the year ending June 1995 and' 96.

Even this was not the end of Aokam's troubles. In the second half of 1996, the Keningan area of Sabah where Aokam's plants were located was ravaged by floods. Bridges and road links were destroyed and for several months it was impossible to move any logs into the plant or to move the plywood out. In January 1997, the Malaysian rating agency (RAM) downgraded Aokam's domestic bond issue from A2 to C1. The C1 rating indicated a high probability of default. Prashant realised with a sigh that had the Euro bonds been rated, they too would probably merit a similar rating today. It was small consolation to Prashant that three months ago, Aokam redeemed the domestic bonds on maturity. After the maturity of the domestic bonds, Aokam was no longer rated by RAM.

Finally, Prashant took a quick look at a recent stock price chart. Except for a brief upward excursion in March 1997, Aokam's share price had, in recent months, fluctuated around RM 3.00.

And just as Prashant was about to put down the old prospectus of Aokam wearily, the name of DATUK Samsudin Abu Hassan caught his eyes. The prospectus showed Abu Hassan (39) to be the Executive Vice Chairman of Aokam, with 8.43% stockholding. Now why did that name ring a bell in his mind? After some jogging of memory, he recalled having come across that name somewhere in the internet barely a week ago. Surprised at this coincidence and determined to put his mind to rest, he managed to finally catch up with DATUK Hassan on the information highway, somewhere on Asia.axcess.com.

It was a news report according to which Abu Hassan was settled comfortably in south Africa with Datin Melleney Samsudin, though reportedly still active on the Malaysia corporate scene. According to an interview with Star Business, he was said to have submitted an application to the Malaysian authorities

Warrants and Convertibles | 17.25

for floatation of Merbok Hilir Bhd on the KLSE main board. He was quoted as saying, "when I got out (sold his stake) of Aokam Perdana Berhad, I ventured into Merbok, over four years ago. Merbok is a dynamic company with high growth potential." Prashant also found from the website that Merbok was apparently the world's largest manufacturer of rubberwood medium-density fibreboard and that Abu Hassan held 55% stake in Merbok, while 35% was held by some Japanese. In the year ending June 1997, Merbok was expected to chalk up profits upwards of RM 40 million. Hassan was quoted as saying "90% of our products are for export and our major markets are Japan, South Korea, China, Hong Kong, Taiwan, and Singapore. We also export to Europe, the United States and the Middle East."

Exhibit 1 Redemption at the Option of the Company

At any time after 27th June, 1997 the Company may, having given not less than 30 nor more than 60 days' notice to the Bondholders and the Trustee (which notice will be irrevocable), redeem all or some only of the Bonds (being US\$ 1,000,000 in principal amount or an integral multiple thereof) at the following redemption prices per Bond (expressed as percentages of the principal amount of the relevant Bond): and thereafter at their principal amount in United States dollars, with interest accrued to the date of redemption. However, no such redemption may be made unless the closing price of the Shares (as derived from the daily quotations sheet of KLSE or, as the case may be, the equivalent quotations sheet of an Alternative Stock Exchange) for each of 30 consecutive dealing days the last of which occurs not more than 30 days prior to the date upon which notice of such redemption is published, is at least 140 per cent, of the Conversion Price then in effect. Furthermore, no such redemption may be made unless the closing price, translated into US dollars at the prevailing rate described below, of the Shares for each of 30 consecutive dealing days, the last of which occurs not more than 30 days prior to the date upon which notice of such redemption is published, is at least 140 per cent, of the Conversion Price then in effect translated into US\$ at the rate of RM 2.575 = USS 1. Notwithstanding the foregoing sentence, the Company may redeem all or some only of the Bonds (being US\$ 1,000,000 in principal amount or an integral multiple thereof) pursuant to this condition 8(B) on or after 13th June, 1999 if at least 90%, in principal amount of the Bonds has already been converted, redeemed or purchased and canceled. If there shall occur an event giving rise to a change in the Conversion Price during any such 30 dealing day period, appropriate adjustments for the relevant days approved by the Trustee shall be made for the purpose of calculating the closing price for such days, if no price as aforesaid is reported on the KLSE or, as the case may be, the Alternative Stock Exchange for one or more consecutive dealing days, such day or days will be disregarded in the relevant calculation and will be deemed not to have existed when ascertaining such 30 dealing day period. The 'prevailing rate' for the translation of the closing price of the Share shall be the selling quoted by Bank Negara Malaysia at 9 a.m. (Kuala Lumpur time) on each of the relevant dealing day.

If redeemed during the 12 months beginnings on 27th June in any of the following years:	Redemption price (% of principal amount)
1997	104
1998	103
1999	102
2000	101

17.26 | Derivatives and Risk Management

No notice of redemption pursuant to this paragraph (B) or pursuant to paragraph (D) of this Condition shall be effective if it specifies a due date for redemption falling during the period commencing 45 days and ending 30 days (both inclusive) prior to 13 June, 1999 (and if given shall not be effective). Any notice of redemption given under paragraph (B) or (D) before 13th June, 1999 and specifying a due date for redemption after the 30th day prior to 13th June 1999, shall be without prejudice to the rights of Bondholders under paragraph (C) of this Condition and shall not apply in respect of any Bonds in respect of which the option under paragraph (D) of this Condition shall be or has been exercised.

Upon expiry of any such notice, the Company will be bound to redeem the Bonds to which such notice relates at the price aforesaid applicable at the date fixed for redemption, together with interest accrued to the date of redemption.

Redemption at the Option of the Bondholders

Any Bondholder may, by completing, signing and depositing at the specified office of any of the Paying Agents during normal business hours of such Agent not less than 30 nor more than 45 days prior to the relevant date for redemption a notice of redemption in the form (for the time being current) obtainable from any of the Paying Agents together with the Certificate in respect of the Bonds to be redeemed, require the Company to redeem in US\$ all or some only of the Bonds held by him on 13th June, 1999 at the percentage which results in an amount to be received by the Bondholder (but without prejudice to Condition 7(B)) equals 125.69% of the principal amount of the Bonds.

Any such notice of redemption will be irrevocable unless its revocation is approved in writing by the Company not later than five days prior to the relevant date for redemption of the relevant Bond and will bind the Company to redeem the Bonds to which such notice relates. Certificates will not be returned to Bondholders except in the limited circumstances set out in the Agency Agreement. Not less than 30 nor more than 45 days' prior notice at the commencement of the period for the deposit of a notice of redemption will be given by the Company to the Bondholders.

Financial Results (\$ 000s)	30.06.91	30.06.92	30.06.93
Turnover	29,722	72,431	135,738
Profit/(Loss) before Taxation	15,651	31,042	54,460
Taxation	5,022	1,504	560
Minority Interests			
Profit/(Loss) before Ex-items	10,629	29,538	55,020
Extraordinary Items	4,503	3,536	9,747
Financial Results (\$ 000s)	30.06.91	30.06.92	30.06.93
Profit/(Loss) for Period	15,132	26,002	64,767
Ordinary Total Dividends		3,670	5,780
Ordinary Net Dividends		1,645	2,133
Ordinary Tax Exempt Dividends	2,025	3,647	3,647
			(Contd.)

Exhibit 2 Financial Results

Depreciation	827	2,059	6,478
Total Interest Expense	925	2,265	13,753
Total Fixed Assets	32,649	95,406	136,421
			ŕ
Long Term Investments	13,591	8,337	2,756
Intangible Assets	119,971	121,082	120,936
Total Long Term Assets	166,211	224,825	260,113
Stocks	9,822	14,327	41,538
Trade Debtors & Receivables	8	957	22,605
Cash & Bank Balances	817	32,335	37,835
Other Current Assets	11,898	28,122	54,840
Total Current Assets	22,545	75,741	156,818
Total Assets	188,756	300566	416,931
Trade Creditors & Payables	2,057	496	7,192
Short Term Borrowings	15,451	42,738	71,340
Other Current Liabilities	23,644	13,405	20,506
Total Current Liabilities	41,152	56,639	99,038
Total Long Term Loans	6,524	80,491	78,026
Total Long Term Liabilities	6,524	80,491	78,026
Total Liabilities	47,676	137,130	177,064
Ordinary Share Capital	101,233	101,248	112,211
Total Reserves	39,847	62,188	127,656
Shareholders' Funds	141,080	163,436	239,867
Minority Interests			
Total Liabilities	188,756	300,566	416,931
Purchase of Fixed Assets	33,361	65,321	47,590

Warrants and Convertibles | 17.27

Exhibit 3

Exchange Rates, Exchange Controls and Foreign Investment Regulations in Malaysia

Exchange Rates

The external value of the Malaysian Ringgit is determined by the currencies of countries that are significant trading partners of Malaysia. The Malaysian Ringgit is freely convertible into other foreign currencies (including the US\$) at market rates.

Fluctuations in the exchange rate between the Malaysian Ringgit and the US dollar will affect the US dollar value of the redemption price of the Bonds at maturity, as well as the US dollar equivalent of the Malaysian Ringgit price of the Share on the KLSE.

The following table sets out the high, low and year-end closing exchange rates between the Malaysian Ringgit and the US dollar (in RM per US\$ 1.00) for the years indicated:

RM/USS Exchange Rate		In RM per US\$ 1	1.00
Year	high	low	Year End
1983	2.257	2.365	2.34
1984	2.277	2.440	2.42
1985	2.408	2.601	2.41
1986	2.415	2.690	2.60
1987	2.466	2.604	2.49
1988	2.490	2.710	2.710
1989	2.660	2.766	2.699
1990	2.680	2.732	2.698
1991	2.685	2.794	2.720
1992	2.487	2.722	2.615
1993	2.541	2.698	2.693
1994 (through 2 June, 1994)	2.579	2785	

17.28 | Derivatives and Risk Management

Source: Datastream

Exchange Controls

Malaysia has a liberal and non-discriminatory system of exchange controls. Except as described below and except with respect to transactions with Israel, Serbia, Montenegro and Haiti, Malaysia does not restrict remittances and transfer of money to other countries.

Under applicable exchange control rules, offshore borrowings by a resident from a non-resident in foreign currencies requires the prior approval of the Controller of Foreign Exchange (Central Bank of Malaysia) if the amount in foreign currency exceeds the equivalent of RM 1 million. For foreign borrowings in excess of the equivalent of RM 200,000, the Controller of Foreign Exchange must be informed in writing of the purpose of the loan. In addition, exchange control notification, which may be approved by any commercial bank in Malaysia, is required for remittances and transfers of amounts in excess of RM 50,000 to non-residents. Where such remittance and transfers of amounts pertain to the repayment of foreign borrowings and the payment of interest on such borrowings a commercial bank will need to ensure that such repayment and payment is in accordance with the terms approved by the Controller of Foreign Exchange. The Company has received all necessary approvals from the Controller of Foreign Exchange with respect to the Bonds.

Except as described above, there are no restrictions on the remittance abroad by the Company to non-residents of payments of principal or interest under the Bonds or of dividend payments relating to the Shares.

Foreign Investment Regulation in Malaysia

Acquisition of shares and assets in Malaysia are regulated and monitored by the Foreign Investment Committee (the FIC). The FIC guidelines for the acquisition of shares and assets in Malaysia are policy guidelines and do not have the force of law. Subject to these guidelines, foreigners may acquire shares in a Malaysian company. In the event of an acquisition by any one foreign interest or associated group of 15% or more or an acquisition by foreign interests in the aggregate of 30% or more of the voting shares of a Malaysian company, or in the event the value of the shares acquired by Malaysian or foreign interests exceeds RJVI 5 million, the FIC guidelines require the prior approval of the FIC.

Warrants and Convertibles | 17.29

"Foreign Interests" include voting, equity or other rights in a company or associated group of such rights held by (I) foreign individuals, (ii) companies or institutions incorporated outside Malaysia or (iii) local companies where foreign individuals and/or companies or institutions incorporated outside Malaysia hold more than 50% of the voting rights or have control of the management of the companies by way of joint venture agreements, management agreements, technical assistance agreements or other type of agreements.

Any single person named in a Conversion Notice acquiring Shares pursuant to conversion of the Bonds is required to seek FIC approval if, as a result of such conversion, it would own more than RM 5,000,000 of Shares in value. Otherwise, the Company has received the necessary approval from the FIC with respect to the issue of the Bonds.

Exhibit 4 Share Information

The Shares were first listed on the KLSE in December 1973. The table below sets forth for the period:? indicated the high and low intra-day prices and the average daily trading volume of the Shares and the high and low closing levels of the KLSE Composite Index and average daily trading volume of the KLSE. For information relating to the KLSE and the KLSE Composite Index, see "Regulations and Market for Malaysian Securities". The KLSE Composite Index has been subject to extreme upward or downward movements and the information set forth below is not necessarily indicative of past or future trends.

		Shares		KL Composi	SE te Index	KLSE
	High RM	Low RM	Average Daily Tradiug volume Thousand Units	High	Low	Average Daily Trading volume Million Units
Year Enc	led 31 St Decem	ber, 1989				
Ql	1.61	1.31	248.1	410.46	357.31	24.64
Q2	1.60	1.26	251.5	458.28	411.13	37.10
Q3	1.60	1.27	388.7	496.15	445.12	50.30
Q4	1.51	1.03	487.5	564.65	455.79	56.70
Year End	ded 31 st Decemb	ber, 1990				
Ql	2.50	1.30	2,159.2	622.20	558.07	93.38
Q2	2.05	1.57	541.7	588.09	518.53	36.75
Q3	2.50	1.30	2,509.8	632.22	459.08	51.86
Q4	2.00	1.36	405.6	513.42	464.71	35.05
Year End	ded 31 st Decemb	per, 1991				
Ql	1.94	1.32	906.1	605.85	470.41	87.55
Q2	1.84	1.56	423.4	635.02	572.28	53.12
Q3	1.92	1.39	616.7	615.69	505.47	42.49
Q4	1.88	1.61	225.4	556.73	507.37	37.84
						(Contd.)

Year End	led 31 st Decem	ber, 1992				
Ql	2.42	1.78	1,050.5	619.06	546.63	48.98
Q2	3.36	2.24	938.0	605.71	563.69	43.72
Q3	5.45	3.08	692.9	617.94	556.24	71.84
Q4	8.25	5.00	496.9	660.35	550.59	173.07
Year	Ended 31	st December,	1993			
Ql	11.60	7.35	401.5	644.83	614.28	218.18
Q2	14.20	10.50	726.7	746.29	649.56	510.83
Q3	25.00	13.30	544.9	853.83	714.23	396.83
Q4	31.50	22.80	504.5	1,275.32		
Year	Ended 31	st December,	1994			
Ql	30.50	21.80	644.0	1,332.04		
Q2	21.50*	15.0*	430.1	1,067.75		
(up to a	nd including 2	nd June 1994)				

17.30 | Derivatives and Risk Management

* : adjusted to reflect the bonus issue of Shares announced on 7th May, 1994

Q : Quarters

On 2nd June, 1994, the reported closing price of the Shares on the KLSE was RM 16.50. On 7th May, 1994, the Company announced, *inter alia*, a bonus issue of Shares to Shareholders on the basis of 2 Shares for every 5 existing Shares and the closing price of the Shares on 2nd June, 1994, was an "exbonus" price. On the same day the KLSE Composite Index closed at 955.44.

Exhibit 5: Financial Results

Financial Results (\$ 000s)	30.06.94	30.06.95
Turnover	223,121	249,646
Profit/(Loss) before Taxation	106,459	45,682
Taxation	40	1350
Minority Interests	5	
Profit/(Loss) before Ex-items	106,419	47,037
Extraordinary Items	7,045	11,309
Profit/(Loss) for Period	113,464	35,728
Ordinary Total Dividends	10,255	8,392
Ordinary Net Dividends	408	406
Ordinary Tax Exempt Dividends	9,847	7,986
Depreciation	9,888	23,749
Total Interest Expense	14,436	23,394
Total Fixed Assets	263,634	386,110
Long Term Investments	77	1,268
Intangible Assets	131,826	146,997
Total Long Term Assets	395,537	534,375
Stocks	55,265	87,411
Trade Debtors & Receivables	22,790	34,837
		(Contd.)

Warrants and Convertibles | 17.31

Cash & Bank Balances	332,594	131,609
Other Current Assets	100,454	142,946
Total Current Assets	511,103	396,803
Total Assets	906,640	931,178
Trade Creditors & Payables	1,022	11,111
Short Term Borrowings	98,273	34,989
Other Current Liabilities	28,608	28,564
Total Current Liabilities	127,903	74,664
Total Long Term Loans	405,673	403,897
Total Long Term Liabilities	405,673	403,897
Total Liabilities	533,576	478,561
Ordinary Share Capital	183,819	199,640
Total Reserves	189,245	245,760
Shareholders' Funds	373,064	445,400
Minority Interests	7,217	
Total Liabilities	906,640	931,178
Purchase of Fixed Assets	137,157	146,488

Chapter **EIGHTEEN**

Interest Rate and Currency Swaps

This chapter discusses interest rate and currency swaps which are the most important form of over the counter (OTC) derivatives. It discusses two ways of valuing an interest rate swap. The first method treats it as an exchange of a floating rate bond for a fixed rate bond. This chapter, therefore, discusses the valuation of floating rate bonds as well as fixed rate bonds. This leads to a discussion of par bond yield curves, zero coupon yield curves, and forward interest rates. The other method treats the interest rate swap as a bundle of forward rate agreements. Similarly, a currency swap can be analysed either as an exchange of bonds in two different currencies or as a bundle of forward contracts in currencies.

18.1 THE SWAP MARKETS

Swaps are the most important form of Over The Counter (OTC) derivatives. At the end of 2005, interest rate swaps accounted for over 60% of all OTC derivatives by notional value outstanding and currency swaps for another 3%. Turnover data for OTC derivatives is available only for 2004 and this shows that interest rate swaps accounted for about 26% of all OTC derivatives by turnover and currency swaps for another 1%. Another category of short-term swaps known as foreign exchange swaps accounted for 39% of the turnover of all OTC derivatives. Swaps are used in equity and commodity derivatives as well, but this chapter will not discuss these.

A swap is an agreement between two parties A and B to exchange one set of cash flows for another.



Figure 18.1 A generic swap. Any swap between two parties (A and B) can be represented as an exchange of one set of cash flows for another set of cash flows

Swaps are a very flexible instrument because there is almost no constraint on the nature of the cash flows that are exchanged. The exact amount of the cash flows at future dates need not be known in advance so long as there is a formula or methodology which allows the cash flow to be determined at the date when it has to occur. This means that the cash flow could depend on future interest rates, exchange rates or other market prices.

Interest rate swaps arise when the cash flows are essentially interest payments in the same currency, on the same notional principal calculated using different interest rates. Consider for example, a five year

The McGraw·Hill Companies

18.2 | Derivatives and Risk Management

swap using a notional principal of Rs 100 million. The cash flow from A to B may be a fixed stream of quarterly cash flows (over the next five years or twenty quarters) computed at an annual interest rate of

6% on the notional principal. The cash flow in each quarter will then be $\frac{1}{4} \times 6\% \times 100$ million = 1.5 million.

The cash flow from B to A could be a floating rate payment based on the three month interest rate prevailing at the beginning of the relevant quarter. An example of such a rate is the Mumbai Inter Bank Offer Rate (MIBOR) which is the rate at which banks in India lend to or borrow from each other. There is a different MIBOR rate for each maturity. In the swap in Figure 18.2, the floating rate is set to three month MIBOR which is the rate on a three month loan (or deposit) from one bank to another in Mumbai¹. So if the three month MIBOR at the beginning of the quarter was 4%, the payment at the end

of the quarter from B to A would be $\frac{1}{4} \times 4\% \times 100$ million = 1 million while if in another quarter three

month MIBOR was 8%, the payment would be $\frac{1}{4} \times 8\% \times 100$ million = 2 million.





Currency swaps arise when the two sets of cash flows are in two different currencies. A three-year fixed rate coupon only swap between rupees and dollars on a notional principal of \$10 million could look like this when the exchange rate is Rs 45/\$ and the fixed rates are 5% in dollars and 8% in rupees:





¹ See footnote 2 later in this chapter.

A currency swap is not usually coupon-only; it would normally involve an exchange of principals at the end. In this case, the swap would look like this:



Figure 18.4 Currency swap for three years with quarterly coupon payments and exchange of principals at maturity. Notional amount is \$10 million (equal to Rs 450 million at initial exchange rate). The coupons are both fixed rate—5% in dollars and 8% in rupees

In practice, of course, the two parties would settle only the difference between the two cash flow rather than exchange the two cash flows. For example, in the interest rate swap of Figure 18.2, suppose that MIBOR in a particular quarter is 7.2%. The two cash flows to be exchanged that quarter are:

- From A to B: fixed rate payment of Rs 1.5 million
- From B to A: floating rate payment of $\frac{1}{4} \times 7.2 \% \times 100$ million = Rs 1.8 million

It is much easier for B to pay A the difference of Rs 0.3 million between the two cash flows rather than actually exchange the two cash flows. In practice, therefore, the two cash flows at each payment date are netted and only the differences are settled in all swaps.

This brings up an important fact about swaps—the actual cash flows being exchanged do not matter. The only thing that matters is the the difference between the two cash flows. In fact, only the difference needs to be settled and the absolute magnitudes of the two cash flows is completely irrelevant. This idea can be used to value swaps more easily as discussed later in this chapter.

Swaps are generally designed to be a fair exchange (that is to say have zero value) at inception. At inception, therefore, swap valuation is concerned with finding the payments to be made on one leg to produce a fair exchange against a set of given payments on the other leg. At a subsequent point of time, however, swap valuation is concerned with valuing the swap itself, as both legs have already been negotiated

18.2 USING SWAPS TO HEDGE INTEREST RATE AND CURRENCY RISK

Consider an Indian borrower who has borrowed \$10 million for three years and has to pay interest quarterly at 5% during this period. One way for this borrower to hedge the currency risk is to buy forward contracts for each of the dollar payments that she has to make. Every quarter there is a coupon payment of \$125,000 and at maturity there is an additional redemption payment of \$10 million. If she buys forward contracts for all of these payments, she hedges her dollar obligations entirely into known rupee payments.

18.4 | Derivatives and Risk Management

A simpler way to achieve the hedge is to enter into the swap in Figure 18.4 as party B. This also converts her dollar obligations into rupee obligations. This is because every quarter she will receive \$125,000 from Party A and she will pay \$125,000 to her lenders. These payments cancel out, leaving her without any dollar inflow or outflow. In addition, she will pay Rs 9 million every quarter to Party A and this becomes her know rupee outflow. At maturity, she will receive \$10 million from Party A and she will pay this amount to her lenders leaving her with no net dollar cash flows. She will pay Rs 450 million to Party A and this is her fixed rupee cash outflow. Thus the currency swap is an alternative to a whole bundle of foreign exchange forward contracts in terms of hedging her dollar obligations.

Similarly, an interest rate swap is a way of hedging interest rate risk. Suppose a company has borrowed Rs 100 million for five years under a floating rate loan that requires quarterly interest rate payments based on MIBOR. The company is also worried that its interest bill could shoot up if the floating interest rate rises too much.

One way for the company to hedge this risk is to enter into the interest rate swap of Figure 18.2 as Party A (pay fixed rate and receive floating rate). Under this arrangement, the company receives floating rate from the swap and pays this out to the lenders. The floating rate exposure thus cancels out. The company pays a 6% fixed rate under the swap and this becomes its effective interest cost. In this case, the interest cost is learnt beforehand and is in fact, constant over the entire life of the loan.

It is also possible to hedge the interest rate risk using interest rate futures as discussed in Chapter 4 or using forward rate agreements as discussed in Section 18.7 below. However, the swap market is also much more liquid than a bundle of individual forward rate agreements. For this reason, in practice, hedging of interest rate risk on a floating rate loan, almost always uses a swap.

Because of their higher liquidity, swaps are often used as hedging instruments for managing the interest rate risk in general. For example, swaps can be used to change the duration of a portfolio. It will be shown later in this chapter that the fixed rate side of an interest rate swap has the same duration as that of a bond with the same maturity as the swap that pays coupons at the fixed rate. The floating leg on the other hand has a small duration equal to the period up to the next reset date. This means that entering into an long term interest rate swap as a fixed rate receiver is equivalent to buying a fixed rate asset of long duration. It is at the same time, about incurring a liability which has very low duration. The net impact is to increase the duration of the portfolio.

Similarly, entering into the swap as a fixed rate payer is equivalent to creating a long duration liability and creating a short duration asset, thus reducing the overall duration of the total portfolio. While the same impact of duration adjustment can be achieved with bond futures, it is often more convenient to use swaps in order to achieve it.

Similarly, currency swaps can be used to create or hedge long-term currency exposures without reference to any particular borrowing. If a company has a large amount of recurring inflows in a currency (royalty payments for example), it can use swaps to eliminate the currency risk of these inflows.

Interest rate and currency swaps can also be used to take views on interest rates and currencies respectively. In general, any instrument that can be used for hedging can also be used for speculation.

18.3 INTEREST RATE SWAP AS EXCHANGE OF FLOATING RATE BOND FOR FIXED RATE BOND

The idea that the two cash flows being exchanged can be altered without changing the difference between the two, allows the interest rate swap to be reinterpreted differently.

Consider the interest rate swap of Figure 18.2 once again. Suppose this swap is modified by including an exchange of the notional principal of Rs 100 million at the end of the swap. This makes no difference to the structure of the swap. If A pays B Rs 100 million and B pays A Rs 100 million, these two payments cancel out and there is nothing left to be settled. The only reason to add these payments is to reinterpret the swap. With the exchange of principals included, the swap looks like Figure 18.8. The cash flows of the swap in Figure 18.5 can be reinterpreted as the cash flows from fixed rate bonds and floating rate bonds as shown in Figure 18.6.



Figure 18.5 Interest rate swap with a notional exchange of principal at the end



Figure 18.6 Interest rate swap (with notional exchange of principal at beginning and at end) reinterpreted as exchange of the cash flows of a fixed rate bond for the cash flows of a floating rate bond

A more direct interpretation is that at inception, A gives B a fixed rate bond so that B can get all the cash flows of the floating rate bond and B gives A a floating rate bond so that A can get all the cash flows of the floating rate bond. There are no further exchanges of cash flows. In other words, the interest rate swap is nothing but an exchange of a floating rate bond for a fixed rate bond as shown in Figure 18.7. If this is so, valuing a swap is the same as valuing these two kinds of bonds. For example, if both these bonds are worth par (Rs 100 million), then the swap is a fair exchange because equal values

18.6 | Derivatives and Risk Management

are exchanged. If however, the floating rate bonds are worth Rs 100 million but the fixed rate bonds are worth only Rs 98 million, then the swap is not a fair exchange. It can become a fair exchange if B makes an upfront payment of Rs 2 million to A to compensate for the difference in values.

In practice, interest rate swaps are structured at inception to be an exchange of fair values. This is done by setting the fixed rate of interest at a level which makes their value equal to that of the floating rate bonds and, therefore, makes the swap a fair exchange. However, a party may frequently want to unwind a swap entered into several months earlier. The swap might have been fair exchange at inception, but may no longer be so now. In this case, the two bonds need to be valued and the difference would have to be paid or received as compensation for termination of the swap. In either case, valuing the swap is important and this valuation comes down to valuing the fixed rate bond and the floating rate bond separately.



Figure 18.7 Interest rate swap (with notional exchange of principal at the beginning and at the end) reinterpreted as an exchange at inception of a floating rate bond for a fixed rate bond

18.4 VALUING FLOATING RATE BONDS

In most cases, valuing floating rate bonds becomes quite trivial because the bonds are simply worth par. To see why this is so, consider the five-year floating rate bond that pays quarterly coupons at the floating rate of three month MIBOR². Typically, this means that the coupon payment is made at the end of each quarter using the MIBOR rate prevailing at the beginning of the quarter.

Now consider this bond when four years and nine months have passed and the bond has entered its last coupon period. At this point, the MIBOR rate that would be used for the final coupon at maturity is already known and there is no uncertainty about any of the cash flows. For example, if the three-month

^{2.} In India, many different floating rate benchmarks are used including T-bill yields and the overnight MIBOR rate because of the relative illiquidity of the term money (three-month and six-month MIBOR) market. For example, the floating rate leg of the overnight index swap (OIS) is based on compounding the overnight rate from day to day. For simplicity and consistency with international practices (where the three-month LIBOR is a popular benchmark), this section uses the three month MIBOR to illustrate the valuation principles. The approach works well even if the floating rate reference rate is different from this.

MIBOR prevailing at the end of four years and nine months is 7.2%, the coupon payment at maturity on a Rs 100 bond will be equal to $\frac{1}{4} \times 7.2\% \times 100 = \text{Rs} 1.80$. The total cash flow at maturity including the repayment of the principal of Rs 100 would be Rs 101.80. What is the present value of this cash flow of Rs 101.80 at the end of four years and nine months? The answer is obtained by discounting Rs 101.80

for three months at the three-month interest rate of 7.2%. This is $\frac{\text{Rs } 101.80}{1+7.2\%/4} = \frac{\text{Rs } 101.80}{1.0180} = \text{Rs } 100$. A little reflection would show that this present value of Rs 100 or par value is not a coincidence. Regardless of what the interest rate is at the end of four years and nine months, the value of the bond at that date would be Rs 100 or par. This is simply because the bond is paying interest for the remaining three months of the life of the bond at an interest rate equal to the prevailing three-month rate. This ensures that the bond is worth par as it enters its last coupon period.

Now consider the bond when it enters its last but one coupon period. In other words, consider the bond when four years and six months have passed. Suppose that the bond holder contemplates selling the bond after three months. The previous paragraph has demonstrated that the bond holder can be assured of selling the bond at par at that time. In addition, he will get the coupon for the quarter at the three-month MIBOR prevailing now. The situation is very similar to that at the end of four years and nine months. The only difference is that there the bond holder is expected to get the principal back in the form of a repayment by the issuer of the bond. Here, he expects to get back the principal by selling the bond at par. He is still assured of getting back principal plus three months' interest at the going interest rate. The present value of this must still be at par.

The analysis can be pushed back another quarter now. When four years and three months have passed, the bond holder expects to be able to recover the principal by selling the bond at par after three months. In addition, he would get interest for three months at the interest rate prevailing now for three-month maturity. Again the bond must be worth par at the end of four years and three months.

Arguing in this manner, it is clear that the floating rate bond is worth par at the beginning of every quarter. In particular, the bond is worth par at inception as well.

This analysis guarantees that the bond is worth par at every coupon reset date (beginning of each quarter). Now consider the value on dates in between. Typically, the bond is worth more than par on these dates. For example, at the end of each quarter just before the bond pays its coupon, the value of the bond equals par plus the coupon. After the coupon has been paid, the bond becomes worth par. When the coupon is say two months away, both the coupon and the par value need to be discounted for the remaining two months till the coupon payment date. However, this will usually produce a value above par because the coupon that is due is for the entire quarter while the discounting is only for two months.

The only exception is when interest rates have risen sharply after the last reset date. In this case, the bond will pay a coupon for the entire quarter at the old rate (say 5%) while the discounting is for two months at a much higher rate (say 10%). In this kind of scenario, the bond could be worth less than par. Though the bond is not worth par between reset dates, it is easy to find the value of the bond on any date. This is because the bond is known to be par immediately after the payment of the next coupon and the amount of the next coupon is also known. These cash flows simply have to be discounted at the prevailing interest rate.

All this analysis assumes that the bond has no default risk so that the coupons and redemption can all be discounted at the floating rate which is also assumed to be risk-free. If the bond has some default risk, it needs to pay more than the risk-free floating rate to compensate for this risk. For example, a

18.8 | Derivatives and Risk Management

corporate bond might have to offer MIBOR plus 0.5% to be worth par at inception. Would the bond be still worth par at every reset date? If the company's creditworthiness is expected to remain unchanged all through so that at every stage, the cash flows can be discounted at MIBOR plus 0.5%, then it is still true that the bond would be worth par at the beginning of each coupon period. This is because the next coupon is known to be MIBOR + 0.5% and the discount rate to be used is also MIBOR + 0.5%.

If however, at some stage the company's creditworthiness changes or is expected to change, then the analysis would fail. For example, if the appropriate discount rate changes to MIBOR + 0.75%, the coupon would remain at MIBOR + 0.5% and the bond would be worth less than par. On the other hand, if the appropriate discount rate changes to MIBOR + 0.25%, the coupon would remain at MIBOR + 0.5% and the bond would be worth less than par.

This is not a serious issue for swap valuation because (as discussed in footnote 3 later in this chapter) we regard the floating rate benchmarks like LIBOR and MIBOR to be risk-free. Therefore, we shall continue to assume all through that the floating leg of a swap is worth par at inception and at every reset date.

The fact that floating rate bonds are worth par at the next reset date means that their duration at any point of time is equal to the period remaining up to the next reset date. For example, a floating rate bond that pays interest on January 1 and July 1 has a duration of two months on April 1. This is because, the bond holder is assured of receiving a known value (par value plus known interest) after two months. From a duration point of view, the situation is exactly the same as if the bond holder held a Treasury Bill or Commercial Paper maturing in two months' time.

This makes it clear how a swap changes the duration of the total portfolio. Since a swap is an exchange of a fixed rate bond for a floating rate bond, entering into a swap as a fixed rate receiver is equivalent to buying a fixed rate bond and selling a floating rate bond. The fixed rate bond would have a duration that is less than the maturity of the swap but still relatively long—typically measured in years. The floating rate bond would have a very short duration—typically measured in months. Selling a short duration bond and buying a long duration bond clearly increases the duration of the total portfolio. On the other hand, entering into a swap as a fixed rate payer is equivalent to selling a fixed rate bond and buying a floating rate bond. This act of buying a short duration bond and selling a long duration bond clearly reduces the duration bond selling a long duration bond clearly reduces the duration bond selling a long duration bond clearly reduces the duration bond.

18.5 VALUING FIXED RATE BONDS USING YIELD TO MATURITY (YTM)

In this book, the risk free interest rate has been referred to without worrying too much about what it really is. In valuing bonds, it is necessary to be more precise. The first refinement that needs to be made is that there is a different risk-free rate for each maturity. The risk-free rate for five-year bonds is not the same as that for seven-year bonds and neither of these is the same as the rate on three-month bonds. This situation is typically depicted in the form of a yield curve that shows the interest rates for each maturity.

The most popular yield curve is what is called the par bond yield curve. This answers the question: what interest rate (coupon) must a bond pay for it to be worth par. To say that the five-year par bond yield is 7% means that a bond which pays coupons at the rate of 7% is worth exactly the par value of the bond. It may be readily verified that if all the cash flows of this bond are discounted at 7%, the total of these present values is also par as shown below.

Year	Cash flow	Discount factor	Present value
1	7.00	0.9346	6.54
2	7.00	0.8734	6.11
3	7.00	0.8163	5.71
4	107.00	0.7629	81.63
Total			100.00

This is a general result, true for all par bonds. This means that a par bond of any maturity can be valued by discounting all cash flows at the rate obtained from the par bond yield curve for that maturity. In many situations where a very high degree of accuracy is not required, it is common to assume that any bond (whether par, premium, or discount) of any maturity can be valued by discounting all cash flows at the rate obtained from the par bond yield curve for that maturity. This is often called the Yield to Maturity (YTM) method. Strictly speaking, the YTM method is not completely correct as explained in the next section. However, for many purposes, the extra accuracy obtained by the more rigorous method may not be needed. Discounting all bonds of the same maturity at the same YTM may be adequate in many situations.

All this focussed on risk-free rates. A similar approach can be used for risky fixed rate bonds as well. Just as there is a risk-free par bond yield curve depicting risk-free YTMs for different maturities, there are a series of risky par bond yield curves depicting risky YTMs for different maturities. There is in fact one such curve for each borrower or each class of borrowers of identical creditworthiness. Risky fixed rate bonds can be valued using risky YTMs obtained from the yield curve for that borrower.

Since the floating rate bond is always worth par at inception, valuing a swap is essentially the same as valuing a fixed rate bond. As explained earlier, swaps are usually structured to be fair exchanges at inception. This means that the fixed rate side of the swap is also par at inception. The fixed interest rate that is used to accomplish this is known as the swap rate for that maturity. It is evident then that the swap rates are nothing but a par bond yield curve³.

How would one use the YTM approach to value a five-year swap (with a fixed rate of 6.40%) entered into one year ago? The way to do this is to recognize that the swap is now effectively a four-year swap since one year is already over. Suppose the four-year swap rate today is 6.90%, the cash flows of the fixed rate side of the swap have to be discounted at the YTM of 6.90%. This produces a present value of Rs 98.30 for the fixed rate side of the swap as shown below:

Year	Cash flow	Discount factor	Present value
1	6.40	0.9355	5.99
2	6.40	0.8751	5.60
3	6.40	0.8186	5.24
4	106.40	0.7658	81.48
Total			98.30

^{3.} A few years ago, it was common to regard the swap yield curve as a risky yield curve appropriate to the creditworthiness of the highly rated banks that dominate the swap market. In recent years, the view has gained ground that the swap yield curve is in fact a risk free yield curve. A discussion of these views is beyond the scope of this book.

18.10 | Derivatives and Risk Management

Since the floating side is worth par (Rs 100), the swap is currently profitable to the fixed rate payer—he is paying out only Rs 98.30 worth of cash flows while he is receiving Rs 100.00 worth of cash flows. He will agree to terminate the swap only if he is paid the difference of Rs 1.70 between the these two values. Similarly, the floating rate payer should be willing to pay Rs 1.70 to get out of the swap because he is paying out Rs 100.00 worth of cash flows while receiving only Rs 98.30.

18.6 ZERO RATES AND FORWARD RATES

The YTM method of valuing fixed rate bonds or the fixed rate side of the swap is a good approximation but it is not completely correct. This section analyses why it is not an exact method and describes the correct method of valuing fixed rate bonds.

The correct method of valuing bonds starts with two fundamental principles:

- 1. *The law of one price* as applied to bonds states that a rupee of cash flow at a particular date has only one price regardless of which bond the cash flow comes from. For example, a cash flow of Rs 1,200 three years from today has the same value whether this cash flow is the half-yearly coupon payment on Rs 48,000 of 5% ten year bonds or is the half-yearly coupon payment on Rs 30,000 of 4% five-year bonds or is simply the redemption proceeds of Rs 1,200 of three year zero coupon bonds. To determine the present value of the cash flow, the only relevant information is the date of the cash flow and the amount. The source of the cash flow is irrelevant in determining its value so long as it is a risk-free cash flow.
- 2. *The principle of value additivity* as applied to bonds states that the value of a bond is the sum of the values of each of its cash flows—coupons and redemption.

In this framework, zero coupon bonds are fundamental building blocks in bond valuation. This is because a zero coupon bond has only one cash flow—the final redemption. The price of the zero coupon bond establishes the value of this cash flow and this value can be used to value any cash flow on the same date from any other bond.

For example, suppose the price of a six-month zero coupon bond is 0.975610 and the price of a oneyear zero coupon bond is 0.942596. These correspond to semi-annually compounded⁴ yields of 5% and

6% because
$$\frac{1}{1+0.05/2} = \frac{1}{1.025} = 0.975610$$
 and $\frac{1}{(1+0.06/2)^2} = \frac{1}{1.03^2} = 0.942596$

The price of a one year bond with a 6% coupon paid semi-annually can be determined as follows by listing out the cash flows and multiplying each cash flow by the appropriate zero price. The first cash flow at the end of six months is Rs 3.00 and multiplying this by the zero price of 0.975610 gives Rs 2.9628. Similarly, the cash flow of Rs 103 at the end of one year has a present value of Rs 97.0874. The total comes to Rs 100.0142 which is slightly more than par.

^{4.} Semi annual compounding is the standard convention in bond markets in India and many other countries. For simplicity, this entire section assumes that all interest rates are semi-annually compounded, that all bonds pay coupons semi-annually and that all swaps involve semi-annual payments. In reality, the most common structure of swaps in the US is for semi-annual payments on the fixed rate leg and quarterly payments on the floating rate leg. In India, the floating rate leg of the Overnight Index Swap (OIS) is based on compounding the overnight right on a daily basis. The principles described in this section using semi-annual payments can be used to value swaps with other payment schedules, as well with suitable modifications.

Interest Rate and	Currency Swa	ps 18.11
-------------------	--------------	-------------------

Date	Cash flow	Zero price	Present value
Six months	3.0000	0.975610	2.9268
One year	103.0000	0.942596	97.0874
Total			100.0142

What should the coupon be on a one-year par bond? If the semi-annual coupon on a Rs 100 bond is x, then the present value of the cash flows would be equal to 0.975610 x + 0.942596 (100 + x) = 94.2596 + 1.918206 x. For a par bond, this value must be equal to 100. This gives us:

100 = 94.2596 + 1.918206 x or 5.7404 = 1.918206x $x = \frac{1.918206}{5.7404} = 2.9926$

The semi-annual coupon equals 2.9926 or the annual coupon rate equals 5.9852%. This can be verified as follows:

Date	Cash flow	Zero price	Present value
Six months	2.9926	0.975610	2.9196
One year	102.9926	0.942596	97.0804
Total			100.0000

The difference between the YTM based approach and the zero based approach can now be seen clearly. In the YTM method, all one year bonds can be discounted at the par bond YTM of 5.9852%. However, the yield of the one year zero coupon bond is 6% and not 5.9852%.

To understand what is wrong with the YTM approach, let us look at the par bond yield of 5.9852% from the perspective of zero yields. In this approach, the par bond yield is some kind of a weighted average of the six months and one year yields. The first cash flow of 2.9926 must be discounted at 5% because it occurs after six months and the second cash flow of 102.9926 must be discounted at 6% because it occures after one year. Since the bulk of the cash flow is discounted at 6% the weighted average is very close to 6%, but it is slightly lower because there is a small cash flow that is discounted at 5%. From this point of view, there is nothing fundamental about 5.9852%—it is simply one weighted average of 5% and 6% and it is these rates that are fundamental. Another bond with another coupon will have a slightly different set of weights and therefore a slightly different YTMs are quite close, the YTM method may be a very good approximation, but it is not exact. The exact method is to use zero yields.

In the context of swaps, it is often not necessary to use zero yields because at inception, the fixed leg of a swap is a par bond and the swap yield curve is itself a par bond yield curve. The need to use zero yields, however, does arise in a variety of contexts:

- 1. When a swap is being valued sometime after inception when the fixed rate is no longer the prevailing swap rate for the remaining maturity of the swap.
- 2. When the swap is an off-market swap where the fixed rate is different from the prevailing swap rate and the difference is compensated in some other way.
- 3. When the swap is not a standard swap—for example, an amortizing swap where the notional principal declines over time.

18.12 | Derivatives and Risk Management

In these situations, it is first necessary to extract the zero yields out of the swap yield curve, which is a par bond yield curve. This is done by a process called bootstrapping. The starting point of this process is that the six-month par bond yield and the six-month zero yield are one and the same because when a bond enters its last coupon period, it has only one cash flow left. The last coupon is paid along with the redemption at maturity and therefore it becomes one cash flow at one date.

Suppose for example, that the six-month par bond yield is 5% and the one-year par bond yield is 5.9852% as in the previous example. How can one recover the one-year zero rate of 6% from this data? The six-month par bond yield of 5% is also the six-month zero yield and this implies that the price of the six-month zero is:

 $\frac{1}{1+0.05/2} = \frac{1}{1.025} = 0.975610$. Now the cash flows of the one year par bond can be listed down and partially valued as follows:

Date	Cash flow	Zero price	Present value
Six months	2.9926	0.975610	2.9196
One year	102.9926	?	?
Total			100.0000

The blank cells can be filled up as follows. Since the one-year par bond is worth 100 (par) and the first cash flow at the end of six-months is worth 2.9126, it follows that the remaining cash flow (at the end of one year) must be worth:

$$100 - 2.9126 = 97.0804.$$

Since the cash flow at the end of one year is 102.9926, it follows that the one year zero price is: $\frac{97.0804}{102.0026}$ = 0.942596. The one-year zero yield is then obtained as 6% as shown below: 10

$$\left(1+\frac{r}{2}\right)^{-2} = 0.942596$$
 or $r = 2\left(0.942596^{-\frac{1}{2}}-1\right) = 0.06$

Now suppose that the eighteen-month par bond yield is 6.4732%. The eighteen-month zero yield can be computed using the already known zero prices to fill up part of the following table.

Date	Cash flow	Zero price	Present value
Six months	3.2366	0.975610	3.1577
One year	3.2366	0.942596	3.0508
Eighteen months	103.2366	?	?
Total			100.0000

The two missing cells can be computed in the same way as for the one year bond. Since the eighteen month par bond is worth 100 (par) and the first two cash flows are together worth 3.1577 + 3.0508 =6.2085, it follows that the remaining cash flow (at the end of eighteen months) must be worth 100 - 6.2085 = 93.7915. Since the cash flow at the end of eighteen months is 103.2366, it follows that the eighteen month zero price is $\frac{93.7915}{103.2366} = 0.908510$. The eighteen month zero yield is then obtained as 6.5% as shown below:

$$\left(1+\frac{r}{2}\right)^{-3} = 0.908510$$
 or $r = 2\left(0.908510^{-\frac{1}{3}}-1\right) = 0.065$

Proceeding in similar fashion, the entire zero yield curve can be constructed. The computations can be done even faster by observing that the above computations essentially amount to the following formula:

$$t \text{ year zero price} = \frac{1 - \frac{t \text{ year par yield}}{2} \left(\text{sum of all zero prices up to } t - \frac{1}{2} \right)}{1 + \frac{t \text{ year par yield}}{2}}$$
(18.1)

Eq. (18.1) is called bootstrapping because to compute the *t* year zero price, the formula needs all the preceding zero prices up to $t - \frac{1}{2}$. The process can get started because the six month zero price is given by:

six month zero price =
$$\frac{1}{1 + \frac{\text{six month par yield}}{2}}$$
 (18.2)

For example, in the earlier example, the six-month zero price can be got using Eq. (18.2) as $\frac{1}{1.025} = 0.975610$. After the six-month and one-year zero prices have been computed, the eighteen month zero price can be obtained using Eq. (18.1) as $\frac{1-0.032366 \times 1.918206}{1+0.032366} = \frac{0.937915}{1.032366} = 0.908510$.

Once zero prices are known, it is quite easy to obtain zero yields using the formula

$$\left(1 + \frac{t \text{ year zero yield}}{2}\right)^{-2t} = t \text{ year zero price}$$
(18.3)
$$t \text{ year zero yield} = 2\left([t \text{ year zero price}]^{\frac{1}{2t}} - 1\right)$$

In some contexts it is also useful to proceed from zero yields to forward yields. Consider again, the zero yields of the earlier example: 5% for six months, 6% for one year and 6.5% for eighteen months. Forward rates are essentially risk-neutral expectations of what interest rates will be in future. What is the six-month yield expected to be six months from now? If a person invests for six months now at 5% and then reinvests the proceeds for the next six months at the then prevailing six month rate, she should earn 6% over the whole year. The forward rate f, therefore satisfies the following relationship:

 $\left(1+\frac{0.05}{2}\right)\left(1+\frac{f}{2}\right) = \left(1+\frac{0.06}{2}\right)^2$ implying that the forward rate is 7.0049%. Intuitively, if she earns 5% for the first six months and 7% for the next six months, that is the same as earning 6% for the entire year. The slight difference between 7% and 7.0049% is due to the effect of compounding. What this means is that the one year zero rate that is observed can actually be regarded as a weighted average of the current six month rate and the six month rate that is expected to prevail six months from now. This is a weighted average and not a simple average because of the effect of compounding.

18.14 | Derivatives and Risk Management

Similarly, the six-month rate that is expected to prevail one year from now is given by the relationship

 $\left(1+\frac{0.06}{2}\right)^2 \left(1+\frac{f}{2}\right) = \left(1+\frac{0.065}{2}\right)^3$ implying a forward rate of 7.5036%. Combining this relationship with the earlier relationship for the one year zero rate, it is possible to express the eighteen-month zero yield as:

$$\left(1+\frac{0.06}{2}\right)\left(1+\frac{0.070049}{2}\right)\left(1+\frac{0.075036}{2}\right) = \left(1+\frac{0.065}{2}\right)^3.$$

In other words, the eighteen-month zero yield is a weighted average of the (a) the current six-month interest rate, (b) the six-month interest rate that is expected to prevail six months from now and (c) the six-month interest rate that is expected to prevail one year from now. This kind of relationship is true in general—the zero yield for any maturity is a weighted average of the expected six-month rate over the life of the zero bond. It has already been shown earlier in this section that the par bond yield is a weighted average of the different zero rates that apply to the various coupons that the par bond pays over its life.

Indirectly, therefore the par bond yield is also a weighted average of the expected six month rate over the life of the par bond. In some sense, therefore, it is possible to regard forward rates as fundamental and to regard both zero rates and par bond rates as emerging out of these forward rates through a complex averaging process.

The relationship between zero and forward rates can be stated more conveniently in terms of zero prices. The six-month rate that is expected to prevail *t* years from today which is the forward rate

$$f_{t,t+1/2} \text{ applicable from year } t \text{ to year } \left(1 + \frac{1}{2}\right) \text{ is given by}$$

$$\left(1 + \frac{f_{t,t+1/2}}{2}\right) = \frac{t \text{ year zero price}}{\left(t + \frac{1}{2}\right) \text{ year zero price}}$$

$$f_{t,t+1/2} = 2 \frac{t \text{ year zero price} - \left(t + \frac{1}{2}\right) \text{ year zero price}}{\left(t + \frac{1}{2}\right) \text{ year zero price}}$$

$$(18.5)$$

For example, the forward rate $f_{1,1.5}$ can be computed more easily using this formula as:

$$f_{1,1.5} = 2 \frac{0.942596 - 0.908510}{0.908510} = 2 \times 0.037518 = 0.075036.$$

For simplicity, in this section, we have used semi-annual coupons and semi-annual compounding. However, similar formulas can be obtained with annual, quarterly or monthly compounding. For example, with quarterly compounding, Eq. (18.5) can be modified as follows:

V of final Zero Price Zero Yie ash flow 1.00000 3.31009 00.0000% 0.98372 3.31009 08.1752% 0.96387 3.71379 06.2363% 0.94412 3.87069 94.1877% 0.92332 4.02919 92.1697% 0.94412 3.87069 94.1877% 0.92332 4.02919 92.1697% 0.92332 4.02919 92.1697% 0.92332 4.02919 92.1697% 0.92332 4.20409 92.1697% 0.98266 4.20409 93.1122% 0.882566 4.20409 84.1265% 0.86281 4.36579 82.152% 0.84291 4.31849 82.154% 0.80390 4.41359 82.155% 0.76593 4.49419 78.3030% 0.76593 4.49419 78.3030% 0.74714 4.53529	Coupons PV t last ca 10% 100 00% 100 00% 99 03% 99 03% 88 69% 88 69% 88 69% 88	V of all excep 0.00 3.76 5.81 7.83 7.83 9.88 9.88 9.88 11.88	Final Cash PV of all Flow excep Flow excep 101.6550% 0.00 101.3550% 0.376 101.325% 3.76 102.0100% 5.81 102.0525% 7.83 102.0525% 9.88 102.1225% 11.88 102.1225% 11.88 102.1225% 11.88 102.1525% 11.88 102.1525% 11.88 102.1525% 11.88 102.1525% 11.88 102.1525% 11.88 102.1525% 11.88 102.1525% 13.89 102.1525% 15.87
1.00000 1.00000 00.0000% 0.98372 3.31009 08.1752% 0.96387 3.71379 05.2363% 0.94412 3.87069 04.1877% 0.92332 4.02919 92.1697% 0.92332 4.02919 92.1697% 0.90316 4.11609 92.1697% 0.90316 4.11609 92.1697% 0.90316 4.11609 92.154% 0.88266 4.20409 93.1154% 0.88266 4.20409 94.1223% 0.84291 4.31849 86.1031% 0.84291 4.31849 84.1265% 0.82338 4.36579 82.1548% 0.83390 4.41359 82.1548% 0.803390 4.41359 76.3965% 0.76593 4.49419 76.3965% 0.74714 4.53529	000% 248% 637% 99 846% 878% 878% 878% 878% 878% 878% 878% 87	0.0 1.8 5.8 7.8 9.8 9.8 11.8 13.8	101.6550% 0.0 101.8550% 1.8 101.325% 3.7 102.0100% 5.8 102.0525% 7.8 102.055% 9.8 102.1225% 11.8 102.1225% 13.8 102.1725% 15.8 102.1725% 15.8
00.0000% 0.98372 3.3100° 98.1752% 0.96387 3.71379 96.2363% 0.94412 3.8706° 94.1877% 0.92332 4.0291° 92.1697% 0.92332 4.0291° 92.1697% 0.92332 4.0291° 92.1697% 0.92332 4.0291° 92.1697% 0.92332 4.2040° 93.1154% 0.88266 4.2040° 90.1154% 0.88266 4.2040° 91.1156% 0.88266 4.2040° 92.1697% 0.93316 4.1160° 93.1122% 0.88256 4.2638° 84.1265% 0.82338 4.3657° 82.1548% 0.82338 4.3657° 82.1548% 0.80390 4.4135° 82.1548% 0.76593 4.4941° 78.3030% 0.76593 4.4941° 78.3030% 0.74714 4.5352°	00% 10% 10% 10% 10% 10% 10% 10% 10% 10%	0.00 1.82 5.81 7.83 9.88 9.88 11.88 13.89	101.6550% 0.00 101.8550% 1.82 101.9325% 3.76 102.0100% 5.81 102.0525% 7.83 102.0525% 9.88 102.1225% 11.88 102.1225% 13.89
85.1/52% 0.9058/ 5.715% 96.2363% 0.94412 3.8706% 94.1877% 0.92332 4.02919 92.1697% 0.90316 4.1160% 92.1697% 0.90316 4.1160% 92.1697% 0.90316 4.1160% 92.1637% 0.88266 4.20409 88.1122% 0.88266 4.20409 86.1031% 0.84291 4.31849 84.1265% 0.82338 4.41359 82.1548% 0.80390 4.41359 82.1548% 0.80330 4.41359 82.1548% 0.76593 4.49419 76.3055% 0.74714 4.53529	. 8 8 8 9 9 9 9 9 9 9 9 9 9 9 9 9 9 9 9	1.82 3.76 5.81 7.83 9.88 9.88 11.88 13.89	102.0950% 1.82 101.9325% 3.76 102.0100% 5.81 102.0525% 7.83 102.0526% 9.88 102.1225% 11.88 102.1500% 13.89
94.1877% 0.92332 4.02919 92.1697% 0.90316 4.11609 90.1154% 0.88266 4.20409 91.1154% 0.88266 4.20409 92.1697% 0.88266 4.20409 93.1154% 0.88266 4.20409 94.122% 0.86281 4.26689 86.1031% 0.84291 4.31849 84.1265% 0.82338 4.36579 82.1548% 0.80390 4.41359 82.1548% 0.803309 4.41359 82.1548% 0.76593 4.49419 76.3965% 0.74714 4.53529	23% 23% 46% 99 59% 88 35% 88 88 88	5.813 7.830 9.884 11.887 13.890	102.0100% 5.81 102.0525% 7.83 102.0550% 9.88 102.0550% 9.88 102.1225% 11.88 102.1525% 11.88 102.1576% 13.89 102.15760% 15.87
92.1697% 0.90316 4.11609 90.1154% 0.88266 4.20409 98.1122% 0.88266 4.20409 88.1122% 0.88261 4.26089 84.1265% 0.84291 4.31849 84.1265% 0.82338 4.36579 82.1548% 0.803300 4.41359 82.1548% 0.78485 4.45369 80.22277% 0.78485 4.45369 76.3065% 0.76593 4.49419 76.3965% 0.74714 4.53529	13% 16% 16% 13% 13% 13% 13%	7.830 9.884 11.887 13.896	102.0525% 7.836 102.0950% 9.884 102.1225% 11.887 102.1500% 13.896 102.1570% 13.896
90.1154% 0.88266 4.20409 88.1122% 0.86281 4.26689 86.1031% 0.86281 4.26089 84.1265% 0.84291 4.31849 84.1265% 0.84291 4.31849 82.1548% 0.82338 4.36579 82.1548% 0.80390 4.41359 82.1548% 0.78485 4.45369 78.3030% 0.76593 4.49419 78.3056% 0.74714 4.53529	88% 8% 5% 8% 5%	9.884 11.887 13.896	102.0950% 9.884 102.1225% 11.887 102.1225% 13.896 102.150% 13.896 102.175% 15.877
88.1122% 0.86281 4.26089 86.1031% 0.84291 4.31849 84.1265% 0.84291 4.31849 82.1548% 0.82338 4.36579 82.1548% 0.80390 4.41359 80.2227% 0.78485 4.45369 78.3030% 0.76593 4.49419 78.3030% 0.776593 4.49419	88 9% 5% 8 8 8 8 8 8 8 8 8 8 8 8 8 8 8 8 8 8	11.887 13.896	102.1225% 11.8879 102.1500% 13.8969 102.1775% 15.873
86.1031% 0.84291 4.31849 84.1265% 0.82338 4.36579 82.1548% 0.80390 4.41359 80.2227% 0.78485 4.45369 78.3030% 0.76593 4.49419 76.3965% 0.74714 4.53529	% % %	13.8969	102.1500% 13.8969 102.1775% 15.8735
84.1265% 0.82338 4.36579 82.1548% 0.80390 4.41359 80.2227% 0.78485 4.45369 78.3030% 0.76593 4.49419 76.3965% 0.74714 4.53529	5% &		100 17050% 15 873
82.1548% 0.80390 4.41359 80.2227% 0.78485 4.45369 78.3030% 0.76593 4.49419 76.3965% 0.74714 4.53529		15.873	CIDICI 0/C7/11.701
80.2227% 0.78485 4.45369 78.3030% 0.76593 4.49419 76.3965% 0.74714 4.53529	52% 8.	17.84	102.1950% 17.84
78.3030% 0.76593 4.49419 76.3965% 0.74714 4.53529	73% 81	19.777	102.2138% 19.777
76.3965% 0.74714 4.5352%	10% 7	21.697	102.2325% 21.697
	5% 70	23.603	102.2513% 23.603
74.5039% 0.72850 4.5768%	1% 7/	25.496	102.2700% 25.496
72.6807% 0.71058 4.6080%	3% 7.	27.3193	102.2842% 27.3193
70.8782% 0.69286 4.6396%	% 71	29.1218	102.2983% 29.1218
59.0964% 0.67535 4.6717 ⁹	% 6	30.9036	102.3125% 30.9036
57.3358% 0.65805 4.7042%	% 6	32.6642	102.3267% 32.6642
55.5965% 0.64096 4.7372%	5% 6:	34.4035	102.3408% 34.4035
53.8789% 0.62409 4.77069	11% 6.	36.12	102.3550% 36.12

18.16 | Derivatives and Risk Management

$$f_{t,t+1/4} = 4 \frac{t \text{ year zero price} - \left(t + \frac{1}{4}\right) \text{ year zero price}}{\left(t + \frac{1}{4}\right) \text{ year zero price}}$$
(18.6)

To illustrate all these computations, Table 18.1 computes zero and forward rates going out to 10 years from a hypothetical par bond yield curve.

18.7 INTEREST RATE SWAP AS BUNDLE OF FORWARD CONTRACTS

The forward rates that we introduced in the previous section, point to a different way of looking at interest rate swap. Instead of looking at it as an exchange of a fixed rate bond for a floating rate bond, it is possible to look at the swap as a bundle of forward contracts or more precisely, forward rate agreements.

A Forward Rate Agreement (FRA) is an agreement to make a loan or deposit at some future date at an interest rate agreed upon now. For example, an FRA entered into on 1 January 2008 might fix the interest rate for a six month deposit that will be made on 1 January 2009. What would be the fair interest rate for such an FRA?

The answer is, this would be the forward interest rate defined in the previous section. One way to see that this is so, is to observe that one way in which to replicate this FRA is for the borrower to borrow money for eighteen months on 1 January 2008 at the eighteen month interest rate (for zero bonds) prevailing on 1 January 2008 and lend the money out immediately for twelve months at the twelve month interest rate (for zero bonds) prevailing on 1 January 2009 when the twelve month loan matures and requires repayment six months after that, when the eighteen-month loan matures.

Consider the example in the previous section where the eighteen-month zero price was 0.908510 and the twelve-month zero price was 0.942596. To receive one rupee of cash on 1 January 2009 the amount that the borrower needs to borrow on 1 January 2008 and lend out for one-year is 0.942596 because a one year loan of this amount will have a repayment of rupee one. The repayment that the borrower

needs to make on June 30, 2009 is $\frac{1}{0.908510}$ times the initial borrowing of or $\frac{0.942596}{0.908510} = 1.037518$.

This means that the interest paid on a six month loan of rupee one beginning on 1, January 2009 is 3.7518% implying an annualized interest rate of 7.5036% which is the forward rate computed in the previous section.

In fact, Eq. 18.4 is simply an algebraic statement of how the forward rate agreement can be replicated using two zero coupon bonds maturing at the beginning and at the end of the period of the forward loan.

Now suppose on 1 January 2008, a person wants to enter into a whole series of FRAs beginning on 1 January 2008, 1 July 2008, 1 January 2009, 1 July 2009, 1 January 2010 and 1 July 2010 respectively. One way to structure this is a bunch of individual FRAs, each of which has a different forward rate as computed above. Another way to structure the transactions to is have a bundle of FRAs, all of which have the *same* interest rate, which would be a kind of weighted average of the individual forward rates. If the correct weighted average rate is used, the transaction will still be a fair transaction because for some periods this average rate will exceed the true forward rate while for others the average rate will fall below the true forward rate. The meaning of 'correct weighted average' is that the average is so chosen

that these differences cancel out in present value terms and the composite transaction at a single fixed rate becomes a fair transaction.

In the previous section, it was shown that both zero yields and par bond yields are actually weighted averages of the different forward rates that prevail over the life of the bonds. A little reflection shows that the par bond yield is the correct weighted average to use. A par bond pays the same coupon every six months throughout its life. Clearly in some of these half-years, the coupon will exceed the forward rate for that half-year while in some half-years, the coupon will fall short of the forward rate for that half-year. The bond is still worth par because the excess coupons in some half-years cancels out the shortfalls in other half-years in present value terms. This is exactly the effect that is needed for the bundle of FRAs to be transacted at a single fixed rate.

In this framework it is instructive to look at two different ways of hedging a floating rate borrowing. Suppose a company has borrowed Rs 100 million under a floating rate loan that requires semi-annual interest rate payments, and the company is worried that its interest bill could shoot up if the floating interest rate rises too high. The company, therefore, wishes to hedge its interest payments.

One method it can use is clearly to enter into forward rate agreements for each half-year that fixes the interest rate at the forward rate for that half-year. This would produce an interest cost that is predetermined today, but is not constant over time as the interest bill each half year would effectively be the forward rate for that half-year. Another method of hedging is to enter into an interest rate swap with Rs 100 million notional where the company pays a fixed rate and receives a floating rate. Under this arrangement, the company receives the floating rate from the swap and pays this out to the lenders. The floating rate exposure thus cancels out. The company pays a fixed rate under the swap and this becomes its effective interest cost. In this case, the interest cost is not only known beforehand, but is also constant over the entire life of the loan. The swap market is also much more liquid than a bundle of individual forward rate agreements. For this reason, in practice, the hedging of interest rate risk on a floating rate loan almost always uses a swap.

The advantage of looking at a swap as a bundle of FRAs is that this is a very powerful approach that can be used to understand more exotic swaps. This book, however, does not deal with these exotic swaps and this more sophisticated approach is not really needed. It is far more convenient to treat a swap as an exchange of a fixed rate bond for a floating rate bond.

18.8 CURRENCY SWAPS AS EXCHANGE OF BONDS

Just as an interest rate swap can be considered as an exchange of a fixed rate bond for a floating rate bond, a currency swap can be considered as an exchange of a bond in one currency for a bond in another currency. This is easy only for the standard currency swap that includes an exchange of principals at the end and not for coupon only swaps in which there is no exchange of principals at the end. In an interest rate swap, an artificial exchange of principals could be introduced at the end without any effect on the values because this exchange being in the same currency and on the same date, simply cancels out. In currency swaps, an exchange of principals changes the value of the swap because two notional principals in different currencies that have the same value today, need not have the same value at some date in the future.

A coupon-only swap can also be valued by regarding it as an amortizing bond that combines interest payments and principal redemptions into equal periodic instalments. This kind of bond can also be valued using the zero yield curve in each currency. In fact, the zero yield curve allows any set of known

18.18 | Derivatives and Risk Management

cash flows to be valued by simply discounting each cash flow at its appropriate zero rate and adding up all the present values. By doing this for both currency legs, it is possible to value the swap as well. For simplicity, however, this section focuses on standard currency swaps that include an exchange of principals at the end.

Consider once again the currency swap of Figure 18.4 reproduced once again as Figure 18.8 for convenience. This swap can be considered to be an exchange of the following:

- 1. \$10 million of three-year fixed rate dollar denominated bonds with a coupon of 5% payable quarterly.
- 2. Rs 450 million of three-year fixed rate rupee denominated bonds with acoupon of 8% payable quarterly.

It is clear that these two bonds would produce the cash flows depicted in Figure 18.8.



Figure 18.8 Currency swap for three years with quarterly coupon payments and exchange of principals at maturity. The notional amount is \$10 million (equal to Rs 450 million at initial exchange rate). The coupons are both at a fixed rate—5% in dollars and 8% in rupees

Suppose for a moment that the exchange rate at inception of the swap is Rs 45/\$ and that three-year par bond interest rates in India and the US are 5% and 8% compounded quarterly. In this case, the two bonds being exchanged are both par bonds in their respective countries and the exchange is clearly a fair exchange.

Many currency swaps are indeed of this kind where two fixed rate par bonds in two different currencies are exchanged. The interest rate swap market provides the par bond yields⁵ in different currencies and therefore no additional information is required to construct a currency swap in this manner.

It is also possible to have a swap that exchanges a fixed rate bond in one country for a floating rate bond in another currency. For example, the dollar leg of this swap could simply pay a three-month dollar LIBOR instead of paying the fixed rate of 5%. This essentially makes that leg a floating rate dollar bond and since floating rate bonds are worth par at inception (and at every reset date) this still makes it a fair swap.

^{5.} The rates of 5% and 8% would thus be the rates on the fixed leg of a swap with quarterly payments in the US and India respectively.

It is also possible to swap a floating rate in one country for the floating rate in another currency. For example, the dollar leg of the swap could be three-month dollar LIBOR instead of the fixed rate of 5% and the rupee leg of the swap could be three-month rupee MIBOR instead of the fixed rate of 8%. Again, both bonds would be par bonds and the exchange would be a fair exchange.

It is possible to use the exchange of bonds idea to value swaps with other structures as well. For example, consider an amortizing swap in which the notional principal reduces periodically. By extracting the zero yields out of the swap yield curves in each currency, it is possible to value these fixed rate payments in both currencies and determine the fixed rates that would make the bonds par bonds and thus make the swap a fair swap.

It is also possible to use the same zero yields approach to value off market swaps or swaps entered into the past when interest rates were different from what they are now.

18.9 CURRENCY SWAPS AS BUNDLE OF FORWARD CONTRACTS

Just as it was possible to regard an interest rate swap as a bundle of forward rate agreements (FRAs), it is possible to regard a currency swap as a bundle of currency forward contracts. This is easiest to see when both legs of the swap are at a fixed rate.

An Indian borrower who has borrowed \$10 million for three years at 5% interest payable quarterly can hedge the currency risk by buying forward contracts for each of the dollar payments that he has to make—the coupon payment of \$125,000 each quarter and the redemption payment of \$10 million at maturity.

A simpler way to achieve the hedge is to enter into the swap in Figure 18.8 as party B. Every quarter he will receive \$125,000 from Party A and he will pay \$125,000 to his lenders. At maturity, he will receive \$10 million from Party A and he will pay this amount to his lenders leaving himself with no net dollar cash flows at any time. His true debt service cost is the Rs 9 million he pays every quarter and the Rs 450 million that he pays at maturity to Party A. He has a fixed rupee cash outflow.

Thus the currency swap is an alternative to a whole bundle of foreign exchange forward contracts in terms of hedging his dollar obligations. There is a difference though. While using forward contracts, the forward rate for the first few quarters will likely be close to the current exchange rate of Rs 45/\$ and thus the first few quarterly coupons of \$125,000 will translate into rupee costs of only a little more than Rs 5.625 million, instead of the Rs 9 million that is involved in the swap. However, the three-year forward rate is likely to be significantly higher than Rs 45/\$—say Rs 49/\$—and the final redemption payment of \$10 million will then cost him say Rs 490 million in the forward contract instead of the Rs 450 million in the swap. Thus the swap is not equivalent to the bundle of forward contracts on a date-by-date basis. The two are equivalent only on an aggregate present value basis.

In the currency forward market, nearby maturities are quite liquid but distant maturities are quite illiquid and bid-ask spreads become very high. Thus forward contracts are not a very attractive way of hedging long-term borrowing. The swap market on the other hand, is quite liquid out to 10 years or even beyond. Just as the interest rate swap yield curve could be used to compute implied forward rates, it is possible to use the swap yield curve in two different currencies to compute the implied forward rates for different maturities. This is essentially an application of the covered interest parity formula which says that forward exchange rates are determined entirely by interest rate differentials—the currency with the higher interest rate must depreciate at an annual rate equal to the interest rate differential.

The McGraw·Hill Companies

18.20 | Derivatives and Risk Management

Chapter **Summary**

A swap is an agreement between two parties to exchange one set of cash flows for another. They are generally designed to be a fair exchange (that is to say have zero value) at inception. At inception, therefore, swap valuation is concerned with finding the payments to be made on one leg to produce a fair exchange against a set of given payments on the other leg. At a subsequent point of time, however, swap valuation is concerned with valuing the swap itself.

Interest rate swaps arise when the cash flows are essentially interest payments in the same currency, on the same notional principal calculated using different interest rates typically a floating rate on one leg and a fixed rate on the other leg. This can be interpreted as an exchange of a floating rate bond for a fixed rate bond. The floating rate bond can be shown to be orth par at inception on every reset date. The fixed rate bond can be valued using a YTM taken from a par bond yield curve. It can be more accurately valued using zero yields extracted from the par bond yield curve. An interest rate swap can also be viewed as a bundle of forward contracts or forward rate agreements. The relevant forward rates can also be extracted from the swap yield curve.

A currency swap can be considered as an exchange of a bond in one currency for a bond in another currency. These bonds may be either fixed rate or floating rate bonds. To value the bond, it is sufficient to value the two bonds separately using the swap yield curve in the respective currencies. A currency swap can also be viewed as a bundle of forward contracts on the currencies involved.

Swaps are attractive mechanisms for hedging interest rate and exchange rate risk, as well as for taking positions on interest rates and currencies.

Suggestions for Further Reading

A good introduction to the role of swaps in corporate risk management is provided in:

Goodman, Laurie S. (1990) "The use of interest rate swaps in managing corporate liabilities", Continental Bank Journal Of Applied Corporate Finance, 1990 Winter 2(4), 35–47.

The following paper gives an overview of swaps, their pricing and their economic rationale:

Litzenberger, Robert.H (1992) "Swaps: Plain and Fanciful", The Journal of Finance, 47(3), 831-850.

CASE ON SWAPS

Case 18-1 ARVIND MILLS LIMITED: DOLLAR-RUPEE SWAP

In August 1998, Arvind Mills was contemplating an offer that it had received from ICICI Bank for swapping \$10 million of its \$200 million floating rate dollar bond into fixed rate rupee debt. The terms appeared significantly less favourable than that of a swap that it had done with the same bank two weeks earlier.

Unlike in major international currencies, the swap market in rupees was extremely thin and there was no market rate against which to compare the offer that it had received. Arvind Mills had been scouting around for swap deals for more than a month, and it had already accepted two quotations earlier in the month for an aggregate amount of \$15 million. The offer now under consideration was the third serious quotation that it had received so far. The company wondered how it should evaluate the swap offer and decide whether it was worth hedging the dollar debt at this price.

Background

Arvind Mills Limited, the flagship company of the Lalbhai Group was founded in 1931 in Ahmedabad. In the late 1980s, the company began to transform its business portfolio with a new strategic vision of becoming a global textile conglomerate. Faced with declining prospects in the traditional textile business, Arvind embarked upon the manufacture of denim fabrics in 1987. Over the next ten years, it added capacity at a scorching pace and rapidly became one of the top ten denim producers in the world. With the new product line came a sharp international focus as Arvind exported about 65% of the denim it made. Arvind described its international orientation as follows: "We source raw material from China, raise finance in Europe and America, get technology from Germany and Switzerland, manufacture in India and Mauritius, market to the industry's biggest names in Asia, Europe and America, and publish our Balance Sheet in India."

In the domestic market, Arvind moved into downstream denim products. Its jeans brands—Lee, Flying Machine, Newport and Ruf & Tuf—covered the entire range from the premium segment to the bottom end of the economy segment. Similarly, it moved downstream from shirting into ready-to-wear shirts. The Arrow brand was a market leader in the premium category, and the company also launched Ruggers, a casual-wear brand, and Excalibur, a range of easy-to-care formals in the mid-market segment.

As a result, by the late 1990s, Arvind had one of the widest range in textiles. Nevertheless, denim remained the mainstay of the business as shown below:

Reproduced with permission.

Prepared by Professor Jayanth R. Varma

Case Material of the Indian Institute of Management, Ahmedabad is prepared as a basis for class discussion. Cases are not designed to present illustrations of either correct or incorrect handling of administrative problems.

Copyright ©1999 by the Indian Institute of Management, Ahmedabad.

Arvind Mills: Product Portfolio					
	1997-98		1996-97		
25. 	Rs crore	Percent	Rs crore	Percent	
Cotton Textiles	724.01	81	721.67	84	
Denim	540.67	61	542.88	63	
Cotton shirting	83.00	9	80.86	10	
Other fabrics	3.41	28.09	3		
Cotton yarn	96.93	11	69.84	8	
Branded Garments	145.18	16	117.30	14	
Telecommunications	25.35	3	17.06	2	
Total sales	894.54	100	856.03	100	
Of which Exports	396.79	44	405.07	47	

18.22 | Derivatives and Risk Management

The Asian Crisis and the Denim Glut

In July 1997, Thailand devalued its currency ending a fourteen year long peg to the US dollar. This was the first of a series of currency and financial crises in other Asian countries like Korea, Indonesia and Malaysia. Contagion from this crisis soon affected countries in Latin America and Eastern Europe and at times threatened to cascade into a global financial crisis.

When the Asian crisis occurred, India was in the midst of its own economic downturn which had begun in 1995–'96. Several industrial sectors slipped into recession and a crisis of investor confidence led to sharp drops in stock markets and property markets and a virtual drying up of the primary equity market. At the same time, exports slackened significantly after several years of impressive growth. Currency depreciation in East and South East Asia threatened the competitiveness of Indian exports and made even the domestic Indian markets vulnerable to cheap imports. The Reserve Bank of India permitted the rupee to depreciate moderately during the second half of 1997 (see Exhibit 1). For more than a year the rupee had remained stable against the dollar while the dollar appreciated against most other currencies. At the same time, Indian inflation had been significantly higher than that of its trading partners and a stable exchange rate relative to the dollar implied a significant increase in the real effective exchange rate (see Exhibit 2). A modest depreciation of the currency was therefore thought to be a useful correction.

In early 1998, as the rupee continued to come under pressure, the Government veered around to the view that a further fall of the currency was unwarranted. On January 16, 1998, the Reserve Bank mounted a spirited defence of the rupee by raising the bank rate from 9% to 11%, increasing the repo interest rate from 7% to 9%, stepping up reserve requirements, and cutting refinance facilities. These measures were successful in shoring up the currency for the time being and some of these were subsequently rolled back.

For Arvind, the Asian crisis served to exacerbate the problems of over-capacity that had emerged in the global denim industry. In the last five years new denim capacities had mushroomed in Mexico and Latin America while South East Asian capacities too increased substantially particularly in Indonesia and the Philippines. As a result, the global denim capacity of 4 billion meters far outstripped the global demand of 3 billion meters. As global denim prices declined, Arvind faced falling price realisations in its main product. Moreover, the relative stability of the rupee meant a sharp decline in competitiveness of

Indian companies relative to other Asian manufacturers. For example, at one point, there was a significant spurt in cheap denim supplies out of Indonesia.

Depressed demand and prices for denim coupled with rising cotton prices led to a squeeze in margins, but this was partly compensated by powerful growth in all other product lines. The company achieved an operating profit margin of 18% of operating revenues in 1997–'98 as compared to 22% in 1996–'97 (see Exhibits 3, 4, 5 and 6 for key financial statements of the company). The first quarter of 1989–'99 also showed similar trends with higher denim volume sales coming at the expense of realisations, both in the export and domestic markets. The outlook for the immediate future was that margins would continue to be under pressure with denim prices remaining weak and cotton prices staying firm.

Deepening Crisis in Mid-1998

In the middle of 1998, the global crisis deepened further and the Indian economy was subjected to severe shocks of its own. In April 1998, the Japanese yen came under severe speculative attack setting off alarm bells around Asia. As the yen weakened dramatically between April and June, analysts feared that China would be forced to devalue and that this would set off a further round of devaluations around the region. India might then need to make a substantial correction to its own exchange rate to remain competitive in global markets. Towards the end of June, the yen retraced part of its decline thanks to joint intervention by Japan and the United States, but doubts persisted about the Japanese economy and its currency.

Meanwhile, in May 1998, India conducted a series of nuclear tests to which the United States and other countries responded with economic sanctions. These sanctions led to fears that foreign capital inflows into the currency might dry up putting severe pressure on its balance of payments. The rupee came under renewed attack and attempts by the Reserve Bank to defend the currency were not entirely successful.

The situation worsened in June 1998 as foreign investors reacted negatively to the Central Government's Annual Budget. The Budget was perceived to be insufficiently reform oriented and the subsequent back sliding on some of the budget proposals reinforced this perception. Foreign investors turned heavy sellers in the stock markets. By July, many analysts felt that the balance of payments would be under intense pressure on various fronts.

Around this time, Arvind Mills became concerned about the large dollar debt that it had raised in 1996 and 1997 to fund its expansion plans. In November 1996, it had raised an 8 year amortising dollar loan at 110 basis points over Libor, while in July 1997, it had raised \$125 million for 7 years at 130 basis points over Libor. The floating rate note of July 1997 had a bullet repayment on July 18, 2004 and paid interest half-yearly on 18th June and 18th December each year. In addition, Arvind had taken some other foreign currency loans from banks and financial institutions (see Exhibit 7 for a break-up of the borrowings of the company). With operating margins under pressure, the company could ill afford to absorb any large foreign exchange losses. It started making inquiries about possible ways of hedging its dollar exposure. It also moved to a posture of maintaining a long dollar position on its trade transactions by leaving export receivables uncovered while taking forward cover for its import payables.

In August 1998, Russia devalued its currency and defaulted on its domestic debt. This sparked another round of crisis in emerging market currencies around the world and financial markets even in developed countries reeled under the shock. Renewed fears were expressed about the rupee as well. On August 20, the RBI moved strongly to defend the rupee by raising the Cash Reserve Ratio (CRR) for the banking system from 10% to 11% and increasing the short term repo interest rate from 5% to

18.24 | Derivatives and Risk Management

8%. The rupee responded strongly to these measures and closed at 42.85 to the dollar—a gain of about 80 paise from the previous close of 43.65.

August 1998 also witnessed the huge success of the Resurgent India Bonds (RIBs), a dollar bond offered by the Indian government exclusively to non-resident Indians. Opinion on this was divided with several analysts criticising the RIBs as being a very high cost borrowing while others pointed out that the \$4 billion raised through RIBs would strengthen the balance of payments and bolster the reserves at least in the short run.

The Swap Deals

Arvind's first swap transaction was on August 3, 1998 when it concluded a three year swap with Deutsche Bank for \$5 million. Arvind Mills agreed to pay 15.82% fixed rate in rupees while receiving dollar Libor for the duration of the contract (see Exhibit 8A). Though the swap rate of 15.82% appeared high, Arvind went ahead because that was the only quotation that it had received by then. Two weeks later, it was able to do a similar deal with ICICI bank for \$10 million at a swap rate of 15.20% which appeared quite attractive in comparison to the Deutsche Bank deal (see Exhibit 8B).

On August 28, Arvind had to decide on the third swap quotation that it had received. This offer, also from ICICI Bank was for \$ 10 million at a swap rate of 15.75% (see Exhibit 8C). The fact that the swap rate was almost as high as the first deal with Deutsche Bank was quite disturbing. On the other hand, Arvind knew that it had hedged only \$15 million out of its \$200 million of dollar debt.

In the absence of a market reference point for determining the reasonableness of the swap terms, Arvind could analyse the problem from several different angles. Had anything changed materially between August 3 and August 28 for Arvind to reject a swap deal at a slightly lower swap rate than what it had accepted on August 3? How pressing was Arvind's need to hedge its dollar debt? Could it afford to wait and hope for better quotes in the weeks to come? At what rate would the rupee have to depreciate over the next six years for the swap deal to "break-even" for Arvind in terms of net borrowing costs? How should Arvind factor in the uncertainty over future Libor rates in the analysis? In comparing the proposed swap to earlier deals, how should its longer maturity be taken into consideration? Exhibit 9 provides some information on key interest rates and exchange rates on the three deal dates.

In any case, Arvind would need to decide quickly as the validity of the swap quotation was quite limited—possibly only a few hours.

Exhibit 1 Rupee–Dollar Exchange Rate



:--



Exhibit 2

Exhibit 3 **Arvind Mills Limited: Financial Highlights**

					(KS. III CIDIES
	1997-98	1996-97	1995-96	1994-95	1993-94
Sales & Earnings :					
Sales & Operating Income	928	863	709	569	416
Exports Sales	397	405	338	215	165
Domestic Sales	531	451	362	351	232
Operating Profit (EBIDT)	159	187	103	81	96
Profit before Extraordinary					
items, other adjustments					
& Taxation	110	145	11	106	70
Profit after Taxation	101	127	114	106	70
Dividends	25	45	45	40	24
Dividends %	25%	45%	45%	41%	41%
Assets Employed :					
Fixed Assets - Gross					
- Net	1017	722	810	585	633
	447	492	340	365	262
Capital Work-in-progress	945	390	152	119	43
Investments	197	217	231	278	218
Net Working Capital	853	925	932	896	636
Total Assets	2717	2117	1762	1633	1159

The McGraw Hill Companies

18.26 | Derivatives and Risk Management

Financed by :					
Net Worth	1158	1065	1047	1013	724
Share Capital	139	101	100	100	70
Reserves & Surplus	1019	964	947	913	654
Borrowings	1559	1052	715	620	435
Total Funds	2717	2167	1762	1633	1159
EPS (Rs.)	10.05	12.67	11.40	10.65	10.01
Book Value per Share (Rs.)	115.22	105.85	104.83	101.81	103.41

Arvind Mills Limited: Balance sheet

		(KS. III C10
SOURCES OF FUNDS :	As at 31.03.98	As at 31.03.97
(1) Shareholders' Funds		
(a) Share Capital	139.05	100.55
(b) Reserves and Surplus	1050.98	978.15
TOTAL	1190.03	1078.70
(2) Loan Funds		
(a) Secured Loans	923.74	943.86
(b) Unsecured Loans	635.03	109.05
	1558.77	1052.91
	2748.80	2131.61
APPLICATION OF FUNDS :		
(1) Fixed Assets		
(a) Gross Block	1017.04	810.02
(b) Less : Depreciation	295.32	225.21
(c) Net Block	721.72	584.81
(d) Capital Work-in-Progress	944.54	390.37
	1666.26	975.18
(2) Investments	197.49	216.68
(3) Current Assets, Loans & Advances		
(a) Inventories	222.50	229.36
(b) Sundry Debtors	228.86	137.91
(c) Cash and Bank Balances	48.62	233.26
(d) Other Current Assets	16.96	13.64
(e) Loans and Advances	526.63	502.90
	1043.57	1117.07
Interest Rate and Currency Swaps | 18.27

Less :		
Current Liabilities and Provisions		
(a) Liabilities	147.79	134.67
(b) Provisions	42.75	57.03
	190.54	191.70
Net Current Assets	853.03	925.37
(4) Miscellaneous Expenditure (to the extend		
not written off or adjusted)	32.02	14.38
TOTAL	2748.80	2131.61

Exhibit 5 Arvind Mills Limited: Profit & Account

(Rs. in Crores)

INCOME :	1997-98	1996-97
Sales and Operating Income	928.32	863.15
Other Income	27.72	31.49
Total	956.04	894.64
EXPENSES :		
Raw Materials	439.13	361.85
Purchase of Finished Goods	21.07	8.13
Employees' Emoluments	47.82	41.62
Others	293.29	250.03
Interest (Net)	3.90	15.45
Depreciation	72.73	57.77
(Increase)/Decrease in Stocks	(32.37	14.77
Total	845.57	749.62
Profit before Extraordinary items,		
other adjustments and Taxation	110.47	145.02
Less : Extraordinary items	1.18	12.20
Other adjustments due to change in Accounting Policy	_	48.36
Less : Transferred from General Reserves	-	(48.36)
Profit before taxation	109.29	132.82
Provision for taxation	8.25	5.42
Profit for the year	101.04	127.40
Balance as per last year's Balance Sheet	35.46	38.33
Transferred from Debenture Redemption Reserve	99.50	_

18.28 | Derivatives and Risk Management

Transferred from Investment			
Allowance (Utilised) Reserve	0.81	5.72	
Subtotal	236.81	171.45	
Transfer to Debenture			
Redemption Reserve	46.74	66.52	
Interim Dividend on Preference Shares	0.43	-	
Tax on Interim Dividend	0.04	-	
Proposed Dividend	25.14	44.97	
Tax on Proposed Dividend	2.51	4.50	
Transfer to General Reserve	10.00	20.00	
Balance Carried to Balance Sheet	151.95	35.46	
Subtotal	236.81	171.45	

Exhibit 6 Arvind Mills Limited: Cash Flow Statement

(Rs. in Crores)

		For the year ende 31/3/98	d For the year of 31/3/97	ended
Α.	Cash Flow from Operating Activities:			
	Profit before Extraordinary Items, Other			
-	Adjustments and Taxation	110.47	145.02	
	Adjustment for depreciation and other items	79.68		
	Operating Profit before Working Capital Changes	177.25	210.09	
	Net Change in Working Capital	(176.53)	40.83	
	Income Taxes Paid	(9.84)	(17.08)	
	Cash Flow before Extraordinary Items	(9.12)	233.84	
_	Extraordinary Items:	(17.64)	(14.38)	
3 	Net Cash Flow from Operating			
	Activities		26.76	219.46
B.	Cash Flow from Investing Activities			
-	Purchase of fixed assets	(596.73)	(424.93)	
-	Sale of Fixed Assets	1.01	16.12	
	Sale of Undertaking	0.00	11.73	
	Net Change in Investments	19.86	5.80	
	Change in Loans and Advances	2.21	95.82	
	Interest Income	77.18	101.68	
-	Dividend Income	4.75	1.99	
	Net Cash Flow from Investing			
_	Activities	(4	91.72)	(191.79)

Interest Rate and Currency Swaps | 18.29

C. Cash Flow from Financing Activities			
Issue of Preference Share Capital	38.41	-	
Change in Borrowings	409.68	297.60	
Interest Paid	(107.10)	(130.25)	
Premium on Redemption of Debentures	(0.03)	(0.04)	
Dividends paid	(45.40)	(44.87)	
Tax on dividends paid	(4.54)	-	
Net Cash Flow from Financing Activities		291.02	122.44
Net Increase in Cash & Cash Equivalents		(227.46)	150.11
Add : Cash & Cash Equivalents as on 1/4/1997		279.74	129.18
Cash & Cash Equivalents as on 1/4/1998		52.28	279.29

Exhibit 7 Arvind Mills Limited: Break-up of Borrowings

		(Rs. in Crores)
	As at 31.03.98	As at 31.03.97
Secured	Loans	
Debentures		
Debentures	254.76	44.76
Loans From Financial Institutions		
Rupee Loan	94.80	75.51
Foreign Currency Loan	452.13	300.82
Total	546.93	376.33
From Banks		
Cash Credit and other facilities	122.05	122.59
Add : Interest accrued and due thereon	-	0.18
Total	122.05	122.77
Total Secured Loans	923.74	943.86
Unsecur	ed Loans	
Unsecured Debentures	7.00	7.00
Fixed Deposit		
From Public	14.66	0.07
From Others	2.66	2.19
Total	17.32	2.26
Loans from bank		
Rupee Loans	50.74	30.17
Foreign Currency Loans	544.30	54.16
Total	595.04	84.33
Loans from institutions	4.22	5.22
Loans from others (Rupee Loans)	11.45	10.24
Total Unsecured Loans	635.03	109.05

18.30 | Derivatives and Risk Management

Exhibit 8A **Term Sheet for First Swap Deal** Trade Date 03 August 1998 Effective Date 05 August 1998 Termination Date 18 June 2001 **Fixed Amounts** Fixed Rate Payer Arvind Mills Limited Fixed Rate Payer Currency Amount INR 212,825,000.00 Fixed Rate Payer Payment Dates 18 Jun and 18 Dec each year Fixed Rate 15.82% per annum Fixed Rate Day Count Fraction Actual/365 (Fixed) London, New York and Bombay **Business Days Floating Amounts** Floating Rate Payer Deutsche Bank AG Floating Rate Payer Currency Amount USD 5,000,000.00 Floating Rate Payer Payment Dates 18 Jun and 18 Dec each year Floating Rate for Initial Calculation 5.71875% per annum Floating Rate Option **USD-LIBOR-BBA** Designated Maturity 6 months Spread None Floating Rate Day Count Fraction Actual/360

Exhibit 8B Term Sheet for Second Swap Deal

Trade Date	11 Aug 98	
Effective Date	13 Aug 98	
Termination Date	18 Jun 2001	
Amortisation	Not Applicable	
Party A details		
Party A fixed rate payer		Arvind Mills Ltd
Party A fixed rate payer	currency amount	INR 427,900,000
Fixed rate payer payment	t dates	18 Jun and 18 Dec each year
Fixed rates		15.2% pa payable half yearly
Fixed rates day count fra	ction	ACT/365 (FIXED)
Party B details		
Party B floating rate paye	er	ICICI Banking Corporation Limited
Party B floating rate pay	er currency amount	USD 10,000,000
Floating rate payer paym	ient dates	18 Jun and 18 Dec each year
Floating rate option		USD - LIBOR - BBA
Designated Maturity		6 months
First floating period		13.08.98 - 18.12.98
		5.69167%
Exchanges in principal		18 Jun 2001

Interest Rate and Currency Swaps | 18.31

Spread	None
Rest dates	The first da
Floating rate day count fraction	ACT/360
Compounding	Inapplicab
Business days for resets	LONDON
Business days for payments	Mumbai an
Calculation agent	ICICI Bank

None The first day in each calculation period ACT/360 Inapplicable LONDON Mumbai and New York ICICI Banking Corporation Limited, Mumbai

Exhibit 8C Term Sheet for Third Swap Offer

Trade Date	28 Aug 98	
Effective Date	01 Sep 98	
Termination Date	18 Jun 2004	
Amortisation	Not Applicable	
Party A details		
Party A fixed rate payer		Arvind Mills Ltd
Party A fixed rate payer	r currency amount	INR 425,500,000
Fixed rate payer payme	nt dates	18 Jun and 18 Dec each year
Fixed rates		15.75% pa payable half yearly
Fixed rates day count fi	raction	ACT/365 (FIXED)
Party B details		
Party B floating rate pay	yer	ICICI Banking Corporation Limited
Party B floating rate page	yer currency amount	USD 10,000,000
Floating rate payer payer	ment dates	18 Jun and 18 Dec each year
Floating rate option		USD - LIBOR - BBA
Designated Maturity		6 months
First floating period		01.09.98 - 18.12.98
		5.6250%
Exchanges in principal		18 Jun 2004
Spread		None
Rest dates		The first day in each calculation period
Floating rate day count	fraction	ACT/360
Compounding		Inapplicable
Business days for reset	S	LONDON
Business days for payn	nents	Mumbai and New York
Calculation agent		ICICI Banking Corporation Limited, Mumbai

The McGraw·Hill Companies

18.32 | Derivatives and Risk Management

Key Interest Rates and Exchange Rates on Deal Dates			
	August 3, 1998	August 11, 1998	August 28, 1998
US Treasury yield: 3 year	5.44%	5.30%	4.89%
US Treasury yield: 5 year	5.46%	5.30%	4.94%
US Treasury yield: 7 year	5.51%	5.40%	5.09%
US T-Bill yield: 6 months	5.25%	5.13%	5.01%
USD Libor: 6 months	5.63%	5.59%	5.44%
USD Swap Rate: 3 years	5.88%	5.81%	5.55%
USD Swap Rate: 5 years	5.92%	5.87%	5.68%
USD Swap Rate: 7 years	5.96%	5.94%	5.78%
GOI Bonds yield: 2-3 year	11.46%	11.37%	11.53%
GOI Bonds yield: 5-6 year	11.92%	11.84%	11.91%
Indian Commercial Paper Yield (new issues during the weeks			
ended August 7, 14 & 28 respectively)	9.50-10.00%	9.25-10.25%	12.50-14.00%
Spot Exchange Rate INR/USD	42.55/56	42.86/88	42.54/55
Forward Premium INR/USD:			
3 months	6.60%	7.50%	11.76%
Forward Premium INR/USD: 6 months	7.46%	8.50%	10.99%
Arvind Mills estimated medium term borrowing cost	16.50	% (compounded annua	lly)

Exhibit 9 Key Interest Rates and Exchange Rates on Deal Dates

Chapter **Nineteen**

Caps, Floors and Swaptions

Caps and floors are widely used in floating rate instruments to limit the range of fluctuations in the floating rate coupons. Caps and floors are also traded separately. This chapter discusses the popular Libor Market Model (LMM) that uses the Black-Scholes formulas to value these options on the floating rate. Swaptions are options on the swap rate that are traded in their own right. They are also encountered in the form of callability provisions in many bonds. The popular Swap Market Model (SMM) that uses the Black-Scholes formulas to value these options on the swap rate these options on the swap rate are also discussed. This chapter highlights the differences between caplet volatilities, cap volatilities, and swaption volatilities as well as the market implied correlations that are contained in these quantities.

19.1 CAPS AND FLOORS

Chapter 18 showed that a borrower can use swaps to hedge the risk that an interest rate on a floating rate faces loan would go too high. By entering into a pay fixed swap, he can turn his liability into a fixed rate loan. But this also eliminates the potential benefit to the borrower when interest rates decline and the floating rate loan becomes a low cost borrowing.

Similarly, a floating rate lender can protect herself from declining interest rates by entering into a swap as a fixed rate receiver. But that removes the potential for her to profit when interest rates rise and the floating rate investment becomes very attractive.

Both the floating rate borrower and lender are keen to limit their downside (very high rates and very low rates respectively) while retaining some upside. Interest rate caps and floors are designed to do precisely this.

A floating rate borrower would want an arrangement where he pays floating rate subject to a cap of say 10%. This means that even if interest rates rise to 15% or 20%, he pays only 10%, but if interest rate falls, he pays only the lower rate. Of course, in finance, there is no free lunch, and the borrower who wants a cap will have to pay for it in some form. The price could be an upfront payment, it could be a higher interest rate (he may agree to pay half percent more than the floating rate benchmark when the cap is not operative) or the price may be in terms of a floor as discussed below.

Similarly, the floating rate lender might want an arrangement whereby she receives floating rate subject to a floor of say 3%. This means that even if interest rates fall to 1% or 2%, she receives at least 3%, but if the interest rate rises, she receives the higher rate. Again the lender who wants a floor will have to pay for it in some form. The price could be an upfront payment, it could be a lower interest rate (she may agree to half percent less than the floating rate benchmark when the floor is not operative) or the price may be in terms of a cap as discussed above.

Many floating rate bonds have caps or floors as part of the instruments themselves. But it is also' possible to buy the cap or floor separately from the market. A borrower can issue a floating rate bond without any caps or floors and then go to a bank and buy a cap. If he desires to cap his interest cost at 10%, he would enter into a contract with the bank under which whenever the floating rate benchmark exceeds 10%, the bank pays him the excess times the notional amount of the loan. If the floating rate is

19.2 | Derivatives and Risk Management

less than 10%, the bank pays nothing. The cap is in this case, said to have a strike of 10%. The bank would charge a price for this cap. The borrower still pays the full floating rate to the lender, but whenever interest rates exceed 10%, he receives the excess from the bank and his net interest bill is capped at 10%. For example, if interest rates rise to 12%, he will pay 12% to the lender, receive 2% from the bank and his net cost will be only 10%. If interest rates are only 7%, the bank pays him nothing and his interest bill is simply the 7% that he pays to the lender.

Similarly, a lender can go to a bank and buy a floor with a strike of 3%. In this case, the contract with the bank would provide that whenever the floating rate benchmark is less than 3%, the bank pays her the shortfall times the notional amount of the loan. If the floating rate is more than 3%, the bank pays nothing. The bank would charge a price for this floor.

It may happen that the borrower, while buying a cap from a bank, might sell a floor to the same or another bank to recoup part or whole of the cost of the cap. Similarly, the lender who buys a floor might sell a cap to recoup part or whole of the cost of the floor.

A capped floating rate bond can be regarded as an uncapped bond plus a short cap (a cap sold by the bond holder to the bond issuer). It is therefore possible for an investor buying a capped floating rate bond to go to a bank and buy a cap to turn the bond into an uncapped bond. For example, if she has bought a bond capped at 10% and also bought a cap with a strike of 10% from the bank—if interest rates rise to 12%, she receives only 10% from the bond buy and 2% from the bank under the cap. Her total income is 12% – exactly what it would have been under an uncapped bond. If interest rates are only 7%, the bond pays 7% and the cap pays nothing so that he again receives the coupon of an uncapped bond.

More complex trades are possible. A borrower might have issued floating rate bonds with a cap of 10% but might decide later that he is prepared to pay up to 12%. In this case, the borrower can sell a cap with a strike of 10% to uncap the original bond and then buy a cap with a strike of 12% to recap it at 12%. If interest rates rise to 11%, he will pay 10% to the lenders under the original capped bond and he will pay an additional 1% under the cap that he has sold. He will receive nothing under the cap that he has bought. His total cost is 11%. If interest rates rise to 13%, he pays 10% to his lenders, pays 3% (13% – 10%) under the cap that he sold and receives 1% (13% – 12%) under the cap that he bought. His total cost is 10% + 3% – 1% = 12%. Thus, the situation is as if the bond had been capped at 12% instead of 10%.

In short, separately traded caps and floors allow caps and floors to be added, removed or modified in an existing floating rate bond. Whether caps and floors are part of the original bond or they are traded separately, they need to be valued. This is the topic that the next section deals with.

19.2 THE LIBOR MARKET MODEL (LMM) FOR CAP AND FLOOR VALUATION

To value a cap, it is useful to decompose it into a series of caplets, one for each coupon period. For example, consider a Rs 100 million five-year floating rate bond that pays quarterly coupons equal to three-month MIBOR¹ capped at 10%. In this case, there are 20 quarterly coupon periods. In each of these quarters, MIBOR has to be compared with the strike rate of the cap and the payment, if any has to be made at the end of the quarter. Each of these twenty quarterly payments is called a caplet.

¹ See footnote 2 in Chapter 18 for a discussion on the MIBOR market.

Caps, Floors and Swaptions | 19.3

The payoff of each caplet is:

- zero if MIBOR is below the strike
- 1/4 times Rs 100 million times the excess of MIBOR over the strike if MIBOR is above the strike

Algebraically, this is $\frac{L}{4}$ max (0, r - k) where L is the notional amount (Rs 100 million), r is the interest rate (MIBOR) and k is the strike rate of the cap.

The payoff of the caplet is exactly the same as that of $\frac{L}{4}$ call options on MIBOR whose strike rate is k. This suggests that the caplet can be valued using the Black-Scholes formula for valuing options by treating MIBOR as the underlying asset for the option. It is actually more convenient to use the Black

formula that is used for valuing options on futures or forward contracts. To use this model, two parameters are required:

- 1. The forward interest rate (which is the risk-neutral expectation of the MIBOR that would prevail on the caplet reset date)
- 2. The volatility of this MIBOR

Assuming that MIBOR has a log normal distribution, the Black formula can then be used to value the

caplet. The only catch is that the discounting has to be done till $\left(t + \frac{1}{4}\right)$ since the caplet payoff happens on that date, but the maturity of the option itself is only till *t* since the MIBOR is observed on that date. The formula is thus given by:

$$C_{t,t+1/4} = \frac{L}{4} P\left(0, t + \frac{1}{4}\right) (f_{t,t+1/4} N (d_1) - kN (d_2)$$

$$d_1 = \frac{\ln \frac{f_{t,t+1/4}}{k} + \frac{1}{2} \sigma^2 t}{\sigma \sqrt{t}}$$

$$d_2 = d_1 \sigma \sqrt{t}$$
(19.1)

where $C_{t,t+1/4}$ is the price of the caplet, *t* is the beginning of the coupon period (July 1, 2010), $P\left(0,t+\frac{1}{4}\right)$ is the price today (time 0) of the $\left(t+\frac{1}{4}\right)$ year zero coupon bond maturing at the end of the coupon period (September 1, 2010), $f_{t,t+1/4}$ is the forward Mibor rate and *k* is the strike rate of the cap. Comparing Eq (19.1) with the Black formula of Eq (10.20), it will be observed that the discounting factor has been changed from e^{-rt} to the price of a zero coupon bond. Moreover this bond matures at

 $\left(t+\frac{1}{4}\right)$ which is the date of the cash flow instead of at *t* which is the date of option maturity.

This formula gives the price of one caplet. The cap is valued by valuing each caplet and adding up all the values.

All this assumed that the floating rate payments are made four times in the year. If the payments are actually *m* times annually, the formula becomes:

The McGraw·Hill Companies

19.4 | Derivatives and Risk Management

$$C_{t,t+1/4} = \frac{L}{m} P\left(0, t + \frac{1}{m}\right) (f_{t,t+1/m} N(d_1) - kN(d_2)$$

$$d_1 = \frac{\ln \frac{f_{t,t+1/4}}{k} + \frac{1}{2}\sigma^2 t}{\sigma\sqrt{t}}$$

$$d_2 = d_1 - \sigma\sqrt{t}$$
(19.2)

Just as a cap is a series of caplets, each of which is a call option on interest rates, a floor is a series of floorlets, each of which is a put option on interest rates. The payoff of each floorlet is:

- zero if MIBOR is above the strike
- ¹/₄ times the notional amount times the excess of the strike over MIBOR if MIBOR is below the strike.

Algebraically, this is $\frac{L}{4}$ max (0, k - r) where L is the notional amount, r is the interest rate (MIBOR)

and k is the strike rate of the floor. This is the payoff of $\frac{L}{4}$ put options on Mibor. Using an analysis similar to that for caplets, the floorlet can be valued as:

$$F_{t,t,+1/4} = \frac{L}{m} P\left(0, t + \frac{1}{m}\right) (kN \ (d_2 f_{t,t+1/m} \ N \ (-d_1))$$

$$d_1 = \frac{\ln \frac{f_{t,t+1/4}}{k} + \frac{1}{2} \sigma^2 t}{\sigma \sqrt{t}}$$

$$d_2 = d_1 \sigma \sqrt{t}$$
(19.3)

where $F_{t,t+1/m}$ is the price of the floorlet, t is the beginning of the coupon period, $P\left(0, t + \frac{1}{m}\right)$ is the price of the $\left(t + \frac{1}{m}\right)$ year zero coupon bond maturing at the end of the coupon period, $f_{t,t+1/m}$ is the

forward Mibor rate and k is the strike rate of the floor.

This formula gives the price of one floorlet. The floor is valued by valuing each floorlet and adding up all the values.

This procedure of valuing caps and floors is known as the Libor Market Model (LMM) because it was developed in the context of Libor as the floating rate benchmarks and it was already being used by market practitioners before its theoretical justification was fully developed. The core idea of the LMM is (a) to assume that the floating rate (Libor or Mibor) has a log normal distribution and (b) to use risk-

neutral valuation with the $\left(t + \frac{1}{4}\right)$ year zero yield as the risk-free discount rate.

19.3 CAPLET AND CAP VOLATILITIES

The only unknown parameter in using the LMM to value caps and floors is the volatility to be used. This can be estimated from historical data on the evolution of the forward rate. Alternatively, given market prices of caps or caplets, the implied volatilities can also be worked out as in other option pricing contexts.

Since Eq. (19.1) is valuing each caplet separately, it is necessary to have a volatility estimate for each caplet. To value a cap, therefore, requires knowledge of all the caplet volatilities. For example, in a tenyear swap with quarterly floating rate payments, there are 40 caplets and each caplet could potentially have a different volatility. To value the cap, all the 40 caplet volatilities are needed.

In the cap market, dealers do often quote these caplet implied volatilities, also known as spot volatilities. It is also common, however, to quote flat volatilities. The flat volatility is a single volatility number that is used for each of the caplets so that the sum of these caplet values is the same as the cap value. What is being done is to compute some kind of a weighted average of all the individual caplet (spot) volatilities. When this average is used in place of the true caplet volatilities, some caplets will be under-priced and some will be over-priced. But these errors will cancel out and the value of the cap as a whole is correct. In some ways, this is similar to valuing a bond using YTM instead of using the appropriate zero yield for each cash flow. The YTM is also a weighted average of these zero yields and here too, while some cash flows are under-valued and some are over-valued, these errors cancel out and the value of the bond as a whole is correct. The advantage of the flat volatility is also the same as the advantage of the YTM—it is much more convenient to use a single volatility than to use dozens of different volatilities for a long maturity cap.

To illustrate the computations, consider the swap yield curve in Table 18.1 of the previous chapter. If the cap volatility (σ) is 20% and settlement is semi-annual (m = 2), the price of a cap for a notional (L) of 1,000,000 with a strike (k) of 5.25% is 12,754 as shown in Table 19.1. It will be observed that much of the value of the cap is accounted for by the later caplets. This arises out of two reasons. First, while the cap is out of the money, it is less so for distant coupons than for nearby ones. Because of the rising yield curve, the forward rate for distant coupon periods is closer to the strike. Second, the option value increases with maturity and the distant caplets have a longer option maturity.

Since floors and caps are puts and calls on the same underlying (the forward rate), there is a put-call parity between floors and caps. To buy a cap and sell a floor at the same strike is the same as entering into a pay fixed swap with a fixed rate equal to the strike. If the floating rate is less than the strike, the sold floor produces an outflow equal to the difference. If the floating rate is more than the strike, the bought cap produces an inflow equal to the difference. The result is that the cash flow is always equal to the difference between the strike rate and the floating rate, exactly as in the case of the swap.

This put-call parity implies that the implied volatility for a caplet is the same as the implied volatility for a floor at the same strike. Similarly, the cap volatility is the same as the floor volatility for the same strike. In the market these common volatilities are referred to as caplet and cap volatilities, but they can be used to price floorlets and floors as well.

Just as in other option markets, there can be a smile in cap volatilities. In other words, the implied volatility for an at-the-money cap can be different from the implied volatility for an out-of-the-money cap or an in-the-money cap. In liquid cap markets, the volatility smile can be obtained from a number of liquid strikes and this smile can be interpolated to price caps at other strikes.

	۰.
	1
	٩.
	æ
	7
-	
	71
1000	
100	-11
1000	
	40
	1
-	
10 m	
_	٠.
100 C C C C C C C C C C C C C C C C C C	1
-	
ALC: NO.	1
	1
	4
	ч
-	
_	۰.
A REAL PROPERTY.	-
Contraction of the local division of the loc	
-	
-	
-	1
	3
	1
	1
	1
-	1
-	
i	
·	
in	
-m	
-mu	
-mp	
-mp	
-mp.	
-wor	
-mou	
-mou	
Graw-	
Graw-	
Graw-	
-Graw-	
-Graw-	
cGraw-	
cGraw-	
CGraw-	
AcGraw-	
NcGraw-	
McGraw-	
e McGraw-	
e McGraw-	
e McGraw-	
he McGraw-	
he McGraw-	
The McGraw-	
The McGraw-	

t	t+1/m	Swap Yield	P(0,t+1/m)	f(t,t+1/m)	L/m	d ₁	\mathbf{d}_2	$N(d_1)$	N(d ₂)	Caplet
										Value
0.0	0.5	3.3100%	0.98372	3.3100%	500,000	-infinity	-infinity	0.0000	0.0000	0
0.5	1.0	3.7100%	0.96387	41182%	500,000	-1.6462	-1.7876	0.0499	0.0369	56
1.0	1.5	3.8650%	0.94412	4.1847%	500,000	-1.0339	-1.2339	0.1506	0.1086	283
1.5	2.0	4.0200%	0.92332	4.5054%	500,000	-0.5019	-0.7469	0.3079	0.2276	888
2.0	2.5	4.1050%	0.90316	4.4640%	500,000	-04319	-0.7148	0.3329	0.2374	1,083
2.5	3.0	4.1900%	0.88266	4.6444%	500,000	-0.2295	-0.5457	0.4092	0.2926	1,608
3.0	3.5	4.2450%	0.86281	4.6020%	500,000	-0.2071	-0.5535	0.4180	0.2900	1,731
3.5	4.0	4.3000%	0.84291	4.7218%	500,000	-0.0963	-0.4705	0.4616	0.3190	2,128
4.0	4.5	4.3450%	0.82338	4.7443%	500,000	-0.0532	-0.4532	0.4788	0.3252	2,323
4.5	5.0	4.3900%	0.80390	4.8451%	500,000	0.0230	-0.4013	0.5092	0.3441	2,655
Total (cap value									

19.4 SWAPTIONS AND CALLABLE BONDS

A swaption is an option to enter into a swap. For example, a swaption might give the right at the end of two years to enter into a three-year pay fixed swap with a fixed rate of 5.5%. The underlying in this case is a three-year swap, the option itself is European with a maturity of two years and the strike rate is 5.5%. This can be regarded as a call option on the swap². The corresponding option to enter into a receive-fixed swap can be regarded as a put option on the swap.

A swaption as discussed above, might appear quite unfamiliar and complex. However, all callable bonds contain embedded swaptions because the callability feature can be regarded as a swaption. Looked at in this way, swaptions are quite familiar and straightforward instruments. It is therefore, worthwhile establishing the equivalence between swaptions and callability.

Consider the following instruments:

- 1. Instrument A is a seven-year bond that is callable exactly at the end of two years. The coupon rate is 6.25% and the coupons are paid semi-annually. This is a 'European' call feature since the bond can be called at exactly one date.
- 2. Instrument B is a European swaption which gives the right at the end of two years to enter into a five-year receive-fixed swap with a fixed rate of 6.25%. The swap coupons are paid semi-annually.
- 3. Instrument C is a European swaption which gives the right at the end of two years to enter into a five-year pay fixed swap with a fixed rate of 6.25%. The swap coupons are paid semi-annually.
- 4. Instrument D is a non-callable seven-year bond (same maturity as bond A). The coupon rate is 6.25% and the coupons are paid semi-annually.
- 5. Instrument E is a non-callable two-year bond (maturing on the call date of bond A). The coupon rate is 6.25% and the coupons are paid semi-annually.

The following equivalences can then be established showing how the swaptions essentially replicate the callability feature of the bond.

$$A+B = D$$

$$A+C = E$$
(19.4)
values that the collection hand can be made into a non-collection of the

Consider the first equivalence that the callable bond can be made into a non-callable bond of the same maturity by adding a receive-fixed swaption. The equivalence is established as follows:

- At the call date two years from now, suppose bond A is not called. The bond cannot be called at any later date since the call feature is European. It then becomes identical to the non-callable bond D. The holder of A + B is free not to exercise the swaption B. The holder of A+B then ends up essentially with D establishing the equivalence in this case.
- At the call date, suppose bond A is called. The holder of A + B uses the redemption proceeds of bond A to buy a floating rate bond and simultaneously exercises the swaption B. The position is then described by Figure 19.1.

The holder of A+B now receives floating coupons from the floating rate bond which cancel out the floating rate payments to be made to the swap counterparty. The only cash flow that is left is the 6.25% fixed rate payment received from the swap counterparty. This is the same as the coupon on the original bond. At maturity, the redemption of the floating rate bond at par replicates the redemption of the original bond.

² It is conventional to regard the fixed rate payer as being long the swap and the fixed rate receiver as being short the swap. Under this convention, the right to enter into a pay fixed swap can be regarded as a call option on the swap. The right to enter into a receive fixed swap can be regarded as a put option on the swap.

The McGraw·Hill Companies

19.8 | Derivatives and Risk Management



Figure 19.1 Using a recieve fixed swaption to turn a callable bond into a non-callable bond of the same maturity. When the bond is called, the preceds are invested in a floating rate bond and the swaption is exercised. The coupons from the floating rate bond are used to make the floating rate payments to the swap counterparty. The fixed rate payment from the counterparty replicates the fixed coupon of the original bond

Thus, regardless of whether the bond is called or not, the holder of A + B is guaranteed to receive the same cash flows as he would have got from the non-callable bond D of the same maturity. This is why we can assert that A + B = D. The careful reader might object that we have actually shown only that A + B is worth at least as much as D because the analysis assumed that the swaption is exercised if and only if the bond is called. To demonstrate that A + B is worth exactly the same as D, it is necessary to show that this is optimal.

This is easily done because the issuer of bond A would call the bond only when it is cheaper to refinance the bond. The bond would be called if and only if the prevailing 5-year interest rate on the call date is less than the 6.25% that it is paying on the existing bond. The swaption should also be exercised if and only if the prevailing 5-year interest rate on the exercise date is less than the strike rate of 6.25%. Thus the circumstances under which it is optimal for the issuer to call the bond are exactly the same as those under which it is optimal to exercise the swaption. This completes the analysis and shows that A + B = D.

To demonstrate the other equivalence that A + C = E, the two scenarios are again considered, but the swaption C is the mirror image of swaption B. Each swaption is exercised exactly when the other is *not* exercised. The analysis is therefore as follows:

- At the call date two years from now, suppose bond A is called. The bond then becomes identical to the short maturity bond E. The holder of A + C is free not to exercise the swaption C and thus he ends up essentially with E establishing the equivalence in this case.
- At the call date, suppose bond A is not called. The holder of A+C exercises the swaption C. The position is then described by Figure 19.2.

The holder of A + C now receives the fixed coupons from the original bond A which cancel out the fixed rate payments to be made to the swap counterparty. The only cash flow that is left is the floating rate payment received from the swap counterparty. At maturity, the holder also gets redemption of the original bond at par. The cash flows thus replicates the cash flows of a floating rate bond. The floating rate bond is worth par at inception. The value of A + C at the call date is therefore par which is the same as the redemption proceeds of the short maturity bond E.

Caps, Floors and Swaptions | 19.9



Figure 19.2 Using a pay fixed swaption to turn a callable bond into a short maturity bond maturing on the call date. When the bond is not called, the swaption is exercised. The coupons from the original bond are used to make the fixed rate payments to the swap counterparty. The remaining cash flow is the floating rate payment from the counterparty. The result is the same as holding a floating rate bond. Since this floating rate bond is worth par, the holder obtains the same value as theredemption proceeds of the short maturity bond

Some readers might object that the floating rate bond worth par is not the same as having the same amount in cash. To see that there is no difference, consider what the holder would have done with the cash from bond E. Suppose she would have bought a four year bond paying the prevailing coupon on the date of maturity of bond E. The holder of A + C can enter into a four year receive fixed swap with a second swap counterparty at the swap rate prevailing at the same date. The floating rate payment received from the first swap counterparty would cancel out the floating rate payment to be made to the second swap counterparty. The only cash flow that would be left would be the fixed rate payment received from the second swap counterparty and the holder is again in the same position as if he had bought the four year fixed rate bond. At the end of four years the same process can be repeated because the floating rate bond depicted in Figure 19.2 is worth par on every reset date.

Thus, regardless of whether the bond is called or not, the holder of A + C is guaranteed to receive the same cash flows as she would have got from the short maturity bond E maturing on the call date. This is why we can assert that A + C = E. Mirroring the case of the A + B = D equivalence, it is easy to demonstrate that the swaption C is exercised if and only if the bond is not called. The issuer of bond A would call bond only when it is cheaper to refinance the bond. The bond would be called if and only if the prevailing five-year interest rate on the call date is less than the 6.25% that it is paying on the existing bond. The swaption should also be exercised if and only if the prevailing five-year interest rate of 6.25%. Thus, the circumstances under which it is *not* optimal for the issuer to call the bond are exactly the same as those under which it is optimal to exercise the swaption C. This completes the analysis and shows that A + C = E.

In reality, many bond callability features are not European. Some bonds are callable anytime (American call) while many are callable at every coupon date. An option that can be exercised at some specific dates (every coupon date for example) are called Bermudan to distinguish them from both European and American options. This choice of nomenclature is appropriate because Bermuda is an island in the Atlantic Ocean between America and Europe. Similarly, there are Bermudan swaptions that can be exercised at every reset date. Just as it was shown above that a European swaption replicates a European call provision in a bond, it can be easily seen that a Bermudan swaption replicates a Bermudan call provision in a bond.

The McGraw·Hill Companies

19.10 | Derivatives and Risk Management

19.5 THE SWAP MARKET MODEL FOR VALUING SWAPTIONS

The fact that the call provision in a bond is essentially a swaption, means that swaption valuation is an essential part of the valuation of a large class of bonds. The swaption is, however, very different in nature from the cap. The cap was decomposed into a number of caplets eachs of which turned out to be a call option. Thus the cap was a basket of options each of which could be individually exercised or allowed to lapse. On every coupon date, the cap holder can compare the cap rate with the prevailing floating rate and decide whether or not to exercise the caplet. The European swaption confronts the holder with only a single point decision. At the exercise date, the holder either enters into a swap or he does not. The decision to enter into the swap commits the swaption holder into a whole series of cash flows over the life of the swap. Unlike the case of the cap, it is not possible to cherry pick some of these dates, take the cash flows on those dates and discard the cash flows on other dates. The swaption is thus, not a basket of options but an option on a basket of cash flows.

However, the Black-Scholes approach can still be used by considering the payoff of the swaption on the exercise date. Consider for example, the two-year swaption to enter into a pay-fixed five-year swap with a strike rate of 6.25%. On the exercise date, if the five-year swap rate is more than 6.25%, it is worthwhile to enter into this swap. For example, if the five-year swap rate is 7%, then the swaption allows the holder to pay only 6.25% fixed. In the absence of the swaption, he would have had to pay 7% to enter into the same swap. There is thus a gain of 0.75% every year (or with semi-annual payments, 0.375% every six months) for the next five years.

To proceed further, let us assume that the payments occur *m* times a year for *n* years and focus on the *i*'th cash flow occurring at time $T + \frac{i}{m}$ where *T* is the maturity of the option and *n* is the maturity of the underlying swap. If *L* is the notional amount of the swap, the cash flow on this date is

$$\frac{L}{n}(s_T - X, 0)$$
 (19.5)

where s_T is the swap rate prevailing at time T and X is the strike rate of the swaption. This is the payoff of $\frac{L}{-}$ call options on the swap rate.

Assuming that the swap rate is log normal, the Black formula can be used to value this cash flow. The only catch is that the discounting has to be done till $T + \frac{i}{m}$ since the cash flow happens on that date, but the maturity of the option itself is only till T since the swap rate is observed on that date. The Black formula thus becomes:

$$\frac{L}{m} P\left(0, T + \frac{i}{m}\right) (s_0 N (d_1) - kN (d_2)$$
(19.6)

where

$$d_1 = \frac{\ln \frac{S_0}{k} + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}$$
$$d_2 = d_1 \sigma\sqrt{T}$$

Caps, Floors and Swaptions | 19.11

where T is the exercise date of the swaption, $P\left(0, T + \frac{i}{m}\right)$ is the price today (time 0) of the $T + \frac{i}{m}$ year zero coupon bond maturing on the *i*'th payment date, s_0 is the forward swap rate and *k* is the strike rate of the swaption. The forward swap rate can be computed, using the yield curve today as explained later in this section.

The value of the swaption can be obtained by using Eq. (19.6) for each of the payment dates and

adding up the values. It is evident that for each date *i*, there is a different discount factor $P\left(0, T + \frac{i}{m}\right)$,

but the undiscounted value $\frac{L}{m}$ ($s_0 N(d_1) - kN(d_2)$) is the same for all the dates. Therefore, the swaption value requires adding up all the discount factors and multiplying the sum by the undiscounted value $\frac{L}{m}$ ($s_0 N$)

$$\frac{-}{m}$$
 (s₀ N).

This leads to the swaption valuation formula under what is known as the Swap Market Model (SMM):

$$LA (s_0 N (d_1) - k N (d_2))$$
(19.7)

where

$$A = \frac{1}{m} \sum_{i=1}^{mn} P\left(0, T + \frac{i}{m}\right)$$
$$d_1 = \frac{\ln \frac{s_0}{k} + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}$$
$$d_2 = d_1 \sigma\sqrt{T}$$

In this expression, A is the present value of an annuity that pays Rs 1 per year (in m equal instalments each year) for n years.

Similarly, the value of a receive-fixed swaption involves put options on the swap rate and can be expressed as:

$$LA \ (kN \ (-d_2) - s_0 \ N \ (-d_1)) \tag{19.8}$$

where

$$A_{1} = \frac{1}{m} \sum_{i=1}^{mn} P\left(0, T + \frac{i}{m}\right)$$
$$d_{1} = \frac{\ln \frac{s_{0}}{k} + \frac{1}{2}\sigma^{2}T}{\sigma\sqrt{T}}$$
$$d_{2} = d_{1} - \sigma\sqrt{T}$$

It remains to determine the forward swap rate. The swap rate is nothing but a par bond yield and it can be readily determined from forward zero yields. This in turn, can be easily computed from the current zero yields in a manner very similar to how the forward rates were computed in Chapter 18. There is, however, a simpler method of computing the forward swap rate directly.

If the forward swap rate is s_0 , the present value at time 0 of all the fixed rate payments is $s_0 A$ where A is the same as in Eq. (19.7) or Eq. (19.8). To value the floating rate payments, we use the fact that the

The McGraw·Hill Companies

19.12 | Derivatives and Risk Management

floating rate payments would be worth par at time *T* if there were an additional principal payment at time T + n. With that additional principal payment, the cash flows would be the same as for a floating rate bond which is worth par at inception. If value is par at time *T*, then the value at time 0 is P(0, T). The value at time 0 of the additional principal payment at time T+n is P(0, T+n). Therefore the value at time 0 of the floating rate payments (only coupons) without the additional principal payment at time T+n is P(0, T) - P(0, T+n). Equating the present values of the fixed and floating payments gives the relationship $s_0A = P(0, I) - P(0, T+n)$ which gives us the forward swap rate as:

$$s_0 = \frac{P(0,T) - P(0,T+n)}{A}$$
(19.9)

To illustrate all these computations, consider once again the swap yield curve of Table 18.1 in the previous chapter. Consider two-year swaptions (T = 2) where the underlying is a five-year swap (n = 5), the strike rate (k) of the swaption is 5%, the notional (L) is 1,000,000 and swap payments are semiannual (m = 2). If the swaption volatility (σ) is 10%, then the swaption will be valued as shown belows.

The first step is to compute A from the zero prices in Table 18.1 using Eq. (19.7). The value of A is 4.07262.

t+l/m	Zero price
2.5	0.90316
3.0	0.88266
3.5	0.86281
4.0	0.84291
4.5	0.82338
5.0	0.80390
5.5	0.78485
6.0	0.76593
6.5	0.74714
7.0	0.72850
Sum =	8.14525
$\frac{\text{Sum}}{m} =$	4.07262

The next step is to compute the forward swap rate s_0 using Eq. (19.9) as follows:

$$s_0 = \frac{P(0,T) - P(T+n)}{A} = \frac{P(0,T) - P(0,+7)}{A} = \frac{0.92332 - 0.72850}{4.07262} = 4.784\%$$

With these values in hand, the value of a pay-fixed swaption can be worked out using Eq. (19.7) as follows:

07262
07202
784%
0.24221
0.38363
40431
35063
367.5

Caps, Floors and Swaptions | 19.13

A	4.07262
<i>s</i> ₀	4.784%
d_1	-0.24221
d_2	-0.38363
$N(d_1)$	0.59569
$N(d_2)$	0.64937
Value of receive fix	
swaption using Eq (19.8)	16, 182.53

The value of a receive-fixed swaption can be worked out using Eq. (19.8) as follows:

19.6 SWAPTION VOLATILITIES

To complete the swaption valuation using Eq. (19.7) or Eq. (19.8), the volatility of the swap rate is needed. This can be estimated from historical data on the evolution of the swap rate. Alternatively, given market prices of swaptions, the implied volatilities can also be worked out as in other option pricing contexts.

Since receive-fixed and pay-fixed swaptions are puts and calls on the same underlying (the swap rate), there is a put-call parity between these two kinds of swaption. To buy a pay-fixed swaption and sell a receive-fixed swaption at the same strike is the same as entering into a forward swap (starting at the swaption exercise date) with a fixed rate equal to the strike. If at the exercise date, the swap rate is less than the strike, the sold swaption would be exercised. If the swap rate is more than the strike, the bought swaption would be exercised. The result is that in either case the swap would be entered into.

Put-call parity implies that the implied volatilities for receive-fixed and pay-fixed swaptions at the same strike are the same.

Just as in other option markets, there can be a smile in swaption volatilities. In other words, the implied volatility for an at-the-money swaption can be different from the implied volatility for an out-of-the-money swaption or an in-the-money swaption. In liquid markets, the volatility smile can be obtained from a number of liquid strikes and this smile can be interpolated to price swaptions at other strikes. However, at-the-money swaptions are more liquid and implied volatilities of these swaptions are more commonly quoted than for other strikes.

What is the difference between swaption and cap volatilities for comparable maturities? The cap gives a basket of options on forward rates while the swaption gives an option on a basket (the swap rate is essentially a basket or weighted average of forward rates). The volatility of an average is lower than the volatilities of the individual rates unless the individual rates are highly correlated. The gap between the cap and swaption volatilities is, therefore, a measure of the correlation between forward rates.

The difference between implied volatilities of caps and swaptions is a measure of the 'implied correlation' between forward rates. Just as the implied volatility is the market's expectation of what the volatilities will be in future, the implied correlation is the market's expectation of what the correlation will be in future. Just as implied volatilities give information that is not always present in historical volatilities, implied correlations give information that is not always present in historical correlations. Just as options are vehicles for trading volatility (taking views on future volatility and trading on these views), cap and swaption markets together provide an opportunity to trade correlations (taking views on future correlations and trading on these views).

19.14 | Derivatives and Risk Management

19.7^{*} RISK NEUTRAL VALUATION OF INTEREST RATE DERIVATIVES

There are some complications in using risk-neutral valuation for interest rate derivatives like caps and swaptions. This optional section deals with some of these issues.

The first difficulty lies in treating the interest rate as the underlying since it is the bond price that is the price of a traded asset. The interest rate itself is not the price of a traded asset and, therefore, it might appear that risk-neutral valuation cannot be used directly. But this difficulty is more apparent than real. In the context of the caplet valuation example mentioned earlier in this chapter, consider the following strategy to deal with the caplet for the period from 1 July, 2010 to 30 September, 2010:

- Today, buy zero coupon bonds with a face value of Rs 4 maturing on 1 July, 2010
- On 1 July, 2010, invest the proceeds of these bond (Rs 4) at the three month MIBOR prevailing on 1 July, 2010
- Today, borrow part of the cost of the bonds bought today by selling zero coupon bonds with face value of Rs 4, maturing on 1 September, 2010

The cash flows on 1 September, 2010 would be as follows:

- *Inflow:* Rs 4 plus interest on Rs 4 for one quarter at the three-month MIBOR that prevailed on 1 July, 2010
- Outflow: Rs 4

The net cash flow is simply interest at MIBOR for a quarter on Rs 4 which is MIBOR $\times \frac{1}{4} \times 4 \times =$

MIBOR. In other words, the cash flow from this strategy on 1 September, 2010 equals the MIBOR that prevailed on 1 July, 2010. Thus though MIBOR is not a traded asset, there is a trading strategy that produces a cash flow at the end of the coupon period equal to the MIBOR that prevailed at the beginning of the coupon period. There are no cash flows in the intervening period because the redemption proceeds on 1 July, 2010 are promptly invested at MIBOR for the quarter. The strategy thus has only two cash flows: a cash outflow equal to the price of the initial portfolio and a cash inflow on 1 September, 2010 equal to the MIBOR that prevailed on 1 July, 2010.

The principle of risk-neutral valuation tells us that the cost of the initial portfolio is equal to the discounted present value of the risk neutral expectation of MIBOR.

$$\left[\left(t + \frac{1}{4} \right) \text{ year zero price} \right] E^*(r_t) = 4 \left[t \text{ year zero price} - \left(t + \frac{1}{4} \right) \text{ year zero price} \right]$$
$$E^*(r_t) = 4 \left[\frac{t \text{ year zero price} - \left(t + \frac{1}{4} \right) \text{ year zero price}}{\left(t + \frac{1}{4} \right) \text{ year zero price}} \right]$$
(19.10)

where *t* is the beginning of the coupon period (1 July, 2010), v_t the three-month MIBOR prevailing at time *t* and $\left(t + \frac{1}{4}\right)$ is the end of the coupon period (1 September, 2010) and as usual *E** denotes the risk neutral expectation.

^{*}Optional section. May be omitted without loss of continuity.

Caps, Floors and Swaptions | 19.15

Comparing this with Eq. (18.6) in the previous chapter shows that the risk-neutral expectation is nothing but the forward interest rate. Thus, everything works out well if the interest rate is treated as if it were a traded asset, and as if the forward interest rate were its forward price.

A similar approach works for the swap rate as well. In fact, the formula for the forward swap rate in Eq. (19.9) is based precisely on inferring the risk-neutral expectation out of the risk-neutral value of traded assets (various zero coupon bonds).

Another complication in interest rate derivatives is the very meaning of the risk-free rate. First, when interest rates are changing, what is the meaning of risk-free rate? Second, since there are several different risk-free rates (one for each maturity) which of these is the risk-free rate? The answer to both of these is the same—there is a different risk-free rate for each maturity. The zero rate for a specific maturity is risk-free for valuing a cash flow at that maturity but not for any other maturity. Provided we use the right zero rate all the time, there is no problem.

However, this implies a further twist when it comes to risk neutral probabilities and expectations. The risk neutral probabilities do depend on which interest rate is taken as risk free. For example, in each caplet valuation there is a different risk free rate and associated with it is a different set of risk neutral probabilities. This is not a problem because each caplet is valued separately and the results are added together. It does not matter at all that the risk neutral probabilities that were used to value one caplet are different from those used to value another caplet.

The swaption valuation involves yet another risk free rate (it can in some sense be thought of as a weighted average of all the zero rates that enter into the annuity factor A in Eq (19.9)). There is another set of risk neutral probabilities associated with this risk free rate and these probabilities are different from those used in cap valuation. This is not a problem for cap valuation or swaption valuation considered separately. It does lead to serious difficulties when considering more complex derivatives whose values depend on both cap and swaption values. These topics are beyond the scope of this book.

19.8^{*} RECONCILING LOGNORMALITY ASSUMPTION FOR CAPS AND SWAPTIONS

The Libor Market Model (LMM) uses the Black formula under the assumption that the floating rate is lognormal while the Swap Market Model (SMM) uses the Black formula under the assumption that the swap rate is lognormal. Both these assumptions appear to be reasonable approximations to the true distribution of these rates and therefore LMM and SMM have proved to be hugely popular.

Taken literally, however, these assumption are not mathematically consistent. The difficulty is that if two quantities are lognormal, their sum or average is not lognormal³. This is not an insuperable difficulty because these assumptions are not intended to be taken literally. In practice, all option markets have a smile which captures the departures from the log normal distribution. When practitioners use a smile, they are acknowledging that the log normality assumption is not exact.

What is important is that the assumption be approximately true so that the Black-Scholes formulas provide useful results, particularly when combined with a suitable smile. It is possible for both the forward rate and the swap rate to be approximately log normal. There is no mathematical impossibility here.

^{*}Optional section. May be omitted without loss of continuity

³On the other hand, the product or quotient is log normal. If two quantities follow the normal distribution, their sums and differences are normal. The lognormal distribution implies that the logarithm is normal. The logarithm of a product is the sum of the logarithms and therefore the product of log normals is lognormal.

The McGraw·Hill Companies

19.16 | Derivatives and Risk Management

Chapter Summary

A cap places an upper bound on the coupon of a floating rate bond, while a floor places a lower bound on the coupon. Caps and floors are also traded by themselves. A cap can be decomposed into a series of caplets (one for each coupon period) and each caplet is a call option on the forward rate for that period. The Black formula can be used to value each caplet separately.

$$C_{t,t+1/m} = \frac{L}{m} P\left(0, t + \frac{1}{m}\right) (f_{t,t+1/m} N (d_1) - kN (d_2))$$
$$d_1 = \frac{\ln \frac{f_{t,t+1/m}}{k} + \frac{1}{2}\sigma^2 t}{\sigma\sqrt{t}}$$
$$d_2 = d_1 - \sigma\sqrt{t}$$

where $C_{t\,t+1/m}$ is the price of the caplet, t is the beginning of the coupon period, $P\left(0, t+\frac{1}{m}\right)$ is the price of the

 $\left(t+\frac{1}{m}\right)$ year zero coupon bond maturing at the end of the coupon period, $f_{t,t+1/m}$ is the forward Mibor rate and k is the set of the flace.

is the strike rate of the floor.

The sum of all the caplet values is the value of the cap. Each caplet should ideally be valued using its own implied volatility, but it is more common in the market to quote a flat volatility which is a weighted average of the caplet (spot) volatilities. Valuing all caplets using this common volatility gives the same value for the cap.

A floor is similarly decomposed into floorlets, each of which is a put option with a value

$$F_{t,t+1/m} = \frac{L}{m} P\left(0, t + \frac{1}{m}\right) (kN(-d_2) - f_{t,t+1/m}N(-d_1))$$

where all the other variables are the same as in the caplet.

The swaption is an option on the swap rate which is a weighted average of the forward rates applicable during the tenure of the swap. The pay-fixed swaption is a call option on the swap rate while the receive-lixed swaption is a put option on the swap rate. The underlying swap commences at the swaption exercise date and runs for its full tenure from that date. Swaption exercise can be European (single exercise date), American (exercise at any time prior to expiry) or Bermudan (multiple exercise dates—typically same as coupon dates).

Swaptions are intimately linked with the callability feature of bonds. For example the holder of a seven-year bond, callable exactly at the end of two years can convert this bond into a non-callable seven-year bond by buying a swaption exercisable at the end of two years into a five-year receive-fixed swap. Alternatively, the callable bond can be converted into a two-year bond by buying a swaption exercisable at the end of two-year pay fixed swap. Thus, swaption valuation enters into the valuation of callable bonds.

Swaptions are valued using the Black formula with the swap rate as the underlying. The payoff on each coupon date is the same and the risk-neutral expectation of this payoff is also the same. However, since these payoffs occur at different dates, they have to be multiplied by different discount factors. This produces a present value of annuity factor in the formula.

$$LA (s_0 N (d_1) - k N (d_2))$$

where

$$A_1 = \frac{1}{m} \sum_{i=1}^{mn} P\left(0, T\frac{i}{m}\right)$$

Caps, Floors and Swaptions | 19.17

$$d_1 = \frac{\ln \frac{s_0}{k} + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}$$
$$d_2 = d_1 \sigma \sqrt{T}$$

where T is the exercise date of the swaption, $P\left(0, T\frac{i}{m}\right)$ is the price today (time 0) of the $T + \frac{i}{m}$ year zero coupon

bond maturing on the z'th payment date, s_0 is the forward swap rate and k is the strike rate of the swaption. The forward swap rate can be computed as:

$$s_0 = \frac{P(0,T) - P(0,T+n)}{A}$$

Swaption volatilities differ from cap volatilities because the swap rate as a weighted average of forward rates has a lower volatility than the individual forward rates, unless the correlation between the forward rates is quite high. The gap between cap implied volatilities and swaption implied volatilities is thus a measure of market implied correlations between the forward rates.

Suggestions for Further Reading

The mathematical justification of the valuation models presented in the chapter is provided by:

Brace, A., Gatarek, D. and Musiela, M (1997), "The market model of interest rate dynamics", *Mathematical Finance*, 7(2), 127-147

A good introduction to the use of swaptions to hedge callable debt is given in:

Brown, Kerlk and Donald J Smith (1990) "Forward Swaps, Swaps Options and the Management of Callable debt", *Journal of Applied Corporate Finance*, 2:4, Winter, 59-71.

Chapter **Twenty**

Derivative Accounting

The purpose of this chapter is to provide an introduction to the accounting treatment of derivatives. Unlike most other assets and liabilities, derivatives are shown at fair value. It is in this context, that the accounting definition of derivatives and the treatment of embedded derivatives becomes important.

To accommodate fair value accounting of the derivative, changes are required in the accounting of the hedged item also when a derivative is used as a hedge. The chapter discusses the notions of fair value hedges and cash flow hedges, and the different accounting implications of these two kinds of hedging strategies.

Hedge accounting is subject to the fulfilment of elaborate documentation requirements, as well as a demonstration of hedge effectiveness. These requirements are also outlined in the chapter.

Hedge accounting is quite complex — the accounting standards that describe this run into about 200 pages, and are supplemented by equally bulky material on implementation guidance. The chapter, therefore, provides only an overview of the complex subject of derivative accounting.

20.1 INTRODUCTION TO DERIVATIVES ACCOUNTING

The accounting treatment for derivatives is in a state of flux. In India, the accounting standard relating to this (AS 30) is in the process of being issued — an exposure draft has been released for comments. Globally, there are two principal accounting standards in this area — FAS 133 in the United States and the International Accounting Standard IAS 39 elsewhere in the world. The draft AS 30 in India is almost identical to IAS 39.

The accounting treatment discussed here, largely follows IAS 39 which is practically the same as the proposed AS 30. However, the differences between FAS 133 and IAS 39 are relatively small and these small differences are not relevant to most of the discussion in this chapter. Each of these accounting standards is more than 200 pages long, and the discussion here, therefore, ignores many of the finer details and covers only the salient features of the accounting for derivatives.

Another complication is that the US and international accounting standards boards have begun a long-term project to completely rewrite the accounting standards on radically different lines and achieve significant simplification. This is touched upon later in this section.

Derivative accounting is extremely complex because it involves two fields —accounting and finance — that have very different conceptual foundations:

• Accounting is focused on transactions as accountants seek to record the various transactions as these occur. For example, if a company places an order for some machinery or submits a bid in response to a tender, accountants do not typically record anything. They wait for the actual transaction (for example, shipments of the goods) to happen before recording anything. If this logic were extended to a forward contract, there would be nothing to record until the expiry of the contract and by then there will be no derivative to record. So accountants have to create a new rule and say that a derivative must be recorded even when it is only a firm commitment and not yet a completed transaction. This creates a further problem if a company uses a forward

20.2 | Derivatives and Risk Management

contract to hedge a business commitment. The problem is that the hedge is recorded using the new rules, but the commitment that it hedges is not even recorded in the books because the business commitment is covered by the old rules. Another set of new rules are needed to reconcile this difficulty. Finance on the other hands, avoids these problems because its focus is entirely on future cash flows and not on completed transactions.

- The accountants' focus on transactions creates another problem in defining hedging. In finance, risk is measured at the level of a portfolio of assets and liabilities. Because of the positive or negative correlation between different components of the portfolio, it is not possible to measure portfolio risk by measuring the risk of individual components. Accounting, however, tries to define hedging at the level of not just individual assets and liabilities, but at the level of individual transactions. Therefore, a company that has an asset and a liability that perfectly hedge each other can go ahead and hedge the asset using a derivative. At the portfolio level, this 'hedge' serves only to create a risk which was absent earlier, but the accountants think that this is hedging. The implementation guidance for IAS 39 (and AS 30) while laying down this position states that the accounting. Exposure is assessed on a transaction basis.' The other side of a transaction view is that a derivative that is a hedge from a finance point of view, because it reduces portfolio risk (the so called 'macro' hedge) need not be a hedge in the accounting view because it is not possible to identify a specific transaction that it hedges.
- Accountants have traditionally recorded transactions at the transaction value (historical cost) and left these values undisturbed until the occurrence of another transaction that alters the value. The advantage of the transaction value is that it is objectively verifiable. There is little scope for disagreement on what a company paid to buy its office building. On the other hand, the current market price of that building is not so easy to determine. Accountants, therefore, like to record the building in their books at the original transaction value.¹ The defining characteristic of a derivative on the other hand is that its value keeps changing as the underlying changes. It is utterly pointless to show a derivative in the books at historical cost. Also, the existence of liquid markets mean that current market value is objectively verifiable for many derivative contracts. For example, there is little scope for disagreement on the value of a foreign exchange forward contract. Accountants therefore responded to the emergence of derivatives with a new set of rules that said that derivatives should be valued at current market value or fair value and not at historical cost. Again this creates problems when a derivative, like an interest rate swap, is used to hedge the company's fixed rate borrowings. The interest rate swap is recorded at fair value and changes in its value, therefore, appear as profits and losses in the accounting books. If the fixed rate borrowing is recorded at historical cost and changes in its value do not show up at all, the accounting profits will not show any hedging benefits. Again, a complex set of rules is required to get around this and reflect the hedging correctly.

There is a proposal currently to solve some of these problems by requiring all financial assets and liabilities to be recorded at fair value (current market value). The US and international accounting standards boards agreed in 2005 that²: 'The long-term objectives, assuming that technical and practical hurdles can be overcome, are ... [t]o require that all financial instruments be measured at fair value with

¹ They may depreciate the building on a straight line basis or using some other formula, but the depreciated value is still based on the original cost and not on its current value.

² <u>http://www.fasb.org/proiect/financial instruments.shtml</u>

Derivative Accounting | 20.3

realized and unrealized gains and losses, recognized in the period in which they occur [and] simplify or eliminate the need for special hedge accounting'. The first drafts of this framework are expected in 2008 and the final standard is probably years away.

This proposal would eliminate a great deal of the complexities involved in defining what hedging is and how to account for it, because most of the time derivatives are used to hedge financial assets and liabilities. This is clearly true for banks, securities firms and other financial institutions, as most of their assets and liabilities are financial. But even in non-financial corporations, almost all interest rate hedging is the hedging of financial assets and liabilities. A large part of corporate exchange risk hedging is also hedging of financial assets like receivables, payables, and foreign currency debt, but a significant part is hedging of future transactions. Most commodity price hedging would, however, not be hedges of financial assets and liabilities. So the shift to fair value accounting of all financial assets and liabilities would simplify hedge accounting, but would not eliminate it completely.

20.2 DERIVATIVES AND FAIR VALUE ACCOUNTING

20.2.1 Definition of Derivatives

The accounting definition of derivatives is very close to the finance definition, but it is not identical. A derivative is defined for accounting purposes as a contract with four characteristics:

- 1. *Underlying:* Its value changes in response to the change in an underlying (a specified interest rate, financial instrument price, commodity price, foreign exchange rate, index of prices or rates, credit rating or credit index, or other variables)
- 2. *Leverage:* The initial net investment required is either nil or lower than what would be required for other types of contracts, expected to have a similar response to the underlying. For example, futures contracts and forward contracts at the prevailing forward price involve no net investment. (The margin deposited with the exchange is regarded as a separate asset and not as the initial investment in the derivative). An option contract does require investment of option premium, but this is substantially lesser than the notional value. A currency swap might involve an initial exchange of principals in different currencies, but since these amounts have the same value, there is no net investment. However, deep in the money options or forward contracts at rates much better than the forward price would involve a large initial net investment and may fail to be derivatives under the accounting definition. Similarly, a pre-paid forward contract may not be a derivative for accounting purposes.
- 3. *Future Settlement:* It is settled at a future date.
- **4.** *Cash Settlement:* The international (and proposed Indian) standard requires that the derivative be cash-settled or be settled by the delivery of a financial instrument. The US standard is similar in that, it requires that the derivative can be either readily net-settled or involves an underlying, whose delivery is not substantially different from net settlement. Under either definition, financial derivatives are derivatives but physically settled commodity derivative contracts that are not traded in an exchange or in an active OTC market may not be covered.

20.2.2 Embedded Derivatives

An embedded derivative is a derivative that is embedded in a host non-derivative contract, so that some of the cash flows of the combined instrument vary in a manner similar to that of a free-standing derivative. The key requirement is that a separate instrument with the same terms as the embedded derivative would meet the definition of a derivative.

20.4 | Derivatives and Risk Management

If the combined instrument is not subject to fair value accounting, there would be an anomaly, in that, a free-standing derivative would be subject to fair value accounting, but the embedded derivative is not so subject. To eliminate this anomaly, accounting standards require that unless the combined instrument is itself subject to fair value accounting (both for balance sheet, and profit and loss purposes), the embedded derivative should be separated from the host instrument and accounted for as a derivative.

A strict imposition of this requirement would imply that a host of well established simple financial instruments would contain embedded derivatives and would have to be accounted for separately. To reduce the need to do this, accountants have carved out a complex set of exceptions to the principle that embedded derivatives should be separated. The exception applies if 'the economic characteristics and risks of the embedded derivative are closely related to the economic characteristics and risks of the host contract.' This test is not very easy to apply and the accounting standards are complicated by a whole set of examples and counter-examples that attempt to clarify this difficult concept.

A simple example of this exception is a physically-settled contract, entered into by an Indian entity to buy crude oil at a price fixed in US dollars. Since global trade in crude oil takes place in dollars, the typical forward contract for crude oil is denominated in US dollars. Technically, this contract contains an embedded currency derivative. However, this is not required to be separated out, as it is regarded as being closely related to the contract to buy crude oil.

Another example of the exception is floating rate debt instruments with a cap or floor. Exposure to interest rates is the principal economic characteristic of the floating rate bond and the cap is an embedded derivative, whose principal economic characteristic is also an exposure to the same interest rates. Therefore, the cap need not be separated from the floating rate bond and need not be accounted for as a derivative.

But there are exceptions to the exception. This is because it is possible to design a floating rate instrument, where the embedded cap or floor dominates it to such a degree that the tail starts wagging the dog. For example, when the interest rate is 6%, the cap may be set at 2% but the interest may be paid on thrice the face value of the bond. So the accounting standards say that the exception for caps does not apply if the cap is set lower than the market rate of interest, or the cap is levered.

This system of exceptions, and exceptions to exceptions is inherently messy but there is no simplification possible, so long as there is a desire not to do fair value accounting for simple instruments.

20.2.3 Fair Value Accounting

Accounting standards require that all derivative contracts be accounted for at fair value in the balance sheet, and for profit and loss purposes. This means that they must be marked to market: any change in the fair (market) value of a derivative becomes a part of the profit or loss for that period. If a derivative was bought for Rs 100 million and the value falls to Rs 92 million at the end of the quarter, the loss of Rs 8 million will be reported as part of the quarterly profit or loss. If the fair value of the derivative rises to Rs 106 million at the end of the quarter, the gain of Rs 6 million will also be reported as part of the quarterly profit or loss.

To see the implication of this, consider the case of an Indian company that imports \$25 million every quarter and therefore buys \$100 million forward (\$25 million each three month, six months, nine months, and twelve months forward) at Rs 45/\$ to hedge its anticipated imports of raw material during the next 12 months. Suppose the company does not adopt the hedge accounting discussed in the next section. Assume for simplicity, that both rupee and dollar interest rates are zero.

Suppose at the end of the first quarter, the exchange rate has moved to Rs 44/\$. Of its \$100 million hedge, the first \$25 million that was bought three months forward would have expired and would have been used to pay for the imports of the quarter. However, \$75 million of the hedge would still be outstanding. Fair value accounting requires that this be revalued at quarter-end prices. This would produce a loss of Rs 75 million due to the depreciation of the US dollar against the rupee by Rs 1/\$. This loss will be reflected not only in the balance sheet but also in the profit and loss for that quarter.

The purpose of hedging the raw material imports was to bring greater certainty into the company's profits. The accounting impact is, however, quite the opposite because in the very first quarter, the profits are being impacted by the losses on forward contracts, covering an entire year's imports. In fact, after hedging, the quarterly profits become more uncertain than they would have been without hedging as shown in the hypothetical example in Table 20.1.

This example assumes that quarterly revenues are Rs 3,500 million, raw material costs are \$25 million, and other costs are Rs 2,000 million. It considers \$100 million hedged at Rs 45/\$ and considers three exchange rate scenarios (Rs 44/\$, Rs 45/\$, and Rs 46/\$). Without hedging, the range of profits from the lowest (Rs 350 million) to the highest (Rs 400 million) is Rs 50 million. After hedging four quarters of imports, the range of profits from lowest (Rs 300 million) to highest (Rs 450 million) is Rs 150 million. This is because the benefits of a cheaper dollar will impact raw material costs partly in future quarters, while it impacts the entire forward contract this quarter.

Table 20.1: In the absence of hedge accounting, hedging future imports can make profits more volatile. Without hedging, the range of profits from the lowest (350) to the highest (400) is Rs 50 million. After hedging four quarters of imports, the range of profits from lowest (300) to highest (450) is Rs 150 million. This is because the benefits of a cheaper dollar will impact raw material costs in future quarters, while it impacts the entire forward contract this quarter. This example assumes that revenues = Rs 3, 500 million, raw material = \$ 25 million, other costs = Rs 2,000 million. It considers \$100 m hedged at Rs 45/\$ and considers three exchange rate scenarios (Rs 44/\$, Rs 45/\$ and Rs 46/\$). It is assumed for simplicity that both rupee and dollar interest rates are zero

Exchange rate Rs/\$	44.00	45.00	46.00
Without Hedging			
Revenues	3,500	3,500	3,500
Imported raw materials \$ 25 m at prevailing exchange rate	1,100	1,125	1,150
Other costs	2,000	2,000	2,000
Profit before tax	400	375	350
Average Profit	375		
Range of Profits (highest - lowest)		50	
With \$100m Hedged at Rs	45/\$		
Revenues	3,500	3,500	3,500
Imported raw materials \$ 25 m at Rs 45/\$	1,125	1,125	1,125
Other costs	2,000	2,000	2,000
Hedging Gain/Loss on remaining \$ 75 m	-75	0	75
Profit before tax	300	375	450
Average Profit		375	
Range of Profits (highest - lowest)		150	

20.6 | Derivatives and Risk Management

Conceptually, why the problem crops up is very clear. The hedge (forward contract) is an asset which is shown at fair value in the books of account because it is a derivative. The hedged item (future raw material purchases) is not a liability in the accountants' definition and is not recorded in the their books at all. In finance theory, the hedge appears to be a perfect hedge because implicitly, finance theory regards future raw material purchases as a liability and moreover all assets and liabilities are implicitly considered at fair value.

Table 20.2 shows that in finance theory, the unhedged position has significant risks because the value of raw material purchases has a range from Rs 4,400 million (lowest) to Rs 4,600 million (highest) of Rs 200 million when the exchange rate varies from Rs 44/\$ to Rs 46/\$. The hedged position on the other hand is risk-free as it has a value of Rs 4,500 millions regardless of the exchange rate.

Table 20.2: In the finance theory view, which recognizes future raw material imports as a liability and values it at market prices, hedging the entire year's imports of \$25 million per quarter at Rs 45/\$ does eliminate risk completely. The unhedged position has significant risks because the value of raw material purchases has a range from Rs 4,400 million (lowest) to Rs 4,600 million (highest) of Rs 200 million when the quarterend exchange rate varies from Rs 44/\$ to Rs 46/\$. It is assumed for simplicity that both rupee and dollar interest rates are zero

Exchange rate Rs/\$	44.00	45.00	46.00
Without Hedging			
Raw materials imports this quarter \$ 25 m at prevailing exchange rate	1,100	1,125	1,150
Raw material imports next three quarters \$ 75 m at prevailing exchange rate	3,300	3,375	3,450
Total	4,400	4,500	4,600
Average Value		4,500	
Range of Values (highest - lowest)		200	
With \$100m Hedged at Rs 45/\$			
Raw materials imports this quarter \$ 25 m at	1,125	1,125	1,125
Rs 45/\$			
Raw material imports next three quarters \$ 75	3,375	3,375	3,375
m at Rs 45/\$			
Total	4,500	4,500	4,500
Average Value		4,500	
Range of Values (highest - lowest)		0	

To consider the other kind of problem with fair value accounting of derivatives, consider a situation where the derivative hedges an asset that the accountant does record in the books of account — inventories. Suppose a jewellery shop decides to hedge the gold price risk of its finished goods inventories by selling gold futures. The accountant does recognize the inventory of jewellery as an asset, but values it at cost

or realizable value, whichever is lower. If the gold price rises³, the realizable value of the inventory does go up, but it remains valued at cost because the cost is less than the realizable value.

Suppose that the jewellery inventory contains 10 kg of gold bought at an average cost of Rs 900 per gram. If the fabrication of the jewellery amounts to 15% of the gold cost and the jeweller adds an additional mark-up while selling it the accountant would, then, value the inventory at $1.15 \times 10,000 \times 900 = \text{Rs} 10.35$ million. When the jeweller sells gold futures, there would be mark to market gains or losses, as the gold price fluctuates. For example, if the gold price rises to Rs 950 per gram, the futures position of 10 kg would incur a loss of $50 \times 10,000 = \text{Rs} 0.5$ million. This loss would be reflected in the quarterly profit and loss account. The inventory would continue to be shown at cost and, therefore, there is nothing to offset this loss. The offset would come in future quarters as the inventory of jewellery is sold at higher prices reflecting the higher gold price.

From the finance theory perspective, the jewellery inventory is effectively valued at market price and the changes in this asset exactly offset the mark to market gains or losses on the futures position. The hedged position is thus risk-free.

Accountants are not prepared to embrace the finance theory view and regard future forecasted transactions as assets or liabilities, let alone record them at fair value. Yet, they recognize that the accounting treatment illustrated in Table 20.1 does appear quite perverse, as the hedge instead of making the, earnings loss volatile, makes them more so. This forces them to move part of the way in the direction of finance theory to achieve reasonable reporting of profits when companies resort to hedging. This is the complex subject of hedge accounting, which is discussed in the next section.

20.3 HEDGE ACCOUNTING

The accounting difficulties highlighted in the previous section arise from the fact that fair value accounting is used for the derivative (the hedge) but not for the hedged item. Accountants may either not recognize the hedged item in the books at all (for example, future raw material imports), or they may recognize it but not use fair value (for example, inventory).

There are two alternative solutions to this problem.

- 1 Fair value or mark to market accounting can be adopted for the hedged item as well. This is basically what is done for what are called fair value hedges.
- 2 The mark to market gains and losses of the derivative can be isolated and prevented from affecting the profits of the company. In this approach, the derivative is still shown in the balance sheet at fair value, but the mark to market gains and losses are deferred. Mark to market gains and losses are allowed to hit the profits of the company only when the offsetting gains and losses on the hedged items are also realized. This is basically what is done for cash flow hedges. This solution works because most companies care about the impact of the accounting on profits and are less worried about the impact on the balance sheet.

The distinction between fair value hedges and cash flow hedges is not just a question of how the accounting is done. It is also a question of the nature of the hedge itself. In a fair value hedge, as the

³ Even if the gold price falls by a modest amount, it is possible that the profit mark-up of the jeweller is sufficiently high that the realizable value of the jewellery at the lower gold price still exceeds the cost of the jewellery. In that case, the jewellery is effectively valued at cost for modest changes in the gold price in either direction.

20.8 | Derivatives and Risk Management

name suggests, what is hedged is an asset or liability whose fair value can fluctuate. In the case of a cash flow hedge, what is hedged are cash flows whose values are uncertain. Since fair value is basically discounted future cash flows, it might appear that cash flow hedges and fair value hedges are really the same. This is not true because the discounted present value of cash flows can change due to changes in the discount rate as well. In the context of interest rate risk in fact, cash flow hedging and fair value hedging can be antithetical to each other, as the following pair of examples illustrates.

Consider a company that has borrowed money using ten-year fixed rate bonds. There cannot be a cash flow hedge for this liability at all because the cash flows are fixed and there is no risk about the quantum of these cash flows. It is however, possible to have a fair value hedge of this fixed rate debt. This is because as the interest rate changes, the present value of the fixed rate liability changes. A receive-fixed interest rate swap can be used to provide a fair value hedge of this debt because the value of this swap changes in the opposite direction as interest rates change. However, the effect of this swap is that it effectively turns the fixed rate debt into a floating rate debt. The fixed rate received from the swap is used to pay the coupon on the borrowing and the effective interest cost for the company is the floating rate paid on the swap. The result is that while the fair value of the liability becomes certain, the cash flows become uncertain.

Conversely, consider a company that has borrowed money using ten-year floating rate bonds. There cannot be a fair value hedge for this liability at all because the floating rate bonds reprice to par at every reset date and there is no risk about the fair value of this debt⁴. It is, however, possible to have a cash flow hedge of this floating rate debt because as the interest rate changes, the cash flows to pay the interest on the floating rate liability do change. A pay-fixed interest rate swap can be used to provide a cash flow hedge of this debt because this swap turns the floating rate debt into effectively a fixed rate debt. The floating rate received from the swap is used to pay the floating coupon on the borrowing and the effective interest cost for the company is the fixed rate paid on the swap. The result is that while the cash flows become uncertain, the fair value of the liability becomes uncertain.

The whole concept of cash flow hedging appears strange to finance theorists, but can appear quite sensible to corporate finance managers. Corporate finance managers welcome cash flow certainty because they argue that the interest bill has to be paid out of the profits of the company. Floating rate debt can lead to a big rise in the interest bill without any offsetting increase in the profits of the company. Thus, floating rate debt appears risky to most finance managers, though its fair value is almost fixed. Fixed rate debt appears safe to most finance managers though its fair value is volatile.

This paradox is easily resolved by probing the implied assumption that the profits of the company do not rise when interest rates rise. In terms of finance theory what is implied is that the operating business of the company behaves like a fixed rate asset — the profits are like a fixed coupon. Well, profits are uncertain but from a hedging perspective what matters is the sensitivity of profits to interest rates. Most operating businesses do not experience any significant change in profits when there is a rise in interest rates — they behave like fixed rate assets in this sense. In many consumer durable and capital goods businesses in facts, profits tend to fall when interest rates rise because the customer finds it more expensive to finance the purchase — these businesses actually behave like inverse floaters⁵.

⁴ There is a small amount of uncertainty about the fair value between reset dates but this can be ignored for practical purposes.

⁵ An inverse floater is a bond whose coupon falls when interest rates rise and vice-versa. For example, the coupon may be defined to be 16 %–MIBOR, where MIBOR is the floating rate benchmark.

In short, operating businesses are not like floating rate assets at all; they are closer to being fixed rate assets. Then the portfolio perspective of finance theory has no difficulty in accepting that a fixed rate liability is safer because it hedges the fixed rate asset represented by the operating business of the company. The next chapter discusses cash flow hedges from the perspective of corporate risk management and explains why they are more important than value hedges for non-financial businesses.

Accountants, however, cannot adopt this interpretation because the future profitability of the business is not an asset that is recognized in their books. The only way in which they can legitimize hedging floating rate debt into fixed, is by bringing in the notion of cash flow hedges, which is difficult to reconcile with finance theory.

The co-existence of fair value and cash flow hedges creates other paradoxes. Companies can hedge fixed rate debt into floating rate debt using fair value hedges and they can hedge floating rate debt into fixed rate debt using cash flow hedges. A financial institution which has a large portfolio of fixed rate and floating rate assets and liabilities can do both at the same time. It can enter into a pay-fixed swap and claim that it is hedging. It can also enter into a receive-fixed swap and claim that it is also hedging. In other words, a bank can turn some fixed rate assets in its portfolio into floating rate and claim that it is hedging. This is because accountants have a transaction view of hedging and not a portfolio view.

In fact, even worse things can happen. The following is an example from the implementation guidance for the international accounting standard IAS 39 which is also reproduced as implementation guidance for the proposed Indian accounting standard AS 30:

'An enterprise has a fixed rate asset and a fixed rate liability, each having the same principal amount. Under the terms of the instruments, interest payments on the asset and liability occur in the same period and the net cash flow is always positive because the interest rate on the asset exceeds the interest rate on the liability. The enterprise enters into an interest rate swap to receive a floating interest rate and pay a fixed interest rate on a notional amount equal to the principal of the asset and designates the interest rate swap as a fair value hedge of the fixed rate asset. Does the hedging relationship qualify for hedge accounting even though the effect of the interest rate swap on an enterprise-wide basis is to create an exposure to interest rate changes that did not previously exist?'

'Yes. IAS 39 does not require risk reduction on an enterprise-wide basis as a condition for hedge accounting. Exposure is assessed on a transaction basis and, in this instance, the asset being hedged has a fair value exposure to interest rate increases that is offset by the interest rate swap.' The transaction view of hedging which accountants adopt, can indeed appear very puzzling and paradoxical to finance theorists accustomed to a portfolio view of hedging.

Accountants are, however, forced to make an exception to their transaction view of things when it comes to what can be hedged. Companies do hedge commitments that are not recorded in the books of account. They also hedge transactions that are forecast to take place, but are not firmly committed. The accounting standards allow firm commitments as well as highly probable forecast transactions to be hedged.

20.3.1 Accounting for Fair Value Hedges

Accounting for fair value hedges is not overly complicated because all that it requires is a fair value accounting of the hedged item. The precise implication of the fair value accounting depends on how the hedged item would have been accounted for in the absence of any hedging:

20.10 | Derivatives and Risk Management

- 1. If the hedged item is a financial instrument which is 'Held for Trading (HFT)', then it is anyway marked to market regardless of whether there is any hedging in place or not. For example, a bank that buys and sells securities as part of its normal trading business is required to use fair value accounting for these securities under the accounting standards that apply to financial instruments in general. For HFT items, the existence of a fair value hedge makes no difference to the accounting. Regardless of the hedging, the derivative is subject to fair value accounting and so is the hedged HFT item.
- 2. Financial assets that are 'Held to Maturity (HTM)' are valued at cost in the absence of hedging. The HTM category consists of non-derivative financial assets with fixed or determinable payments and fixed maturity that an entity has the positive intention and ability to hold to maturity'. If an HTM asset is covered by a fair value hedge, then the asset becomes subject to fair value accounting. Instead of being valued at cost, as it would be if unhedged, it is valued at fair market value and valuation gains and losses are reported as part of the profit or loss for the period.
- 3. In between the HFT and HTM categories is a category of assets that are referred to as 'Available For Sale (AFS)'. AFS assets are neither intended for short-term trading, nor are they intended to be held to maturity. They may be sold under some circumstances. In the absence of hedging, AFS assets are shown at fair value in the balance sheet, but valuation changes are not shown as part of normal profits. Valuation gains and losses are deferred by being shown as a separate reserve or as Other Comprehensive Income (OCI)⁶. When the AFS asset is sold or ceases to exist, the cumulative deferred valuation gains or loss flows into the profits. If an AFS asset is covered by a fair value hedge, then the valuation gains and losses are not deferred but are allowed to flow into profits where they offset the valuation gains and losses of the derivative with which it was hedged.
- 4. Most non-financial assets and liabilities are valued at cost in the absence of hedging. The impact of hedging these items is similar to that of hedging HTM financial assets.

One important complication in fair value hedging is that most hedges have some basis risk. To deal with the problem of basis risk, it is possible to designate a derivative as a hedge of only a component of the changes in value of the hedged item. For example, suppose the company has bought a fixed rate corporate bond and it is being accounted for as an HTM asset. Suppose this bond is hedged using an interest rate swap. The corporate bond can change in value due to two principal reasons: (a) the bond will drop in value if interest rates rise and increase in value if interest rates fall and (b) the bond will drop in value if the creditworthiness of the borrower worsens and rise in value if the creditworthiness improves. The interest rate swap hedges only the interest rate risk of the corporate bond and does not affect its credit risk at all. The way to deal with this basis risk is that when designating the interest rate risk. When basis risk is addressed in this way, the interest rate swap (being a derivative) is of course fully marked to market, but the corporate bond is marked to market only for valuation gains and losses, arising out of interest rate risk.

Suppose that the bond was bought at par at the beginning of the quarter. A month later, an improvement in creditworthiness of the issuer increased the market price of the bond to 104. By the end of the quarter,

⁶ Other Comprehensive Income (OCI) is the term used in US accounting. International accounting standard uses an Investment Revaluation Reserve or some such similar reserve account. Also under certain conditions referred to as 'impairment', the cumulative valuation *loss* of an unhedged AFS asset may be taken out of reserves or OCI, and reflected in the profits.

Derivative Accounting | 20.11

a rise in general interest rates caused the bond to decline to 103. When the company prepares its quarterly accounts, it must first decompose the total valuation gain during the quarter of (103 - 100 = 3) into two components (a) a valuation loss due to rise in interest rates (103 - 104 = -1) and (b) a valuation gain due to improvement in creditworthiness (104 - 100 = 4). Since the hedge covered only the first component, only this valuation loss would be adjusted in the accounting value of the corporate bond. In other words, the bond will be shown in the books at (100 - 1 = 99). The valuation gain of 4 due to creditworthiness will not be reflected in the accountant *reduces* the accounting value of a bond whose market price has *risen*. Such are the vagaries of hedge accounting.

20.3.2 Accounting for Cash Flow Hedges

The idea of cash flow hedge accounting is that valuation gains and losses of the derivative are deferred while the derivative is still shown in the balance sheet at fair value. When a derivative is designated as a cash flow hedge, its accounting treatment is similar to that explained for AFS assets earlier — valuation gains and losses are deferred by being kept in a separate reserve (Hedging Reserve Account) or in Other Comprehensive Income.

At some point in the future when the hedged cash flow impacts the profits of the company, the deferred valuation gains and losses of the hedge are taken out of the reserve account and moved into the profits. The hedged cash flow and the valuation gains and losses of the derivative hit the profits at the same time and can offset each other. At all points of time, however, the derivative continues to be shown at fair value in the balance sheet, but most companies care about the impact of the accounting on profits and are less worried about the impact on the balance sheet.

Basis risk creates a complication in cash flow hedging just as it does in fair value hedging. In fair value hedging, the solution adopted is to divide the valuation gains and losses of the hedged item into two components representing the hedged risk, and the unhedged or basis risk, and apply hedge accounting only to the hedged risk. In cash flow hedging, the method used is the mirror image of this — it is the valuation gains and losses of the *derivative* that are split into an 'effective' component that hedges a cash flow risk and an "ineffective" component that can be thought of as basis risk. The effective component flows into the valuation gains and losses of the derivative is deferred but the ineffective component flows into the profits directly.

The exact implementation of this idea is a little tricky. The exact procedure is that the Hedging Reserve Account is adjusted to the lesser of the following (in absolute amounts):

- 1. The cumulative gain or loss on the hedging instrument from inception of the hedge; and
- 2. the cumulative change in fair value (present value) of the expected future cash flows on the hedged item from inception of the hedge.

Any remaining gain or loss on the hedging instrument is the ineffective component that flows into profits directly.

Consider the following example based on an item discussed in the implementation guidance of IAS 39 as well as the proposed AS 30. An enterprise hedges the anticipated future sale of 24 tonnes of pulp by selling 24 tonnes of pulp forward in a contract that is cash-settled based on the spot price at a commodity exchange. There is a basis risk because the anticipated future sale is in the local market which is different from the commodity exchange. At the end of the quarter, the spot price of pulp has

20.12 | Derivatives and Risk Management

increased in the local market by 100 and on the exchange by 80. Ignoring interest and storage costs for simplicity, assume that spot and forward prices are identical, and that the change in local spot price is also the change in the present value of the anticipated cash flow from the future sale of pulp.

Then despite the evident basis risk, the entire change of 80 in the value of the forward contract is 'effective' because it is less than the change in the value of the hedged cash flow of 100. Therefore, the entire change of 80 in the value of the forward contract will be parked in the Hedging Reserve Account and nothing will flow into profits in that quarter. In the balance sheet, the forward contract has to be marked to market by the full 80 of change in its value. However, since the hedged cash flow is not marked to market in any case, the extra 20 of change in its fair value is not recorded anywhere in the accounting books.

However, in the above example, if the numbers are reversed so that the exchange spot price rises by 100 and the local spot price rises by 80, then the change in the value of the forward contract has to be split into an 'effective' component of 80 (the change in the present value of the hedged cash flow) and the remaining 'ineffective' component of 20. The effective component of 80 will go to the Hedging Reserve Account and the ineffective part of 20 will flow into profits in that quarter. In the balance sheet, the forward contract has to be marked to market by the full 100 of change in its value.

The above analysis assumes that the basis risk is small enough for the enterprise to conclude that there is a highly effective hedge relationship though the relationship, is not perfect. This and other requirements of hedge accounting are discussed in the next section.

Another complication in cash flow hedging is that it is possible to split an option value into a time value and intrinsic value, and designate only the intrinsic value component as a hedging instrument. Similarly, it is possible to split the value of a forward contract into a component relating to the spot price and the component relating to the forward premium (interest and storage costs) and designate only the spot price component as the hedging instrument. When such a splitting is done, the component that is not designated as a hedging instrument (the time value or the forward premium) flows into the profits directly, and only the designated component is subject to cash flow hedge accounting, as discussed above.

Take an example of a company that bought a one-year at-the-money put option at the beginning of the year to hedge a future sale of a commodity, when the spot price was 1000 and the implied volatility was 30%. Assume as in the earlier pulp example, that interest and storage costs are zero, so that spot and forward prices are identical and there is no need for discounting. Suppose the underlying of the option contract is identical to the commodity to be sold and there is no basis risk due to this. There is, however, a significant basis risk in using an option to hedge a cash flow which has no option characteristics. The anticipated sale of the commodity is like a short forward position which has a delta but no gamma, theta, or vega. The put option that has been bought has gamma, theta, and vega and moreover has a delta that is different from ± 1.00 .

Accountants allow the put option to be used as a hedge of this future sale by breaking up the option into its time value and intrinsic value and regarding only the intrinsic value as a hedge. With the data given above, the Black-Scholes formula gives an option value of 119.24. Since the at-the-money put option has zero intrinsic value, the entire option value of 119.24 is time value. Now suppose at the end of the quarter, the spot price has declined to 900. Clearly, the time to maturity is now only 0.75 years. Assuming no change in the implied volatility, the Black-Scholes formula gives an option value of 156.10. The option now has an intrinsic value of 100 and therefore the time value of the option is now 56.10.

Why has the time value dropped by 63.14 from 119.24 to 56.10? There are two reasons for this. First, by sheer efflux of time, the option loses value because there is a shorter time to maturity. The
Derivative Accounting | 20.13

Black-Scholes formula shows that if the spot price had remained at 1000, even then the option value would have dropped by 15.88 from 119.24 to 103.36. This is the impact of theta. What accounts for the balance of 63.14 - 15.88 = 47.26? There is a second reason for change in time value and this is that the option has a delta different from -1.00. An at-the-money put option has a delta close to -0.50 and as the option moves into the money, the option delta moves closer to -1.00 because of the option gamma. So when the spot price falls by 100, the option value rises by a little more than 50. Since the intrinsic value changes by 100, the time value declines by a little less than 50. This explains the remaining 47.26 fall in the time value of the option.

In the accounting treatment, the change in the intrinsic value is allowed to offset the change in the expected future cash flow in accordance with the principles of cash flow hedging. The impact of the delta mismatch, as well as the impact of gamma, theta, and vega (if any) all flow directly into profits. In this example, therefore, the change of 100 in the intrinsic value is deferred by being parked in the Hedging Reserve Account. The change of 63.14 in the time value of the option is a loss that is charged to the profits of the company.

It may appear strange that an option that was bought for 119.24 has appreciated to 156.10, but in the accounting books, the option produces a loss of 63.14. The point is that over the life of the option, the *entire* original time value of 119.24 is bound to be lost. At expiry, the option might have an intrinsic value but zero time value. Over the life of the option therefore, the entire time value will have to be charged to profits. The only impact of the movement of the spot price as well as the implied volatility is to change how this loss is distributed across quarters.

These do not affect the aggregate amount of the time value that is bound to be lost. It is easy to understand this when the option premium (particularly the time value component of the option premium) is interpreted as an insurance premium. Just as an insurance premium has to be amortized over the life of the insurance policy, so also the time value of the option premium has to be amortized over the life of the option. The only difference is that the insurance premium is amortized in equal instalments, while the option time value is amortized in an uneven manner depending on how market prices move over time.

To make this point even clearer, consider what would happen if the spot price rose to 1100 instead of falling to 900. The option moves out of the money and its intrinsic value remains zero. The Black-Scholes formula indicates that the option value drops to 65.70 because it is now out of the money. The entire drop in value is attributable to the change in the time value because the intrinsic value is unchanged. Thus this entire amount of 119.24 - 65.70 = 53.54 is a loss that would be charged to the profits during the quarter.

This example illustrates the proposition that regardless of whether the spot price rises or falls, an *at-the-money* option produces an accounting loss when it is used as a cash flow hedge *if the volatility is unchanged*. If either (a) the money is not at-the-money or (b) the implied volatility rises, an option can sometimes produce an accounting profit in some quarters prior to maturity. However, over the entire life of the option, there is bound to be a loss equal to the original time value.

20.4 REQUIREMENTS FOR HEDGE ACCOUNTING

Application of hedge accounting (whether for fair value hedges or for cash flow hedges) is permitted only if the following conditions are satisfied:

20.14 | Derivatives and Risk Management

- At the inception of the hedge there is formal designation and documentation of the hedging relationship and the entity's risk management objective and strategy for undertaking the hedge. That documentation should include identification of the hedging instrument, the hedged item or transaction, the nature of the risk being hedged and how the entity will assess the hedging instrument's effectiveness in offsetting the exposure to changes in the hedged item's fair value or cash flows attributable to the hedged risk.
- 2. The hedge is expected to be highly effective in achieving offsetting changes in fair value or cash flows attributable to the hedged risk, consistently with the originally documented risk management strategy for that particular hedging relationship.

The tricky part here is assessing hedging effectiveness. This has to be done both prospectively and retrospectively. At inception and in subsequent periods, the company must make an assessment as to whether the hedge is expected to be highly effective in future. This assessment can be made by examining the past data on value changes in the hedging instrument and the hedged item. As discussed in Chapter 4, in finance theory the statistical correlation is the principal measure of hedge effectiveness. Accountants accept this measure but they also rely on another criterion.

They require that the changes in the derivative should be between 80% and 125% of the changes in

the value of the hedge item. Since $\frac{1}{180} = 1.25$, the range of 80% to 125% is symmetrical in a multipli cative sense and this range is therefore consistent with a symmetrical distribution for the logarithm of the value. This criterion is closely related to the correlation. Unless the R-square or squared correlation is above 0.80, the 80%–125% range will be violated quite often.

In each period the actual performance of the hedge must also be examined. Again, the 80%-125% range is a common measure of hedge effectiveness. The problem in measuring hedge effectiveness post facto is that it is possible for a random deviation to cause the ratio to fall outside the 80%-125% range. This is particularly likely when the change in the derivative is very small. For example if the derivative moves by only 1%, and the hedged item moves by 1.3%, the ratio is 130%. However, the 0.3% deviation between the hedge and the hedged item may be economically insignificant and few people may believe that the hedge has been ineffective.

As has already been pointed out, in case of fair value hedges only the changes in value of the hedged item arising out of the hedged risk is subjected to hedge accounting. Hedging effectiveness is measured relative to this change in value. For example, if only the interest rate risk of a bond is hedged and not the credit risk, hedging effectiveness will measure whether the change in the derivative are within the 80%–125% range of changes in the value of the bond attributable to interest rate changes. Similarly, in case of cash value hedges, only the intrinsic value of an option may be designated as a hedge. In this case changes in the intrinsic value must be within the 80%–125% range of changes in the value of the hedged item.

Chapter Summary

Accounting is normally transaction oriented and tends to emphasize transaction values (historical cost) rather than fair value. To account for derivatives in a meaningful way, accountants have to adopt elements of fair value accounting and recognize financial commitments that are not completed transactions. This is pushing accountants to consider fair value accounting for all financial instruments.

Derivative Accounting | 20.15

If derivatives are accounted for at fair value while other items are at historical cost, accounting profits will fail to reflect the benefits of hedging. Accountants have therefore introduced the notion of hedge accounting in which the derivative and the item that it hedges are accounted for in a similar manner. For this purpose, hedges are classified as fair value hedges (that bring certainty to the present value of future cash flows) and cash flow hedges (that bring certainty to the future cash flow hedges of interest rate risk are sometimes antithetical to fair value hedges of the same risk.

In order to adopt hedge accounting, the enterprise must have stringent documentation describing the hedging policy and its implementation. It must also demonstrate hedge effectiveness both at inception and on an ongoing basis. The requirement is that the changes in the derivative should be between 80% and 125% of the changes in the value of the hedge item.

Suggestions for Further Reading

The principal accounting standards dealing with derivatives worldwide are:

Financial Accounting Standards Board (1998), amended 2006), "Statement No. 133, Accounting for Derivative Instruments and Hedging Activities", <u>http://www.fasb.org/pdf/fas 133 .pdf</u>

Institute of Chartered Accountants of India (2007), "Exposure Draft, Proposed Accounting Standard (AS) 30, Financial Instruments: Recognition and Measurment", <u>http://www.icai.org/icairoot/announcements/announ808.pdf</u>

International Accounting Standards Board (1998 amended 2005), "IAS 39 Financial Instruments: Recognition and Measurement"

Recent initiatives to change the accounting standards are described in:

Financial Accounting Standards Board (2007), "Project Updates: Financial Instruments", <u>http://www/fasb.org/project/financial instruments.shtml</u>

20.16 | Derivatives and Risk Management

CASES

Case 20-1 CASE ON HEDGE ACCOUNTING FANNIE MAE: ACCOUNTING FOR INTEREST RATE CAPS

In the last quarter of 2002, Fannie Mae changed its accounting for purchased interest caps. The cryptic disclosure that was made then attracted very little attention, but when Fannie Mae's use of hedge accounting came under intense regulatory scrutiny in 2004, questions began to be raised about this accounting change. The question that arose was whether this was a change in accounting principle that could be applied only prospectively as Fannie Mae had done or whether it was a correction of an accounting error which would require retrospective application by correcting and restating the past accounts also.

Background

In 1938, the US government established the Federal National Mortgage Association (better known as Fannie Mae) to expand the flow of mortgage money by creating a secondary market. In 1968, Fannie Mae became a private company operating with private capital on a self-sustaining basis. However, since it is widely regarded as having an implicit support from the US government, it is able to borrow at very fine rates.

Fannie Mae's charter is to provide stability in the secondary market for residential mortgages, respond appropriately to the private capital market, increase the liquidity of mortgage investments and improve the distribution of investment capital available for financing residential mortgages.

Fannie Mae carries out two main lines of business. In its portfolio investment business, Fannie Mae purchases mortgage loans, mortgage-related securities, and other investments from lenders, securities dealers, investors, and other market participants. It also invests in liquid high-quality non-mortgage securities principally as a source of liquidity and an investment vehicle for surplus capital. The portfolio investment business is funded primarily by borrowing money in the domestic and international capital markets through the sale of debt securities. Fannie Mae profits from the difference, or "spread," between the yield on portfolio investments and the interest paid on borrowings.

In its credit guaranty or securitization business, Fannie Mae does not buy the mortgage loan. Instead, it repackages the mortgage loans as Mortgage Backed Securities (MBS) that are sold to other investors. Fannie Mae assumes the credit risk on these mortgages and guarantees the timely payment of scheduled principal on the MBS. It earns a guaranty fee for this service.

Exhibits 1 and 2 provide key financial data for Fannie Mae.

Fannie Mae uses derivatives on a large scale to manage interest rate risk and supplement the issuance of debt in the capital markets. A mix of debt issuances and derivatives can be used to achieve the same interest rate risk management objectives that would be achieved by issuing only debt securities. At December 31, 2002, Fannie Mae had derivatives positions amounting to \$657 billion in notional value as shown in Exhibit 3. Exhibit 4 provides an explanation of the hedging function provided by the different types of derivatives used by Fannie Mae.

Copyright © 2005 by the Indian Institute of Management, Ahmedabad.

Reproduced with permission.

Prepared by Prof. Jayanth R. Varma on the basis of published material. Teaching material of the Indian Institute of Management, Ahmedabad, is prepared as a basis for class discussion. Cases are not designed to present illustrations of either correct or incorrect handling of administrative problems.

Interest Rate Caps

Interest rate caps are used to ensure an upper bound on the interest rate paid on floating rate debt. The periodic interest payments on floating rate bonds are not fixed in advance but are linked to a benchmark short term interest rate. For example, a company that issues a seven year floating rate bond may agree to pay interest every quarter at a rate equal to half percent more than the three-month Libor interest rate that prevailed at the beginning of that quarter. If the Libor rate at the beginning of a quarter was only 4%, the company would pay interest for that quarter only at the rate of 4.5% per annum. Two years later, if Libor has climbed to 8%, it would have to pay interest at the rate of 8.5% per annum. Since short term interest rates can be very volatile, companies are concerned that their interest bill could shoot up to very high levels. They may therefore want to ensure that their interest rate does not go too high.

One way to ensure that the interest payment never exceeds 6.5% would be to buy a Libor cap at 6%. In this contract, the cap seller would at the end of each quarter pay the cap buyer the excess of quarterly interest at the Libor rate over the quarterly interest at the rate of 6%; no payment is made in either direction if Libor is below 6%. So if interest rates do rise to 8%, the cap seller would pay the company quarterly interest at the rate of 2% per annum (excess of 8% over 2%). Deducting this from the 8.5% interest that the company would pay on the floating rate bond, the net interest outgo is only 6.5% as desired.

Interest rate caps are regarded as a collection of caplets—one for each quarter or interest payment period. A seven-year cap with quarterly payments would have 28 caplets—one for each quarter in the life of the cap. The payment in a specific caplet is the computed by applying the excess of Libor in that quarter over the cap rate to the cap principal. More precisely, it is equal to the principal times ¹/₄ times the excess of Libor in that quarter over the cap rate. Each caplet can therefore be regarded as a call option on Libor with a strike rate equal to the cap rate. The cap as a collection of caplets is thus a collection of call options on Libor. Therefore the cap also has many of the features of an option.

Time Value and Intrinsic Value

The value of any option can be broken into two parts—a time value and an intrinsic value. Consider for example a caplet for some future quarter with a cap rate of 6%. If the current interest rate is 7%, one can see that the caplet clearly has a lot of value since at current interest rates, it would give us a payment of 1% applied to the cap principal. This value discounted to its present value today is known as the intrinsic value of the caplet. The value of the caplet may however be significantly more than the discounted value of 1% applied to the cap principal because by the time we reach the relevant quarter, interest rates could well have climbed to 8%, 9% or even 10% and we could be getting payments of 2%, 3% or even 4% under the caplet. If on the other hand, interest rates were to fall as we reach the relevant quarter, the worst that can happen is that the 1% payment that we expected might disappear—the payment cannot be negative. To take a more extreme case, suppose the current interest rate is only 3%. The intrinsic value of the caplet is zero since the interest rates may rise between now and the caplet date. This excess of the caplet value over the intrinsic value is known as the time value of the caplet.

More precisely, the intrinsic value must be computed using what is called the forward interest rate. For example, a company wants to borrow money for the quarter beginning January 1, 2010. Though the loan would be availed of only in 2010, it wants to fix the interest rate right now in 2005. The interest rate that a lender would agree today for that loan to be made in future is called the forward rate

20.18 | Derivatives and Risk Management

in 2005 for the quarter beginning January 1, 2010. If interest rates are expected to rise in future, the forward rate could be much more than the rate of interest for a three month loan in 2005 itself. In practice, forward rates can be derived from long term interest rates and we do not have to go around seeking quotations for future loans.

One implication of using forward rates is that some of the caplets may have positive intrinsic value while some other caplets have zero intrinsic value because the relevant forward rates are different. Some of these forward rates may be less than the cap rate (zero intrinsic value) while others may be higher (positive intrinsic value).

Decomposing each caplet in this way also leads to a decomposition of the total cap price into time value and intrinsic value. The intrinsic value of the cap is the sum of the intrinsic values of the caplets. Similarly, the time value of the cap is the sum of the time values of the caplets.

The intrinsic value of a cap reflects the value arising out of current or expected future short term interest rates being above the cap rate. The time value on the other hand reflects the uncertainty about future short term interest rates. If there were no such uncertainty and we were sure that the future short term interest rates would always be equal to the forward rates, then the time value of the cap would be zero. The price of the cap would then be equal to the intrinsic value. For this reason, the intrinsic value of a cap is often described as the zero volatility value of the cap. It is the value that would arise if the volatility (which is a mathematical measure of the degree of uncertainty) were equal to zero.

However, cap traders often compute intrinsic and time values of a cap using a much simpler but cruder device which we shall call the "trader" method. They do this by comparing the cap rate with the going market interest rate for the same maturity as the cap. In the example above, if the seven year interest rate is 6.7%, then the intrinsic value of a cap at 6% is the present value of a series of payments each quarter equal to 0.7% times the principal.

This method (the trader method) can sometimes give quite incorrect results. For example, suppose the seven year interest rate is only 5.8% and the cap rate is 6%. The trader method says that the intrinsic value of the cap is zero. If we used the caplet method, we might see a very different answer. Suppose for example, the current short term interest rate is only 4%. Why is the seven year rate then as high as 5.8%? The most likely reason is that interest rates are expected to go up sharply. If over the next seven years, the short term interest rate is expected to average about 5.8%, then the market rate of 5.8% makes sense. In other words, the average of the forward rates for all the quarters in the next seven years may well be around 5.8%. Since in the initial quarters the rate is well below 5.8% (only 4% in the-first quarter), the rate in some future quarters will have to be well above 5.8% for the average to be 5.8%. Quite likely, in some quarters the forward rate would be above 6%. In these quarters, the caplet would have a positive intrinsic value. Adding up these positive caplet intrinsic values gives a positive intrinsic value of the cap also in the caplet method. Thus the trader method wrongly tells us that the intrinsic value of the cap is zero when it is in fact positive. In general, the trader method tends to underestimate the intrinsic value and overestimate the time value of the cap.

Hedge Accounting

The decomposition of the caplet value into a time value and an intrinsic value is important for accounting purposes. SFAS 133 which deals with the accounting for derivatives requires all derivatives to be valued at fair value on each balance sheet date. If the market price of the purchased cap rises or falls, the balance sheet would also reflect this rise or fall in value. Corresponding to this, of course the income statement would reflect a profit or loss. If a cap was bought for \$1.5 million and a year later its price is only \$0.9 million, the difference of \$0.6 million is a loss that has to be recorded.

Derivative Accounting | 20.19

Hedge accounting provides a partial escape from having to record this loss in the income statement. The idea in hedge accounting is that the purchased cap is used to hedge the risk of the floating rate loan. In years in which the floating rate is very high, the cap is used to pay a part of the interest bill. Hedge accounting therefore allows a part of the changes in the cap price to be offset against future interest payments.

In accounting terms, the hedge of a floating rate note with a cap is described as a cash flow hedge. The fair value of a floating rate note is always close to its par value and there is no need to hedge its fair value. The objective of the cap is to hedge the variability in the cash flows (future interest payments) of the floating rate note. Changes in the intrinsic value of a cap closely mirror the changes in expected future interest payments of the floating rate note. However, there is no such close relationship between the time value of the cap and the future interest payments of the floating rate note. Therefore, FAS 133 allows only the intrinsic value of the cap to be used as a hedge of the floating rate note. The time value is not regarded as a hedge and quarterly changes in the time value of the cap must therefore be recorded as profits or gains in that quarter.

What happens under hedge accounting is that the fair value of the cap is computed based on market prices or valuation models, and this fair value is then decomposed into intrinsic value and time value. This analysis has to be done when the cap is first purchased and thereafter at every balance sheet date. Changes in the time value of the cap appear as gains or losses in the income statement. Changes in the intrinsic value of the cap do not flow into the income statement—they are recorded as part of the reserves (other comprehensive income) under shareholders equity in the balance sheet.

Fannie Mae's Accounting

Fannie Mae's financial statements for 2002 contained the following cryptic disclosure:

"During the fourth quarter of 2002, we refined our methodology for estimating the initial time value of interest rate caps at the date of purchase and prospectively adopted a preferred method that resulted in a \$282 million pre-tax reduction in purchased options expense and increased our diluted EPS for 2002 by \$.18. Under our previous valuation method, we treated the entire premium paid on purchased "at-the-money" caps as time value with no allocation to intrinsic value. Our new methodology allocates the initial purchase price to reflect the value of individual caplets, some of which are above the strike rate of the cap, which results in a higher intrinsic value and corresponding lower time value at the date of purchase. This approach is more consistent with our estimation of time value subsequent to the initial purchase date. This change does not affect the total expense that will be recorded in our income statement over the life of our caps and has no effect on our non-GAAP core business earnings measure." If one read this disclosure carefully, the following inferences can be made:

- 1. Prior to 2002, Fannie Mae was using the short cut "trader" method of computing intrinsic value at the time of purchasing the cap.
- 2. Prior to 2002, Fannie Mae was using the more involved "caplet" method of computing intrinsic value at subsequent balance sheet dates.
- 3. From 2002, Fannie Mae started using the "caplet" method at the time of purchasing caps also.
- 4. The accounting of caps bought prior to 2002 was not changed.

Regulators reviewing the matter in 2004 were concerned whether this method of handling the change is correct. APB 20 is the principal guidance on the subject in the United States. Under certain conditions, APB 20 allows prospective application of:

20.20 | Derivatives and Risk Management

- Changes in accounting principle (including changes in the methods used to apply these principles)
- Changes in accounting estimates (for example a change in the estimate of the useful life of an asset for computing depreciation)

However, APB 20 distinguishes errors in accounting from the changes in accounting described above. Errors in financial statements result from mathematical mistakes, mistakes in the application of accounting principles, or oversight or misuse of facts that existed at the time the financial statements were prepared. In contrast, a change in accounting estimate results from new information or subsequent developments and, accordingly, from better insight or improved judgement.

Thus, an error is distinguishable from a change in estimate. A change from an accounting principle that is not generally accepted to one that is generally accepted is a correction of an error. Reporting on the correction of an error requires retroactive restatement of prior period financial statements and disclosure of the nature of the error and the effect of the correction on income before extraordinary items and net income.

The question was how Fannie Mae's change of accounting for interest rate caps should have been accounted—prospectively or retroactively.

	Year ended 31December				
	2002	2001	2000		
Interest income	50,853	49,170	42.781		
Interest expense	(40,287)	(41,080)	(37,107)		
Net interest income	10,566	8,090	5,674		
Guaranty fee income	1,816	1,482	1,35 i		
Fee and other income (expense), net	232	151	(44)		
Provision for losses	(128)	(94)	(122)		
Foreclosed property income (expense)	36	16	28		
Administrative expense	(1,219)	(1,017)	(905)		
Special contribution		(300)			
Purchased options expense (Note 1)	(4,545)	(37)			
Debt extinguishments, net	(710)	(524)	49		
Income before federal income taxes and cumulative effect of change in accounting principle	6,048	7,767	6,031		
Provision for federal income taxes	(1,429)	(2,041)	(1,583)		
Income before cumulative effect of change in accounting principle	4,619	5,726	4,448		
Cumulative effect of change in accounting principle, net of tax effect (Note 2)		168			
Net income	4,619	5,894	4,448		
Preferred slock dividends	(99)	(138)	(121)		
Net income available to common stockholders	4,520	5,756	4,327		
Basic earnings per common share	4.56	5.75	4.31		
Diluted earnings per common share	4.53	5.72	4.29		
Cash dividends per common share	1.32	1.20	1.12		

Exhibit 1 Fannie Mae: Income Statement Data

Year ended 31 December

Notes

1. Represents the change in the fair value of the time value of purchased options under FAS 133, "Accounting for Derivative Instruments and Hedging Activities" (FAS 133).

2. Represents the after-tax effect on income of the adoption of FAS 133 on January 1, 2001.

Exhibit 2 Fannie Mae: Balance Sheet Data

	Year ended 31 December					
	2002	2001	2000			
Mortgage portfolio, net	797,693	705,324	607,551			
Liquid assets	61,554	76,072	55,585			
Total assets	887,515	799,948	675,224			
Borrowings:						
Due within one year	382,412	343,492	280,322			
Due after one year	468,570	419,975	362,360			
Total liabilities	871,227	781,830	654,386			
Preferred stock	2,678	2,303	2,278			
Stockholders equity	16,288	18,118	20,838			

Exhibit 3 Derivative Notional Amounts (\$ millions)

	Year ended 31 December				
	2002 2001				
Pay-fixed swaps	168,512	213,680			
Receive-fixed swaps	52,370	39,069			
Basis swaps	25,525	47,054			
Caps and swaptions	397,868	219,943			
Other	12,320	13,393			
Total	656,595	533,139			

Exhibit 4 Fannie Mae: Usage of Derivatives

Derivative Instrument	Hedged Item	Purpose of the Hedge Transaction
Pay-fixed, receive- variable interest-rate swap	Variable-rate debt Anticipated issuance of debt	To protect against an increase in interest rates by converting the debt's variable rate to a fixed rate
Receive-fixed, payvariable	Noncallable fixedrate debt	To protect against a decline in interest rates.
variable interest rates swap		Converts the debt s fixed rate to a variable rate
Basis swap or spread-lock	Variable-rate assets and	To lock in or preserve the spread between
	liabilities	variable-rate, interest-earning assets and
		variable-rate, interest-bearing.

20.22 Derivatives and Risk Management

Derivative Instrument	Hedged Item	Purpose of the Hedge Transaction
Pay-fixed swaption	Variable-rate debt	To protect against an increase in interest rates by having an option to convert floating-rate debt to a fixed rate
Caps	Variable-rate debt	To protect against an increase in interest rates by providing a limit on the interest costs on our debt in a rising rate environment
Receive-fixed swaption	Noncallable fixed- rate debt	To protect against a decline in interest rates by having an option to convert fixed-rate debt to floating-rate debt
Foreign currency swaps	Foreign currency- denominated debt	To protect against fluctuations in exchange rates on non-U.S. dollar- denominated debt by converting the interest expense and principal payment on foreign-denominated debt to U.S. dollar-denominated debt

Chapter **Twenty One**

Corporate Risk Management

Objectives

This chapter picks up the discussion of corporate risk management from where it was left off in Chapter 5. The purpose is to analyse risk management within the framework of modern finance theory which regards shareholder value maximization as the principal objective of the company. The central perspective in this chapter is that of the well-diversified investor of finance theory. The chapter also discusses the perspectives of other stakeholders, particularly in the context of financial distress. This theoretical discussion leads naturally to a re-examination of cash flow hedges and value hedges. Finally, the linkage between risk management and capital structure are explored.

21.1 RISK MANAGEMENT AND SHAREHOLDER VALUE

A company is owned by its shareholders who are residual claimants to its assets and cash flows. In other words, shareholders receive whatever is left over after all suppliers and employees have been paid. This means that the risks assumed by a company are largely borne by its shareholders. It is only in extreme situations like bankruptcy or financial distress that the risks fall on other stakeholders of the company. Any analysis of corporate risk management must, therefore, start with asking what hedging policies would shareholders like the companies to adopt. This is particularly important because finance theory regards shareholder value maximization as the principal objective of the company.

Finance theory provides the surprising answer that shareholders do not care. Finance theory provides three arguments for this. First is that investors should hold diversified portfolios of shares. They should not put all their money into one basket but should spread it across a large number of companies in different industries and geographies. In such a diversified portfolio, many risks disappear. For example, an exporter worries about an appreciation of the domestic currency that makes its products uncompetitive in global markets. An importer worries about a depreciation of the domestic currency that makes its imported goods prohibitively expensive. A shareholder who holds shares in both companies might contemplate the future with a great deal of equanimity. If the rupee appreciates, that is bad for the exporter, but good for the importer and vice-versa. The shareholder may be quite happy if both the exporter and the importer left their currency risk unhedged. The shareholder's diversified portfolio provides a natural hedge against currency risk.

The second reason why shareholders do not care about risk management is that shareholders can hedge any risk that is not diversified. Suppose a shareholder holds a portfolio in which there are a large number of exporters, but few importers. In this case, there is no natural hedge against appreciation of the domestic currency. However, if exporters do not hedge this risk, the shareholder can hedge it by using the currency forward and futures markets.

It might be objected that individual shareholders may not have access to some of these hedging instruments or might face large transaction costs in accessing these markets. For example, in India access to currency derivatives is largely limited to those who can demonstrate an underlying exposure and many shareholders might not be able to meet this test. Even if they could, retail customers face

21.2 | Derivatives and Risk Management

steep transaction costs in the foreign exchange market. This argument becomes less convincing when it is recognized that a large proportion of shares is held by institutional investors who have easy access to derivative markets. For example, a large fraction of the shares of some of India's largest software exporters is held by foreign institutions who have ready access to currency hedges.

The third argument is that shareholders often choose undiversified portfolios precisely because they want exposure to certain kinds of risks. For example, some investors might decide to be overweight on stocks of gold mining companies because they are bullish on the outlook for gold. If the gold company hedges its gold price risk by selling gold forward, it makes its stock less attractive to these investors. Similarly, investors who are bullish on a particular currency might buy stocks of leading exporters in that country. These investors would not then want these exporters to hedge the currency risk.

These arguments imply that shareholders would like detailed information about the risk management policies adopted by different companies so that they can take this into account while choosing their portfolios as well as while designing the hedges that they overlay on these portfolios. So long as there is adequate and comprehensive disclosure about risks and risk management policies, investors may not care much about the precise form of risk management that is adopted.

Modern financial reporting standards require a fairly detailed disclosure of risks and risk management policies. However, even these disclosures may not be sufficient for the shareholders to implement the optimal hedges. This is because the disclosures do not usually go beyond listing the risk factors and providing some data on notional values. Precise data on maturities, probabilities of various contingencies, and precise structures of option-like exposures are typically lacking. Shareholders can ignore this missing information and implement approximate hedges. This may be adequate if their portfolios are diversified or only mildly concentrated.

21.2 LENDERS, EMPLOYEES AND OTHER STAKEHOLDERS

If shareholders do not care much about risk management, then it becomes necessary to examine the perspective of other stakeholders. These stakeholders may bear less risk than shareholders who, as residual claimants, are exposed to the bulk of the risks. But if they are more worried about the risks that they do bear, then they might penalize a company for poor risk management. The company might then choose to adopt better risk management even if shareholders do not reward the company for doing so.

For example, it is common for secured lenders to require that the borrower insure the mortgaged assets against fire and other hazards. The rationale for this is clear. If a fire destroys an uninsured factory that was mortgaged as security for a loan, then the lender suddenly becomes an unsecured creditor at a time when the company's creditworthiness has been damaged by fire losses. The lender thus has a strong desire to demand that insurance be kept in force. When a borrower is in financial distress, it is not unknown for the lender to pay the insurance premium to protect the value of the collateral for the loan.

For the same reason, lenders might want the company to adopt other forms of risk management as well. For example, a bank lending against commodity inventories might require that the commodity price risk be hedged in the futures or forward markets. Unsecured lenders might also impose limits on unhedged currency and commodity price risk to reduce the incidence of financial distress among their borrowers. In other cases, they might not impose any particular risk management policy but might simply charge a higher rate of interest if a company has inadequate risk management.

Employees might be even more risk-averse. While shareholders can hold diversified portfolios of shares, employees do not hold diversified portfolios of jobs. Financial distress on the part of their employer carries with it the threat of unemployment. Certain forms of health and retirement benefits might also be lost in the event of bankruptcy of the employer. Employees will therefore prefer to work for companies with prudent risk management policies. Riskier enterprises can attract highly talented employees only by offering higher remuneration.

Customers and suppliers might also have a stake in the continued survival and financial health of the company. Many durable goods require extensive support, maintenance, and supply of spare parts during their economic life. Bankruptcy of the manufacturer can disrupt this process and make the equipment effectively unserviceable. In one well known instance a couple of decades ago, a car manufacturer on the verge of bankruptcy, found its customers deserting it because of the uncertainty about after sales service and maintenance. In less extreme cases, customers might be willing to pay a higher price for a durable good manufactured by a financially prudent manufacturer. Indirectly, therefore, these customers pay for the higher hedging costs incurred as a result of this prudence.

Suppliers might also impose a risk premium on a company with poor risk management. This situation arises when the suppliers have to incur front-end costs in customizing their products or manufacturing facilities to suit the needs of the company. For example, a supplier would be willing to set up a dedicated production line or logistics facility if it is confident that it can recoup the investment over the life of the supply contract. This requires that the company to whom it is supplying the product, remains financially healthy enough to honour its part of the supply contract. If the supplier is not confident about this, it might either not set up dedicated facilities or charge a risk premium as part of the pricing. In either case, the company faces a higher cost or poorer quality for its raw material and components.

The government as a stakeholder also imposes a risk premium on companies that do not manage their risks well. This happens through the operation of the tax system. The government collects a share of the profits of the company as taxes. It shares in the losses only to a limited extent. In India and many other countries, losses can be carried forward and set off against future profits. In this system, a company that makes a loss does not get an immediate tax refund. All that is available is a potential reduction in the tax bill in future years. This is worth less than an immediate refund for two reasons. First, is the time value of money which makes future tax reductions much less valuable than immediate cash. Second, is that the reduction in future tax bill arises only if the company survives and makes profit in future years.

Therefore, the government shares in losses to a much lower extent than it shares in profits¹. This implies that a risky company which earns huge profits in some years and suffers large losses in other years, ends up paying a much higher effective tax rate on its cumulative profits. A company with better risk management that produces a steady stream of profits has a lower effective tax rate.

The sum and substance of all this discussion is that even if shareholders do not reward a company for sound risk management, other stakeholders are likely to reward the company in the form of lower interest costs, lower salaries, lower prices for raw materials and components, lower taxes and higher prices for the finished goods.

21.3 PLANNING AND CONTROL REASONS FOR HEDGING

The above discussion has focused on outside stakeholders rewarding or penalizing a company for its risk management practices. The behaviour of internal management is equally important in determining the risk management policy.

¹ In some countries, the government allows losses to be carried back to past years to generate a refund of taxes paid in those years. But this facility is subject to numerous restrictions and it remains true that the government's share of losses is less than its share of profits.

21.4 | Derivatives and Risk Management

Risk management simplifies internal planning and control, as explained in Chapter 5:

- Hedging brings greater certainty about the cash flows and makes it easier to plan for cash requirements.
- Hedging eliminates the effects of exchange rate and other risk factors, and makes it easier to evaluate the operating efficiency and performance of different divisions.

Many managers express these ideas in terms of the notion of core competence. The core competence of most companies lies in designing, manufacturing, and marketing their products. Speculating on exchange rates, interest rates, or commodity prices is not part of the core competence of most companies. Many financial investors and speculators would have greater ability to estimate demand and supply, and forecast future prices of most financial assets than companies.

For example, a financial speculator might have greater ability to forecast future gold prices than a gold mining company. The core competence of the gold mining company might be in exploring gold reserves or in surmounting the geological challenges of mining the gold under unfavourable conditions, or in optimizing the recovery of gold from low grade ores. The company might be better off focusing on this core competence and leaving speculation on gold prices to financial speculators. This is achieved by hedging the gold price risk and concentrating all managerial energies on the mining business itself.

21.4 FINANCIAL DISTRESS

A very useful way of analysing risk management practices is to distinguish between large risks and small risks. Whether a company does or does not hedge small risks makes little difference to others. As already pointed out, shareholders may not care because they can diversify or hedge the risks in any case. Other external stakeholders do not bother about small risks because these fall on shareholders as residual claimants. Only large risks can wipe out the shareholders and cascade down to other stakeholders risk lenders, employees, suppliers, and customers.

The people who worry about small risks are the internal managers who are concerned about the impact of these risks on planning, internal control, and performance evaluation. Thus the hedging policy regarding small risks is largely a function of the tastes and styles of management.

Large risks are different because they are of concern to external stakeholders. One way to conceptualize this phenomenon is in terms of financial distress. Large risks have the potential to cause financial distress or even bankruptcy. Financial distress and bankruptcy carry with them a variety of costs:

- The direct costs of bankruptcy including court expenses and fees paid to accountants, lawyers, administrators, and valuers are normally quite small amounting to not more than 10% of the assets of the company in most cases. This is typically dwarfed by the other costs described below.
- Financial distress often severely damages the brand value of the company. Often the only way to protect the value of the brands is to sell the brands quickly before the financial distress erodes the brand value. One reason for this damage was mentioned earlier in this chapter many products require after sales service and if the survival of the company is in doubt, these products become less valuable to the customers. Another reason is that the value of many luxury brands is due to the reputation that the product acquires through intensive advertising. If distress and bankruptcy lead to a cutback or cessation of this advertising, then the brand loses much of its allure. A third reason is that bankruptcy carries with it the stigma of failure and some of this stigma attaches to the brand as well.

- Financial distress often leads to a loss of key employees. The reason for this was also alluded to earlier in this chapter employees fear that bankruptcy could threaten their jobs and therefore find it more prudent to seek other jobs well in advance. In knowledge-intensive businesses, incipient distress leads to such a large exodus of talent that the company is either sold or simply ceases to exist.
- Financial distress damages the company's relationship with key suppliers who may be unwilling to invest in customization of products or processes in the absence of any confidence that the company would survive. This leads to an increase in costs or deterioration of quality of purchased inputs. This was also mentioned earlier in this chapter.
- Financial distress can wipe out the value of many intangible assets. The brand value has already been mentioned earlier in this list, but a similar effect can be observed for many other intangibles. For example, consider a company that has invested in an ongoing research and development programme that is very close to creating new products. If the company remains solvent, this programme is enormously valuable. But financial distress makes it difficult for the company to complete the research and development programme and therefore makes the new products less likely. Even if the products do arise, they are less valuable if the company lacks the financial strength to aggressively market and promote these products. The situation here is worse than that of a brand or patent that can at least be sold. A partially complete research and development programme cannot be sold because it is difficult to transfer the informal or tacit knowledge embedded in this programme to a potential buyer.
- Financial distress can derail the implementation of corporate strategy. Long-term strategic plans often require investment over several years to develop new markets, technologies, or competencies and this investment can be difficult to sustain during periods of distress. When a long-term strategy is abandoned midway, the investment that has already been made is often rendered worthless.

It is evident that the costs of financial distress can be very large. Companies would therefore go to great lengths to avoid financial distress. Risk management is an important element of any such attempt.

21.5 CASH FLOW HEDGES AND VALUE HEDGES

The above analysis of financial distress brings out the importance of cash flow. Many of the costs of distress arise from the inability of the company to sustain the additional investments that are needed to extract value from the investments that have already been made:

- Sustaining the brands
- Completing ongoing research and development programmes
- Implementing long-term strategic plans

All these costs can be reduced if there are adequate cash flows to finance these investments, without recourse to external sources of funding. Cash flow shortfall is also the proximate cause of bankruptcy. Even under highly adverse conditions, companies can survive without falling into bankruptcy if they have adequate cash flows to meet interest payments, debt repayments, and other commitments as they arise. Cash flows, therefore, mitigate many of the distress costs that arise from the threat of bankruptcy:

21.6 | Derivatives and Risk Management

- Direct costs of bankruptcy
- Inability to retain key employees
- Damage to critical supplier relationships

In practice therefore, financial distress costs can be largely avoided if the company manages to maintain adequate cash flows even under adverse conditions. Much of corporate risk management is directed towards this goal of ensuring that cash flows do not drop below a minimum acceptable level, consistent with ensuring the survival of the company and sustaining its strategic plans.

The previous chapter discussed cash flow hedges and value hedges from an accounting point of view. The above discussion casts this in a very different light. Cash flow hedges bring greater certainty to cash flows and are thus very important from the perspective of reducing financial distress. Value hedges on the other hand are less important. They would be immensely important if shareholders were exteremely concerned about risks. But if shareholders are well diversified or have access to hedging opportunities, then value hedges offer little or no benefits.

Much of corporate hedging is therefore directed towards cash flow hedges. Financial institutions are in a different category. The nature of their business is such that their liquidity needs are rarely met out of operating cash flows. Liquidity in a financial business comes largely by selling liquid assets or raising new debt. In either case, the liquidity of a financial institution depends on its solvency (the value of its assets exceeding the value of its liabilities) at every point in time. From the perspective of these institutions, therefore, value hedges are critical while cash flow hedges are relatively unimportant.

The analysis of financial distress suggests moreover, that cash flow hedges in non-financial businesses should be designed to avoid large cash flow shortfalls. Avoidance of relatively small cash flow uncertainties is relatively unimportant as these have little bearing on the chances of financial distress.

21.6 CAPITAL STRUCTURE AND HEDGING

The conceptualization of risk management in terms of avoiding financial distress is also useful in exploring the linkage between risk management and capital structure. Modern finance theory regards the optimal capital structure as being the result of balancing financial distress costs against the tax advantages and incentive benefits of debt. In the absence of distress costs, the optimal capital structure would be weighted heavily towards debt because interest is a tax deductible expense and because debt is a powerful tool for incentivizing and disciplining managers. Companies choose to have a significant amount of equity mainly as a way of reducing the chance of financial distress.

Risk management can thus be seen as a substitute for equity. By carefully measuring exposures to various forms of risks and mitigating them using derivatives and other tools, risk management makes the cash flows of the business more predictable and stable. These stable cash flows can support a larger amount of debt without causing an excessive risk of financial distress.

Large amounts of equity capital on the other hand, are a very simple form of risk management. There is no need to measure risks carefully and to design well calibrated hedging strategies to manage these risks. A large base of equity capital is a general purpose risk management tool. It protects the company against all forms of risks including the risks that the company has not even identified and measured.

Capital structure is thus a very effective and omnibus risk management device. In fact, it is the ultimate risk management device in all companies. In all businesses, there are some risks that cannot be anticipated and measured, and there are some risks that cannot be hedged even after they have been measured. Equity capital is the only defence against these risks.

This suggests a linkage between the risk management practices of a company and its capital structure. A financially strong company with plenty of capital can afford to leave most of its risks unhedged, secure in the knowledge that its large capital is an adequate defence against financial distress. On the other hand, a financially weak company with a high level of debt needs a highly sophisticated risk management system to mitigate the possibility of financial distress. This is often what is observed in practice.

This linkage works the other way as well. A company needs a well-developed risk management system to be able to borrow large amounts of debt. In the absence of such risk management systems, lenders would find the company too risky and would either not lend to it at all, or would lend only at prohibitively high rates of interest.

The next chapter discusses financial institutions which are special cases of this phenomenon. By definition, financial institutions operate with low levels of equity capital, relative to the size of their balance sheet. Risk management is then critical for their survival. Many financial institutions regard risk management as their core competence.

21.7 IS THE RISK DEPARTMENT A PROFIT CENTRE OR A COST CENTRE?

Sometimes a large company with significant derivatives activities regards its treasury as a profit centre and expects the treasury to make profits by trading derivatives. It must be recognized that in doing so, the company is actually becoming in part, a financial company. The treasury in this framework is a business division engaged in the financial services business. The company must then manage this division exactly the same way that a finance company or a financial services division would be managed. The derivative business of a corporate treasury which is managed as a profit centre must be managed no differently than the derivative business of a bank or financial business.

This means that value hedges and mark to market analysis become more important than cash flow hedges and held to maturity analysis. Moreover, all the risk management methods described in the next chapter (including VaR modelling, stress testing, back testing and a strong mid-office) must be applied to this kind of a corporate treasury. The derivative activities of such a treasury would not be confined to hedging but may include speculation, arbitrage, and trading exactly as in the case of a bank or other financial institutions.

The other alternative is to regard the risk department of the corporate treasury as a cost centre whose job is to manage the risks of the business at the lowest possible hedging costs. Just as the company regards the fire insurance premium that it pays as a cost and does not attempt to earn a profit by buying and selling insurance policies, it also regards the hedging costs as a business expense. In this case, the elaborate risk management techniques described in the next chapter are not necessary. It is, however, important to ensure that the derivative activities of such a treasury are purely for the sake of risk management.

Derivative activities of such a cost-centre treasury must be confined to hedging activities and must be subject to well-defined risk management and hedging policies laid down by the board of directors:

1. The minimum and maximum extent of hedging must be laid down. To give the widest amount of freedom to management, the minimum level of hedging may be zero and the maximum level of hedging may be 100%. This means that the company cannot over-hedge by buying derivatives when there is no underlying exposure. Similarly, it cannot enter into transactions that increase the

21.8 | Derivatives and Risk Management

risk rather than reduce it. Within these limits, management has a great deal of freedom. At the other extreme, the policy may require 100% hedging at all times (both minimum and maximum set to 100%). This gives no freedom to operating management at all. In between these extremes, the board of directors may lay down limits which are neither too wide nor too narrow. It may, for example, require that the hedging be between 30% and 100%. This would prevent the company from ever being completely unhedged, but would give a fair deal of freedom to management.

- 2. The hedging policy must lay down the types of derivatives that are permissible. The objective is to ensure that the company does not buy derivatives that it does not understand, or does not have the expertise to analyse adequately. This would avoid derivative disasters like the Proctor and Gamble example discussed in Chapter 1.
- 3. The company must be able to value the derivatives that it buys. In case of plain vanilla derivatives, it might have in-house valuation models for this purpose. For more complex derivatives, it must ensure that the derivative is traded by a at least 4-6 large dealers so that valuation can be obtained at any time by seeking competitive quotes.
- 4. The hedging policy must lay down the counterparty limits for derivative transactions. Large companies may need derivative trades of a large size and unless these are done with large sound banks or other strong counterparties, there is a risk of default by the counterparty.
- 5. A derivative that is bought for hedging purposes would normally be held till maturity. The company must, therefore, make sure that the derivatives that it buys, make sense if held to maturity. A derivative that makes sense only for short-term profit making is not a hedging instrument and is inappropriate for the cost-centre oriented treasury.
- 6. The company must also make sure that the mark to market cash flows of the derivative are acceptable in the light of the liquidity position of the company. It must also ensure that if necessary, the derivative can be unwound prior to maturity. Derivative disasters like Metallgesselschaft discussed in Chapter 1 can be avoided by paying attention to this.

Chapter Summary

Finance theory requires that risk management be analysed in terms of shareholder value maximization, which is often regarded as the principal objective of the company. This leads to the surprising answer that shareholders do not care because they are well diversified and even if they were not, they could hedge risks as easily as the company itself. Moreover, concentrated portfolios often reflect a conscious decision to take exposure to certain risks and corporate hedging is then a distraction.

Large risks however bring with it the risks of financial distress which is of grave concern to other stakeholders like lenders, employees, customers and suppliers who cannot view risks with as much equanimity as well diversified shareholders. Financial distress is very costly and companies would make substantial efforts to reduce the risks of financial distress.

Financial distress is the result of a cash flow shortfall that is severe enough to disrupt the company's strategic plans and investments, or threaten bankruptcy. The function of risk management is thus to reduce the chance of such cash flow shortfalls. Cash flow hedges are, therefore, more important than value hedges.

If risk management is viewed as an attempt to mitigate financial distress, then it is seen to be a substitute for equity capital. Capital structure is a general purpose risk management tool that provides a defence against all forms of risks - even risks that have not been identified and measured. Conversely, risk management is a necessary prerequisite for sustaining high levels of debt without creating an excessive risk of distress. Financial institutions (discussed in the next chapter) are good examples of enterprises that need highly sophisticated

risk management systems to support high levels of leverage. Large companies may allow their treasuries to run a full scale derivative business including speculation, arbitrage, and trading. These treasuries must also be controlled using the same sophisticated risk management systems that are used by banks and financial institutions.

Suggestions for Further Reading

The issues involved in corporate risk management are analysed in:

Doherty, N. A. (1985) Corporate Risk Management: a Financial Exposition, New York, McGraw-Hill.

Froot, K. A. Scharfstein, D.S., and Stein, J.C (1994) "A Framework for Risk Management", *Harvard Business Review*, Nov–Dec 1994, 91–102

Rene M.Stulz (1996) "Rethinking Risk Management", *Bank of America journal of applied corporate finance*, 9(3), 8–24.

The role of derivatives in corporate risk management is explained in:

Stulz, RM (1984) "Optimal Hedging Policies", *Journal of Financial and Quantitative Analysis*, 19, 127-40 Weinberger, David B. (1995) "Using Derivatives: What Senior Managers Must Know", *Harvard Business Review*, Jan–Feb 1995, 33–41.

21.10 | Derivatives and Risk Management

CASES

Case 21-1 ASHANTI GOLDFIELDS COMPANY LIMITED (A)

"Nature endowed us with gold. Our forebears bequeathed us a rich culture based on it. We have created a Pan-African Gold company from our continent's rich natural and human resources." Ashanti Goldfields Company Limited

In early 1999, Ashanti Goldfields Company Limited (AGC) was reviewing its strategies in response to the falling price of gold. The gold price was now dangerously close to the operating cost of AGC's largest mine—the century old Obuasi mine in Ghana. Since the beginning of 1996, the price of gold had fallen steadily and was now approaching levels not seen in the last twenty years. Fortunately, the company's gold hedging strategy had been hugely successful and had insulated it from the worst effects of this price fall (See Exhibits 1,2 and 3 for AGC's financial performance and the impact of the gold hedges). Now as the gold price continued to fall and AGC was on the verge of completing its most valuable gold mine—the Geita mine in Tanzania, it wondered whether this hedging strategy was adequate.

Company Background

Gold has been so intimately associated with the history of Ghana that the country was for long known as the Gold Coast. In the pre-colonial era, the gold of Ghana was the foundation of the prosperity and power of the Ashanti Empire. This empire, which stretched across most of West Africa, was one of the most sophisticated societies of its age. The gold contributed to the demise of the empire as it was a key attraction for colonial powers.

AGC came into being in 1897 when Ghanaian and English entrepreneurs joined hands to mine gold at Obuasi. The company was listed on the London Stock Exchange. Over the next century, as governments came and went, AGC survived as one of the largest economic enterprises in Ghana.

In the late 1960s, AGC was taken over by Lonrho, the UK based pan African conglomerate. Lonrho was founded in 1909 as the London and Rhodesian Mining and Land Company, but its international expansion took place under the flamboyant leadership of 'Tiny' Rowland from the mid 1960s to the late 1980s. Rowland's close personal relationships with many African political leaders (ranging from Mobutu and Ghadaffi to Nelson Mandela), coupled with his willingness to invest in projects that other foreign companies considered too risky helped him create a large conglomerate with diverse business interests throughout sub-Saharan Africa. The Lonrho empire spanned businesses as diverse as mining, hotels and casinos, motor car assembly, construction, food processing and trading spread across fifty countries. At the same time, his unorthodox management practices provoked British Prime Minister, Edward Heath, to describe him as the "unpleasant and unacceptable face of capitalism" in 1972.

In 1968, as Ghana turned to the socialist model of development, AGC was partly nationalised with Lonrho remaining a large minority shareholder. Later, as the country embarked on a World Bank supported

Prepared by Professors Samir K. Barua and Jayanth R. Varma. Written 1999. Revised 2000. Teaching material of the Indian Institute of Management, Ahmedabad, is prepared as a basis for class discussion. Cases are not designed to present illustrations of either correct or incorrect handling of administrative problems. Copyright © 2000 by the Indian Institute of Management, Ahmedabad.

programme of structural reforms, it was gradually semi-privatised in the mid 1990s. A flotation in 1994 and additional capital issues in 1995 reduced the government's stake below 30%. Subsequently, the government sold some of its shareholding to international institutional investors. In 1999, the government of Ghana still held 20% of the company's equity capital. In addition, it held a "Golden Share" (a non-voting special rights redeemable preference share of no par value) which gave the government a veto power over certain important matters like the disposal of a material part of AGC's assets. AGC accounted for 60% of the total Ghana stock market capitalisation, and many foreign investors were attracted

Meanwhile the Lonrho conglomerate also went through a major restructuring after the departure of Tiny Rowland¹ in the mid 1990s. Lonrho transferred most of its non-mining businesses to a separate company, Lonrho Africa, and then changed its name to Lonmin reflecting its new status as a purely mining company. In 1999, Lonmin was the biggest shareholder in AGC with a 32% stake.

The rest of the shareholding of AGC was spread world-wide through listing on six stock exchanges—Accra (Ghana), Harare (Zimbabwe), London (United Kingdom), New York (United States), Sydney (Australia) and Toronto (Canada).

Gold Prices

For millennia, gold has been the store of value par excellence and the benchmark against which other currencies were evaluated. Even after the Second World War, when all major currencies pegged themselves to the US dollar, the dollar itself maintained its link to gold. It was only in the early 1970s that this link was also broken and all currencies became genuine paper currencies. For another decade however, gold remained the ultimate safe haven to which investors fled during periods of inflation as well as during economic and political crisis.

In early 1980, the gold price hit its all time high of \$850/oz.², more than twice the level prevailing in the early 1970s when currencies became freely floating (See Exhibit 4 for a chart of gold price movements). Over the next two years, gold fell sharply to below \$300 before the Latin American debt crisis saw gold recover to over \$500 in early 1983. Then the cycle was repeated again with gold falling below \$300 in 1985 and surging to \$500 in the aftermath of the stock market crash of 1987.

By the late 1980s however, a different pattern seemed to be emerging. Gold seemed to be losing its safe haven status and becoming just another commodity. Gold did rally in the build-up to the Gulf war, but the price rise was muted. In the 1990s, gold seemed to be unaffected by economic and political turmoil and fluctuated within a narrow band ranging from \$300 to \$400. The Mexican peso crisis caused hardly a ripple in the market. Even more shocking to gold bulls was the 1997 Asian crisis and the Russian meltdown of 1998. As the world experienced the worst global financial crisis since the Second World War, the price of gold fell steadily!

The US dollar was seen as having taken over the safe haven function traditionally performed by gold. It was said half-jokingly that of the three gifts (frankincense, myrrh and gold) that the three wise

¹ Tiny Rowland sold most of his stake in Lonrho to a German entrepreneur, Dieter Bock, in 1994 and was ousted from the board shortly thereafter. He was bitterly critical of his successors and seriously contemplated a bid for either Lonmin or Lonrho Africa in early 1998. Tiny Rowland died in July 1998 at the age of 80.

² The troy ounce is 31.104 grams and is the unit in terms of which gold prices are quoted internationally. A metric tonne is 32,150 troy ounces. Gold prices tend to be set in US dollars. When the dollar fluctuates against other major currencies like the Japanese Yen and the Deutsche Mark/Euro, the price of gold in dollars tends to be stable and it is the price of gold in these other currencies that tends to fluctuate.

21.12 | Derivatives and Risk Management

men brought to the infant Christ, gold had remained a gift worthy of a king for two millennia, but the time has now come for a more thoughtful wise man to substitute gold with US Treasury bonds³. For a contrary view of gold prices and prospects, see Exhibit 5.

Demand and Supply

The demand for gold for use in jewellery, dentistry, electronics and other industries is relatively stable and predictable. Similarly, the annual production of gold from the gold mines is also relatively stable and predictable. For several years now, this production has fallen short of the jewellery and industrial demand for gold. Exhibits 6 and 7 provide information about gold production and gold demand.

What makes gold markets so difficult to predict is the demand or supply of gold by investors, speculators and central banks⁴. For example, the central banks and multilateral financial institutions sit on a gold pile of well over a billion ounces of gold, which is several times the annual production of gold in the world (see Exhibit 8). These reserves were built up gradually over the 19th and 20th centuries, but after 1965, these central banks stopped adding to their gold reserves. Between 1965 and 1990, official gold reserves declined by about 7%; in the next eight years, reserves fell by another 6%. But, in the late 1990s, several large central banks were seriously considering the sale of a substantial fraction of their gold reserves.

Gold Hedging Mechanisms

It was in this background that from the 1980s onwards, a large number of gold producers began to consider ways of protecting themselves against a fall in the price of gold. Within a decade, it had become common for most gold mines to use a variety of mechanisms to hedge their exposure to the gold price.

Forward Contracts

The simplest mechanism is for the producer to enter into contracts for sale of gold at specified future dates at specified prices. These forward contracts ensure that the price that would be realised for output in the future is fixed right now and the producer is not affected by fluctuations in the gold price at all.

The other party to a forward contract would be contracting to buy gold at specified future dates at specified prices, but very few ultimate consumers of gold are actually interested in entering into such contracts. So how does a gold dealer offer forward contracts to the producers? The key to this is the gold lease.

The central banks, which hold huge amounts of gold reserves, are obviously keen to earn some return on them. So they are prepared to lend this gold to the big international banks (which are also the largest gold dealers) at a small rate of interest (the lease rate), typically 1-2%. The gold dealer can thus buy gold forward from a producer, borrow gold from a central bank at a lease rate of say 2%, sell the gold spot and invest the proceeds in a safe financial asset at say 6%. It would earn a 4% spread between the lease rate and the dollar interest rate and would therefore be willing to quote a forward price correspondingly higher than the spot price. This excess of the forward price over the spot price is known as contango. The dealer is not exposed to gold price risk at all: the gold it receives under the forward contract would be used to repay the gold loan.

³ Kenneth Gooding, *Financial Times*, December 13, 1997.

⁴ No claim is made that these are three distinct categories. A central bank may be thought of as an investor or even as a speculator!

Spot Deferred Contracts (SDC)

A popular variation of the forward contract is the spot deferred contract (SDC) in which the producer obtains a flexible delivery date. SDCs have several variants, but the concept can be illustrated with a simple example, an SDC with annual rollovers and a final delivery date ten years from today. The delivery price for delivery at the end of the first year is fixed and is equal to the price for a one-year forward contract. However, the producer can choose not to deliver at the end of the first year and roll over to the next year. The delivery price for the second year would be set equal to the original delivery price plus the one-year contango prevailing at the end of the first year. In general, at each rollover date, the delivery price for the next period is set equal to the delivery price for the current period plus the contango prevailing at the rollover date.

In this simple variant of the SDC, rollover is functionally equivalent to closing out the contract and initiating a new one-year forward contract. However, the SDC may have important tax and accounting advantages over the conceptually equivalent closing out strategy. Moreover, more complex forms of the SDC may make the next delivery price depend on the contango (and therefore on the gold lease rate) in more complex ways.

Buying Put Options

A third way for producers to hedge gold price risk is to buy put options on gold. A put option gi,ves the holder the right, *but not the obligation*, to sell gold at the expiry date at the pre-specified "strike price". This means that if the gold price falls, the producer could take advantage of the put option and sell the gold at the strike price. But if the gold price rises, it can discard the put option and sell the gold in the spot market to take advantage of the higher price there. The put option thus provides a floor price for gold, while allowing the producer to benefit fully from any upside in the gold price. In other words the put option is a one way bet on the gold price. There is no free lunch here, because the producer would have to pay a premium to buy the put option. The higher the strike price, the higher this premium would be.

Selling Call Options

Many producers hate to pay a premium for these put options and would like to recover the option premium by some other means. Many of them, therefore, also sell call options to recover part of the premium. A call option is the mirror image of a put option; it gives the holder the right, *but not the obligation*, to buy gold at the expiry date at the pre-specified "strike price". As the seller or writer of the call option, the producer has now committed to sell gold at the strike price. Clearly, the holder would exercise the call option only when the gold price is high and thus the producer would be called upon to sell gold under the call option only when the spot price is higher than the strike price. Selling gold calls, therefore, sets a ceiling on the gold price and leaves the producer fully exposed to any fall in the gold price. The producer receives a premium for this.

The most common practice is to combine selling call options with buying put options. The effect is to set both a floor and a ceiling on the gold price. Within this range, the producer is exposed to the fluctuations in the spot price, but both the upside and the downside is limited. This strategy is known as a range-forward or as a vertical spread.

Many exotic forms of call options are also available. In so-called "turbo" options, or "escalating ounce" calls, the quantity of gold that the holder can buy at the fixed strike price rises when the gold price rises. These options are more valuable and therefore sell at a higher premium.

21.14 | Derivatives and Risk Management

The standard accounting policy for sale of call options is to treat the premium earned as the income of the year in which the option expires. It would therefore boost the revenues and profits of that year. The cash inflow however comes in when the options are sold and is available for use in the business.

Gold Loans

As pointed out earlier, gold dealers offering forward contracts hedge their risks by using gold loans. Some gold producers have also used the gold loan as a way of simultaneously hedging gold price risk and raising finances for developing their gold mines. For example, in 1990, a large North American gold company raised a gold loan of approximately a million ounces of gold (about \$400 million) to finance one of its massive gold mines. The loan was to be repaid by monthly deliveries of gold over several years and was collateralised by the assets of the mine. Needless to say, such large gold loans were available only to the most creditworthy companies.

Early close-outs

A less extreme way of using the gold hedge as a financing mechanism arises when the gold price has fallen significantly from the price prevailing when the hedge was put in place. For example, suppose a producer entered into a two-year forward contract when the spot gold price was \$400; at a 2% annual contango, the forward price would have been \$416. A year later, suppose the gold price has fallen to \$300; at a 2% contango, the forward price is now \$306. The gold dealer would gladly pay a little over \$100 (\$110 less interest for one year) to close out the forward contract. The producer may simultaneously enter into a new one-year forward contract at \$306 to cover its risk. In other words, the producer receives a part of the gold price in advance. It could use this money to finance its capital expenditure for which it would otherwise have to borrow.

Again, standard accounting policy for close-outs would treat the gain as the income of the year in which the forward contract would have matured in the normal course. It would therefore boost the revenues and profits of that year.

Agc's Financing and Hedging Strategy

AGC pursued an ambitious strategy of becoming a pan African gold company. Acquisitions in the mid 1990s included the following:

- Cluff Resources, acquired for \$134 million, owned gold mines in Zimbabwe,
- International Gold, acquired for \$100 million, was a Canadian company that owned the Bibiani gold mine in Ghana
- Golden Shamrock, acquired for \$100 million, owned the Iduaprien gold mine in Ghana as well as the Siguiri mine in Guinea

Many of these acquisitions were financed with debt including a \$250 million convertible issue listed in New York in 1996 and a \$270 million revolving credit facility negotiated in 1999. AGC's debt level was therefore fairly high by global gold industry standards. Moreover, it had large capital expenditure plans to modernise its existing mines and to bring new mines to production. The Geita mine in Tanzania, for example, would need an investment of \$100 million in 1999.

In order to pursue this strategy of aggressive debt financed growth, AGC had to insulate itself from gold price risk. Otherwise, a downturn in the price of gold would not only make some of its projects unviable, but also threaten its ability to service its debt.

From the mid-1990s onwards, AGC, therefore, progressively increased its gold hedging programme, and in March 1999, it covered 8.7 million ounces of forward contracts, SDCs and put options as well as 3.5 million ounces of call options written. The hedging programme thus covered a significant fraction of AGC's reserves of 23 million ounces. Exhibit 9 provides a break-up of the hedge book by maturity and by type of instrument. For comparison, Exhibit 10 provides a summary of the hedge book as of December 1997; it is evident that in a little over a year, the hedge book had grown significantly in size. Moreover, the hedging programme now covered a much longer maturity, going out 15 years instead of 10, and the average maturity (excluding call options) had risen from 4.1 years to 6.7 years.

When gold dealers entered into long term gold hedges with AGC, they had to worry about AGC's creditworthiness. What would happen for example if at the delivery date AGC were unable to honour its commitment because of insolvency or technical problems in its mines? In the case of gold companies with strong balance sheets, gold dealers may have sufficient confidence in the company's credit-worthiness not to ask for additional safeguards. AGC was hardly in this category. It had financed many of its acquisitions with debt, it had large capital expenditure plans for which the financing was not yet tied up, and its Ghana location was perceived to add to the "political risk". Consequently, AGC was regarded in the financial markets as a poor credit risk.

It therefore had a margining agreement with the 17 banks (counterparts) with which it did its hedge trades. Under the terms of this agreement, the entire hedge book would be valued every day using prevailing market prices to determine its "replacement cost". This was the replacement cost from the point of view of the counterparties, i.e., how much would it cost them to create the same positions in the market that they had against AGC. If gold prices were low, this value would be negative implying that the counterparties would be happy to simply see the whole hedge book disappear. In fact, it is AGC that stands to benefit from the hedge book, and the banks do not have to worry that AGC may default.

But if the replacement cost turns out to be positive, it would imply that the hedge book is valuable to the counterparties who stood to benefit from enforcing the hedge contracts. If AGC cannot honour its commitments, then the counterparties would have to spend money buying similar contracts in the market place. Viewed from the other side, AGC would be better off selling gold at spot rather than honouring the hedges that it has entered into. The counterparties would like safeguards to ensure that AGC does not default.

The agreement therefore provided that if the hedge book showed a replacement cost exceeding \$300 million, AGC would have to make a cash deposit (margin) with the counterparties equal to the excess over \$300 million. Put differently, the counterparties were saying that they were willing to extend an implicit credit of no more than \$300 million to AGC. As of March 1999, the replacement cost of the hedge book was a negative \$215 million—in other words, for AGC, the hedge book was not a liability, but an asset (worth \$215 million⁵). Exhibit 11 provides an analyst's view of Ashanti Goldfields and its hedging strategy.

⁵ The further fall in the price of gold since March 1999 would have caused this value to increase.

21.16 | Derivatives and Risk Management

The Scenario In Mid-1999

Two sets of factors seemed to necessitate a review of the hedge programme in mid-1999. First gold prices seemed to be in secular decline. As the global economy recovered quickly from the Asian and Russian financial crises in late 1998 and early 1999, gold prices showed no signs of improvement. Central banks in several countries including Switzerland were actively considering the liquidation of a large part of their gold reserves. For several years, gold had been the poorest performing asset in central bank reserves and these central banks now wondered why they should not sell the gold and switch to higher yielding assets like US treasury bonds. The IMF also considered selling part of its gold to finance debt relief for Heavily Indebted Poor Countries (the HIPC initiative). The gold markets were highly apprehensive of the impact of large central bank and IMF selling. A century ago, it was Bismarck's decision to sell the German silver reserves and switch entirely to gold that marked the demonetisation of silver and its transformation into just another commodity. Could the same happen to gold now?

In May 1999, the Bank of England dropped a bombshell on the gold markets by announcing its intention to sell 415 tonnes out of its 715 tonne gold reserves in a series of auctions of 25 tonnes each starting in July. In response to this announcement, the gold price fell by 10% to a 20-year low of \$258. AGC's share price was also falling, from around \$8 in April to around \$6 in July.

The second set of worries related to AGC's upcoming Geita mine in Tanzania. Most observers thought that the Geita mine was AGC's most valuable asset. It was expected to come on stream in mid-2000 and quickly reach a production level of well over half of the Obuasi mine. In fact, further discoveries of gold in this mine could easily make it as large as Obuasi itself. More importantly, at Geita the cash costs of mining gold was only about \$180 as compared to about \$225 at Obuasi. Arranging the financing for Geita was of critical importance to AGC as the mine would need a capital expenditure of nearly \$100 million in 1999. It hardly seemed prudent to invest a large sum in this mine without hedging the gold price risk.

In this situation, AGC needed to review its financial and hedging strategy.

Required

- 1. Has AGC benefited from its hedging programme? How would it have performed if it had not hedged?
- 2. Is the hedging strategy consistent with the financing strategy and are these strategies consistent with the overall business strategy of AGC?
- 3. Is the hedge book appropriate in terms of its size? Should AGC hedge more or less than it does now? Is the significant increase in the hedge book between December 1997 and March 1999 appropriate?
- 4. Is the hedge book appropriate in terms of its maturity profile? Do you agree with AGC's decision to use hedges going out as long as 15 years?
- 5. Is the hedge book appropriate in terms of its choice of hedge instruments? Do you agree with the increased reliance on options as against forwards? Do you agree with the mix of put and call options?



Exhibit 1

21.18 | Derivatives and Risk Management

Exhibit 2 Group Profit and Loss Account								
				(US \$ million)			
	Qtr ending	Qtr ending	Yr ending	Yr ending	15 mon to			
	31-Mar-99	31-Mar-98	31-Dec-98	31-Dec-97	31-Dec-96			
Gold production ('000 ounces	5)							
Total	402.151	331.414	1547.61	1169.156	1028.597			
Attributable	385.093	319.732	1484.739	1139.832	1024.803			
Turnover	153.70	133.70	600.30	531.30	564.00			
Operating cost	93.70	87.90	379.60	349.30	363.00			
Exceptional operating costs			9.20	4.70				
Royalties	3.20	2.70	12.60	10.60	13.60			
Depreciation	26.40	23.50	108.60	87.80	84.90			
Total costs	123.30	114.10	510.00	452.40	461.50			
Operating profit	30.40	19.60	90.30	78.90	102.50			
Gain on sale of investments				0.20				
Earnings from associated und	lertakings			1.10				
Loss on sale of business			24.00					
Profit before interest	30.40	19.60	66.30	78.90	103.80			
Net interest payable	8.40	4.60	23.20	23.20	22.30			
Profit before taxation	22.00	15.00	43.10	55.70	81.50			
Taxation				2.40	0.10			
Profit after tax	22.00	15.00	43.10	53.30	81.40			
Minority interest	0.50		2.40					
Shareholders' Profit	21.50	15.00	40.70	53.70	81.40			
Dividends			10.90	21.80	39.00			
Retained profit	21.50	15.00	29.80	31.90	42.40			
EPS(US\$)	0.20	0.14	0.37	0.50	0.88			
Production Costs per ounce								
Cash operating costs	211	237	218	254	252			
Royalties	8	8	8	9	10			
Depreciation & amortisation	64	70	68	73	70			
Total costs	283	315	294	336	332			

Attributable production is the total (consolidated) production less the production of subsidiaries attributable to minority interest.

Exhibit 3 Group Balance Sheet							
					(US \$ million)		
	As at						
	31-Mar-99	31-Mar-98	31-Dec-98	31-Dec-97	31-Dec-96		
Fixed Assets:							
Property, Plant & Equipment	1,112.60	1,118.20	1,099.40	1,099.00	960.40		
Intangible assets	111.50		111.50				
Investments	2.20	2.20	2.20	2.20	2.20		
Total	1,226.30	1,120.40	1,213.10	1,101.20	962.60		
Current Assets:							
Stocks	100.70	99.10	99.00	95.70	104.80		
Debtors	71.80	49.90	64.30	54.90	71.20		
Cash	130.20	134.70	112.90	107.50	60.30		
Business held for resale					31.00		
Total	302.70	283.70	276.20	258.10	267.30		
Creditors: due within one yea	r						
Creditors	200.60	198.60	206.40	179.00	151.60		
Borrowings	20.50	28.00	76.00	27.30	23.60		
Total	221.10	226.60	282.40	206.30	175.20		
Net current assets	81.60	57.10	-6.20	51.80	92.10		
Total assets	1,307.90	1,177.50	1,206.90	1,153.00	1,054.70		
Creditors: due after one year							
Creditors	183.30	123.30	217.90	122.90	96.80		
Borrowings	505.70	500.60	414.30	491.50	393.60		
Provisions	21.50	17.50	21.50	17.50	5.20		
Total	710.50	641.40	653.70	631.90	495.60		
Capital and reserves							
Stated capital	540.80	515.60	518.60	515.60	489.50		
Profit and loss account	52.30	19.60	30.80	4.60	68.30		
Shareholder's funds	593.10	535.20	549.40	520.20	557.80		
Equity minority interests	4.30	0.90	3.80	0.90	1.30		
Total	597.40	536.10	553.20	521.10	559.10		
Total Liabilities	1,307.90	1,177.50	1,206.90	1,153.00	1,054.70		

The McGraw·Hill Companies

21.20 | Derivatives and Risk Management



Source: Kitco. Com

Exhibit 4B Gold Prices in January-June 1999

J٤	nuary	Fe	bruary	March		April		May		June	
4	287.00	1	286.90	1	286.70	1	280.55	4	285.30	1	266.85
5	286.40	2	288.55	2	287.70	6	278.50	5	286.00	2	267.15
6	287.65	3	286.85	3	287.90	7	280.15	6	287.95	3	264.75
7	289.95	4	287.25	6	289.50	8	281.00	7	282.40	4	265.30
8	290.90	5	289.40	7	292.95	9	280.60	10	278.00	7	265.60
11	291.15	8	288.85	8	291.80	12	283.20	11	279.45	8	262.35
12	289.80	9	287.00	9	294.00	13	283.45	12	277.55	9	261.30
13	285.45	10	288.20	10	292.50	14	283.40	13	277.85	10	258.60.
14	285.65	11	287.30	13	289.70	15	283.65	14	276.00	11	259.70
15	285.90	12	287.00	14	283.00	16	284.20	17	276.10	14	260.30
18	286.75	15	289.20	15	284.70	19	284.20	18	275.05	15	260.00
19	286.50	16	285.30	16	283.70	20	283.35	19	272.50	16	258.70
20	286.00	17	285.75	17	283.70	21	283.25	20	274.00	17	258.15
21	287.75	18	285.70	20	284.30	22	283.40	21	273.30	18	258.95
22	287.00	19	286.35	21	284.10	23	283.30	24	272.20	21	259.30
25	287.70	22	287.20	22	284.25	26	281.50	25	271.70	22	258.55
26	287.10	23	286.70	23	283.40	27	281.05	26	269.50	23	259.15
27	283.80	24	287.40	24	279.80	28	283.00	27	268.95	24	260.15
28	283.55	25	286.45	27	280.15	29	283.00			25	260.70
29	285.40	26	287.05	28	279.60	30	286.60			28	260.30
				29	279.45					29	262.00
										30	261.00

Exhibit 5

Is Gold More Than Just A Commodity?

While it was fashionable in the financial press to dismiss gold as just another commodity, a contrary view could be found at web sites like <u>www.the-privateer.com</u> or <u>www.gold-eagle.com</u>. A strong distrust of the state and a conspiracy theory of international finance permeate much of this writing. Yet, it is useful to understand their arguments. This note provides a brief summary of the thinking of two gold analysts from this stream.

Bill Buckler

Mr. Buckler maintains the web site <u>www.the-privateer.com</u> and has published a newsletter called the privateer since 1984. His extensive analysis of gold includes the following points:

- "Gold is money. It was rare and prized long before the concept of 'money' was ever discovered. It has been money for nearly 3,000 years; it is money now, and it will continue to be money ... Modern 'money', in all its infinite variety, is nothing more than an IOU, a promise to pay. It is not money for the very simple reason that payment cannot be rendered with a promise. You cannot eat, drink, live in or wear a 'promise'."
- The central banks and governments of the world have fought a losing war against gold since the 1960s.
 - During 1962 to, 68, they tried to cap the gold price at \$35 an ounce. "The First Gold War was lost in April 1968. In August 1971, the last link between Gold and the Dollar was repudiated. In March 1973, the world's currencies "floated". By the end of 1974, Gold had soared to \$US 195 an ounce."
 - During 1975 to, 78 the US and the IMF sold gold in large quantities to depress the gold price. "More important by far, it 'burned' large numbers of small individual investors."
 - When gold prices soared in 1979, reaching a peak of \$850 in January 1980, the government acted again to burn the investors. "The public was lured out of Gold as a substitute for paper by a huge rise in interest rates. The U.S. Prime rate hit 20% in April 1980 and stayed there (with a brief dive in mid-1980) until the end of 1981. There was a rush out of Gold and back to Dollars."
 - When interest rates fell, gold started rising again. The governments then engineered a series of stock market booms (global boom, 1982,87; Japanese boom, 1988,90; US boom, 1994today) to funnel money from gold into stocks.
- Around 1987, governments shifted their strategy and began to wage a hidden war against gold. "The price of Gold cannot be held down by selling the physical metal. The decade between 1970 and 1980 proved that conclusively. In the years since the 1987 crash—when the \$US 400 'glass ceiling' on Gold has been put and kept in place, Central Banks have continued to sell Gold, but only in emergencies. The real mechanism for holding down the price has been different."
 - Governments learned in the 1960s and 1970s that it was impossible to meet an increased demand for Gold with physical Gold. They needed a paper substitute. Gold "derivatives" provided that substitute. With more tradable alternatives to physical Gold, it became far easier to control the Gold price."

21.22 | Derivatives and Risk Management

- Forward sales by gold producers and gold leasing by the central banks was the key mechanism in this process. "The miners' Gold is still in the ground. The Central Bank sometimes lends Gold, or it lends a claim to Gold. These are what our bullion dealer sells. And since most demand for Gold is not a demand for the physical metal but a demand for paper (forward, future, etc) claims to the metal, this mechanism can meet the demand without an undue strain upon the available supply of the physical metal, and the upward pressure on the price of Gold that would cause."
- With these techniques in place, governments were able to confine gold in a narrow trading range in the face of the Mexican crisis, the Japanese deflation, the violent swings in the dollar's exchange rates against the yen and the mark. The purpose of this trading range was:
 - "To encourage the attitude that Gold had lost its former use as a 'hedge' against price inflation and political and economic crises."
 - ➤ "To encourage the attitude that (price) inflation was 'dead'."
- In early May 1999, the shares of gold companies started rising sharply prompting a sharp change in government strategy. "A big rise in Gold stocks is always the precursor to a big rise in Gold itself ... Gold stocks were in imminent danger of becoming the next "hot sector" on the stock market. The sudden quantum leap in Aussie gold stocks (up 15% over May 6-7) made this danger acute. Something had to be done before this newfound interest in Gold stocks began to spill over into an interest in Gold would not suffice. So the Bank of England announced actual Gold Auctions! ... the Gold auctions is not something that the B of E has cooked up over the past week or month. This is a contingency plan that has long been formulated, and has waited for the 'right time' to be brought forward. After almost 20 years of 'covert' Gold price manipulation (forward selling, lending, leasing, short selling etc.), what is happening here is a reversion to the OVERT selling of Gold (not paper claims to Gold).
- Mr. Buckler's analysis can be summed out very simply "In any discussion of the future of Gold, or of the price of Gold, the first thing that must be realized is that Gold is a political metal. In the true meaning of the word, its price is 'governed'. This is true for the very simple reason that Gold in its historical role as a currency is fundamentally incompatible with the modern worldwide financial system." In this view, therefore, the fall in gold prices during 1997–99 did not reflect its fundamentals at all, but was merely the result of government manipulation.

Richard M. Pomboy

Mr. Pomboy's views were expressed in a speech at the Grant's Fall Investment Conference in October 1997 and are available at <u>www.gold-eagle.com</u>. His analysis contains the following points:

- Mr. Pomboy begins with an observation similar to that of Mr. Buckler: "Gold, the traditional store of value which people have used as currency since civilization began and to buy their way out of danger, is now obsolete, demonetized and a relic. It has been replaced by paper, and if you are such a heretic as to believe that paper may have a slight risk to it, all risk can be eliminated through derivatives, which are, of course, more paper. The paper asset mania requires that the principal alternative, gold, be thoroughly discredited."
- "The basic supply/demand equation is extremely bullish with a gap between supply and demand of around 800 or 900 tonnes per year last year and much larger this year ... Current gold price will lead to production cutbacks. At \$320/oz., 50% of the western world's gold production is unprofitable on a full cost basis. At \$320/oz., 25% of the western world's gold production is unprofitable on a cash cost basis."

- Producer forward sales, speculator shorts and some central bank sales fill the deficit currently. "With producers heavily hedged and shorts having large positions which have already pushed up lease rates, that is the cost of borrowing the gold they sell short, it is likely that neither of these groups will have a significant impact on filling the deficit going forward. That means that gross central bank sales of probably 1,500 tonnes or more are needed to keep gold at the current low level."
- There are significant risks in the gold loan mechanism. "The loaned gold ends up being sold into the marketplace by the dealer on behalf of short sellers, producers, etc. and much of it is fabricated into jewelry; The loaned gold has, in large part, permanently disappeared. If one central bank wants its gold back, the dealer can borrow gold from another central bank. But if many or most of the central banks want their gold back, for whatever reason, then lease rates would skyrocket, there would probably be a default and gold would go into backwardation where the spot price was higher than the future. This would wreak havoc with producers who sold forward, short sellers and dealers, all of whom could suffer huge losses. There is increasing evidence that the amount of gold on loan is much greater than generally thought and thus the risks in the market are increasing. As a result of all this borrowed gold the risks in the gold market are very real. No one anticipated the palladium market going into backwardation as the Russians were expected to continuously supply metal to the market. This key assumption failed which could happen in the gold market where the key assumptions are that central banks will continue to lend their gold, and in such amounts that lending rates remain at low enough levels to keep the market in contango. If lending central banks withdrew from this activity or if physical demand for gold is substantially increased by other central banks or investors becoming aggressive buyers then lease rates would become so high that the market would go into backwardation and the spot price would rise dramatically as in the case of palladium."
- The current state of the market is that "Central banks have sold much of what they want to sell. Producers are heavily hedged and have huge gains. Shorts are in record short position. Sentiment at worst levels." Any catalyst for investment demand (EMU problems, problems in paper markets, commodity resurgence, significant Asian purchases) could lead to a quick reversal of all of this as central banks start buying, producer reverse their hedges, and speculators cover their short positions.

World Gold Froduction Since 1040								
(millions of troy ounces)								
Country/Region	1840- 1850	1851- 1875	1876- 1900	1901- 1925	1926- 1950	1951- 1975	1976- 1996	
Australia	0.0	51.7	39.7	54.5	22.9	22.4	82.9	
USSR	7.5	22.6	29.1	24.6	73.2	134.9	133.3	
USA	5.2	58.4	51.0	93.9	66.8	40.4	109.9	
Canada	0.0	2.2	4.8	21.2	85.1	87.8	72.8	
South Africa	0.0	0.0	2.7	178.4	292.3	582.0	430.0	
Other Countries	5.2	19.1	37.0	104.9	159.9	117.9	297.7	
World Total	17.9	154.0	182.3	477.5	700.2	985.4	1126.6	

Exhibit 6A World Gold Production since 1840

Source: Gold Institute

21.24 | Derivatives and Risk Management

Exhibit 6B Major Western World Gold Producers							
		(n	illions of troy ounces)				
Country/Region	1996	1997	1998				
South Africa	16.0	15.8	15.0				
USA	10.2	11.4	11.7				
Australia	9.3	10.1	10.0				
Canada	5.3	5.4	5.3				
World Total		80.0	82.0				
Source: Gold Institute							

Exhibit 7						
	W	orld Gold	Demand			
					(millions	of Ounces)
	1993	1994	1995	1996	1997	1998
India	13.0	13.3	15.3	16.3	23.7	26.2
Pakistan	1.2	1.2	1.4	1.7	2.6	3.2
North Asia	13.5	13.9	13.7	12.0	13.1	10.1
-China	7.2	7.2	7.2	6.8	6.9	6.2
-Taiwan	5.2	5.2	0.5	4.0	4.6	2.9
-Hong Kong	1.2	1.5	1.4	1.3	1.6	1.0
SE Asia & S Korea	11.0	13.2	14.3	14.6	10.2	-3.6
-Thailand	3.1	4.0	3.7	3.4	0.5	0.6
-Singapore	0.7	0.8	0.8	0.6	0.7	0.5
-S Korea	2.9	3.4	3.9	4.0	3.7	-5.2
-Malaysia	0.7	0.8	1.0	1.1	1.0	0.5
-Indonesia	2.6	3.1	3.8	4.1	3.0	-1.3
-Vietnam	1.0	1.1	1.2	1.3	1.4	1.4
Saudi Arabia	6.4	5.6	6.2	5.9	6.4	6.7
Egypt	1.8	2.0	2.2	2.4	3.1	3.4
Gulf States	2.6	2.8	3.4	3.8	4.6	4.6
-UAE	1.1	1.1	1.3	1.7	2.3	2.6
-Kuwait	0.6	0.8	1.1	1.1	1.1	1.1
-Bahrain	0.2	0.2	0.3	0.3	0.3	0.3
-Oman	0.5	0.5	0.5	0.5	0.6	0.5
-Qatar	0.2	0.2	0.2	0.2	0.2	0.2
Turkey	5.1	2.6	4.5	4.9	6.5	5:5
Latin America	3.0	3.1	2.7	3.2	3.4	3.8
Brazil	1.7	1.6	1.7	1.9	1.9	2.1
Mexico	1.4	1.5	1.0	1.3	1.6	1.8
Europe	9.5	9.0	9.5	8.8	9.4	9.9
-Italy	4.0	3.7	3.5	3.4	3.6	3.6
-France	1.5	1.4	1.6	1.5	1.6	1.9
-Germany	2.8	2.5	2.8	2.4	2.4	2.3
-UK	1.2	1.3	1.5	1.5	1.9	2.1
USA	9.5	9.7	10.1	10.7	11.6	13.8.
Japan	7.1	6.8	5.5	4.9	3.4	3.5
World Gold Council Market Total	83.8	83.3	92.1	89.4	98.2	87.2

Source: World Gold Council

	TheN	LG FAU	-Hill C	innana	20								
				maduus	3								
			Ň	orld Off	icial Go	Id Rese	erves 19	950-199	98				
											(mi	llions of	Ounces)
	1950	1955	1960	1965	1970	1975	1980	1985	1990	1995	1996	1997	1998
World	7.999	1066.5	1153.9	1232.9	1176.9	1179.1	1152.1	1147.3	1143.7	1109.3	1108.6	1089.7	1075.1
Institutions	44.2	62.1	70.6	39.3	117.2	159.2	196.6	195.8	204.8	200.6	202.0	199.1	109.9
All Countries	955.5	1004.4	1083.3	1193.5	1059.7	1019.8	955.5	951.5	939.0	908.7	906.6	890.6	965.3
Canada	16.6	32.4	25.3	32.9	22.6	22.0	21.0	20.1	14.8	3.4	3.1	3.1	2.5
USA	652.0	621.5	508.7	401.8	316.3	274.7	264.3	262.6	261.9	261.7	261.6	261.6	261.6
Japan	0.2	0.6	7.1	9.4	15.2	21.1	24.2	24.2	24.2	24.2	24.2	24.2	24.2
Switzerland	42.0	45.6	62.4	8.8	78.0	83.2	83.3	83.3	83.3	83.3	83.3	83.3	83.3
Belgium	16.8	26.5	33.4	44.5	42.0	42.2	34.2	34.2	30.2	20.5	15.3	15.3	9.5
France	18.9	26.9	46.9	134.5	100.9	100.9	81.9	81.9	81.9	81.9	81.9	81.9	102.4
Germany	0.0	26.3	84.9	126.0	113.7	117.6	95.2	95.2	95.2	95.2	95.2	95.2	119.0
Italy	7.3	10.1	62.9	68.7	82.5	82.5	66.7	66.7	66.7	66.7	66.7	66.7	83.4
Netherlands	9.0	24.8	41.5	50.2	51.1	54.3	43.9	43.9	43.9	34.8	34.8	27.1	33.8
UK	81.8	57.5	80.0	64.7	38.5	21.0	18.8	19.0	18.9	18.4	18.4	18.4	23.0
South Africa	5.6	6.0	5.1	12.2	19.0	17.7	12.2	4.9	4.1	4.2	3.8	4.0	4.0

Source: World Gold Council

	The McGraw-	Hilla	odmo	inies													
	0												•	(mil	llions	of Oui	ices)
Forward contrac	cts:	1999	2000	2001	2002	003	2004	2005	2006	2007	2008	2009	010	2011	2012	2013	Total
(i) Spot deferred	contracts:																
Amount hedged ((souno 000,)	460	636	657	570	701	713	473	384	264	169	200	200	200	200	150	5977
Average price (U	IS\$/ounce)	387	383	377	371	374	382	389	400	398	405	413	413	413	413	408	388
(ii) Fixed forward	d contracts:																
Amount hedged ((sound 000,)	78	55	95	45	65	25										363
Average price (U	IS\$/ounce)	414	381	421	376	376	376										397
Put options purc	chased:																
Amount hedged ((souno 000,)	209	140	8	120	-132	160	150	150	330	230	230	230	230	270		2407
Average price (U	IS\$/ounce)	326	323	309	333	279	428	424	427	403	388	392	396	400	404		393
Sub-total																	
Amount hedged ((seouno 000,)	747	831	842	735	634	898	623	534	594	399	430	430	430	470	150	8747
Average price (U	IS\$/ounce)	373	373	375	365	394	390	397	407	401	396	402	404	406	408	408	390
Short term Spot	Deferred contracts:																
Amount hedged ((souno 000,)	50															50
Average price (U	IS\$/ounce)	292															292
Net Call options	written:																
Amount hedged ((souno 000,)	59	191	224	268	665	212	203	351	351	361	179	179	100	100	75	3518
Average price (U	IS\$/ounce)	482	482	435	420	388	401	404	392	392	389	377	377	375	375	375	401
(i) Forward contr	acts have been class	ified ac	cordin	g to tl	ne inte	nded d	date o	f matu	rity.				2	6			
(ii) Options have	been allocated to the	e period	ls in w	hich t	hey ex	pire.											
Source: World Go	ld Council]

21.26 | Derivatives and Risk Management
Corporate Risk Management | 21.27

Exhibit 10 Hedging Commitments as at 31 December 1997				
Forward contracts:				
(i) Spot deferred contracts:				
Amount hedged ('000 ounces)	3,085			
Average price (US\$/ounce)	416			
(ii) Fixed forward contracts:				
Amount hedged ('000 ounces)	1,576			
Average price (US\$/ounce)	411			
Put options purchased:				
Amount hedged ('000 ounces)	1,793			
Average price (US\$/ounce)	367			
Sub-total				
Amount hedged (000 's ounces)	6,454			
Average price (US\$/ounce)	401			
Short term Spot Deferred contracts:				
Amount hedged ('000 ounces)	Nil			
Average price (US\$/ounce)	Nil			
Net Call options written:				
Amount hedged ('000 ounces)	1,636			
Average price (US\$/ounce)	429			

Exhibit 11 Excerpts from Research Report

"We like the company's track record of beating set budgets. Although we forecast a reduction in cash flow and profits in FY2000, we see an increase in FY2001 and following years. Increases in production and improvements in operational efficiencies are aided by several new mines coming into production in FY2000."

"Cash flows from these operations will permit a reduction in debt. The strong hedge book provides excellent revenue protection against a low gold price. We therefore recommend the stock as a **long-term buy.**"

PROS	CONS
Excellent growth potential	High debt
Good track record in cost	Several developing mines
reduction	included in portfolio
Strong hedge book	Little earnings gearing towards gold price
International listing	High costs
Sound management	Emerging market company

"Ashanti's strong hedge book provides protection against weakness in the gold price. In 1998 the average price realised was \$98 above the average spot price. Ashanti's policy is to maintain a high level of revenue protection through gold hedging activities."

"At the current price the share offers long-term value, considering the growth opportunities of the company. We are confident that the company can achieve a 5% reduction in costs and a 5% increase in production, relative to the budgeted figures. This is reflected in our NPV evaluation of \$9.61."

⁶ Standard Equities research report on AGC dated May 2, 1999

⁷ \$7.85 per share

21.28 | Derivatives and Risk Management

"We emphasise that the company is likely to see a reduction in cash flow in FY2000. Only after that are increases in cash flow and profit expected. We also point out the relative low sensitivity of NPV to changes in the gold spot price. This is a positive factor during the current suppressed gold price. However, upside potential during gold rallies is also limited."

"Due to the excellent growth opportunities and track record in cost reduction we recommend the stock as a long-term buy."

Case 21-2 ASHANTI GOLDFIELDS COMPANY LIMITED (B)

"All year we have been praying for the price [of gold] to go up and now it has, we are crying. It is difficult even for well-educated people here to get to grips with this problem ": Nuhu Salifu, trustee, Obuasi branch, Ghana Mineworkers Union. "The gold of AGC /Ashanti Goldfields Company] is a symbol of our national sovereignty": Jerry

Rawlings, President of Ghana

Through most of 1999, the biggest worry for gold companies like Ashanti Goldfields Company Limited (AGC) was the declining price of gold (see Case A for information about AGC and its response to falling gold prices). In July 1999, the price of gold dropped to a 20-year low of \$253 per ounce. With no end in sight to the fall in gold prices, and with its cash cost running at approximately \$225 per ounce, AGC had plenty to worry about.

In September, fifteen major central banks whose actual or planned selling of gold had caused the collapse in gold prices, met in Washington and announced an unexpected moratorium on any further gold sales during the next five years (See Exhibit 1). The impact on the gold market was electric: gold price leapt up beyond the \$300 level to a peak of \$325. Exhibit 2, excerpted from the World Gold Council's *Gold Demand Trends*, gives a vivid account of the tumultuous developments in the gold markets in September 1999. Exhibit 3 provides the gold prices during July to October 1999.

Gold producers like AGC should have been jubilant, but they were not. The key to that was the hedges that they had put in place to protect themselves against a fall in the price of gold (see case A for details of AGC's gold hedge book). In June 1999, the hedge book of AGC was worth \$290 million to AGC. In other words, the counterparties would have been willing to pay \$290 million to AGC to get rid of the hedge book completely. With gold prices at \$325, the position was totally reversed: AGC would have had to pay \$570 million to the counterparties to get rid of the hedge book. In the terminology of the gold trade, the hedge book had a replacement cost of \$570 million in the beginning of October as compared to a negative \$290 million in June—a swing of \$860 million.

Three things had caused this sharp change:

1. The rise of about \$55 in the gold price made a big difference to the value of the total hedge of about 10 million ounces.

Prepared by Professor Samir K Barua and Jayanth R. Varma. Written 1999. Revised 2000. Teaching material of the Indian Institute of Management, Ahmedabad, is prepared as a basis for class discussion. Cases are not designed to present illustrations of either correct or incorrect handling of administrative problems. Copyright © 2000 by the Indian Institute of Management, Ahmedabad. Repreduced with Permission

- 2. A sharp rise in the volatility of gold prices suddenly made the call options written by AGC far more valuable than previously.
- 3. About 15% of AGC's hedges were in the form of exotics like turbo calls. In the worst case scenario, AGC's hedge could rise from 10 million ounces to 15 million ounces as the exotics kicked in.

Under the terms of the agreement with the counterparties, AGC had to make cash margin deposits when the replacement cost exceeded \$300 million. This meant an immediate liability of \$270 million at a replacement cost of \$570 million. AGC's immediate response at damage control was to restructure its hedge book and to seek standstill agreements with the banks to gain time to make the margin calls. On September 30, 1999, AGC stated that:

"Ashanti has restructured 80% of its hedge book to remove the sensitivity of the hedge value to further rallies in the gold price. The restructuring was initiated before the recent rally, as part of the Company's contingency planning, and has covered approximately 9 million ounces (out of Ashanti's 11 million ounces of forwards and put options). It has included converting a substantial component of the forward sale positions into synthetic put options. The Company has already, before the rally, eliminated any exposure to floating leasing rates during the rest of 1999, and the first quarter of 2000. Ashanti's hedge book continues to be actively managed and tightly controlled. Management is satisfied that the hedge portfolio is robust in the current gold market and will generate a realised gold price of US\$380 per ounce in 1999 and about US\$360 per ounce next year."

On October 5, 1999, it added:

"... in light of the continued turmoil in the gold market, the Company continues to monitor its hedging position and has over the last few weeks managed the hedge book so as to reduce the sensitivity of the hedge value to rallies in the gold price."

"...Ashanti's hedge book consists of a number of hedging instruments commonly used by gold mining companies and currently represents a net hedge of 10 million ounces. While the estimated replacement cost of the hedge book (based on current market conditions and US\$317.50 per ounce, being the closing price in New York on 4 October 1999) would be approximately US\$450 million, the delivery by the company at the hedged forward prices of the future production of gold from its operations will generate revenues substantially in excess of the current cash operating costs of approximately US\$214 per ounce."

"...The company continues to have a strong operating performance and asset base and its gold operations will continue to benefit from any increase in the price of gold."

The next day, AGC announced that

"[AGC] has entered into a temporary standstill arrangement with its hedging counterparties in order to give the Company time to work out a more permanent arrangement with its counterparties. Under this arrangement, the counterparties have agreed not to make margin calls or to take certain other actions in respect of the hedging contracts".

With the standstill in place, AGC still needed to find ways to improve its liquidity position very quickly. The problem clearly was one of liquidity, and not of solvency: the same rise in gold prices that devastated the hedge book also made the gold reserves in its mines vastly more valuable. Since the gold reserves were about 23 million ounces while the hedge book was only 10 million ounces, clearly the gains outweighed the losses. The problem was that gold under the ground is not a liquid asset while the margin calls on its hedge book were an immediate liability.

21.30 | Derivatives and Risk Management

Clearly, there would be some buyers for some of its mines, especially the attractive Geita mine in Tanzania. There might be buyers also for the company as a whole. The company's largest shareholder, the UK based mining group, Lonmin was ready to launch a bid for the 68% shares of AGC that it did not already own. This was however conditional on the support of the government of Ghana. Other predators and white knights were also evaluating the possibility of bidding either for AGC or for some of its valuable properties. Any such bidder would however have to strike a deal with three stakeholders:

- 1. Lonmin was not only the largest shareholder, but also exercised considerable management control over the company. In fact, even the chief executive of AGC, Sam Jonah was paid by Lonmin, which charged AGC for these and other services.
- 2. The government of Ghana was the regulator and also a 20% shareholder. More importantly, its golden share gave it a veto power over any asset disposals and other major decisions.
- 3. The group of banks (the hedge book counterparties) which had agreed to the standstill would have to agree to any asset sales. Even a change of management would be difficult without their consent given the large financing needs of AGC.

For AGC itself, all these options were quite painful. Selling its prized assets outside Ghana might preserve its independence, but end its hopes of creating a pan-African gold company and reduce AGC to a rump that would have a very dim future. On the other hand, neither the company nor the government of Ghana was likely to look favourably at the prospect of AGC being taken over by Lonmin or any other potential bidder. Yet why should any investor be willing to provide liquidity support to AGC without seeking management control?

Exhibit 1 The Washington Agreement on Gold

Oesterreichische Nationalbank	Banque Nationale de Belgique	Suomen Pankki
Banque de France	Deutsche Bundesbank	Central Bank of Ireland
Banca d'Italia	Banque centrale du Luxembourg	De Nederlandsche Bank
Banco do Portugal	Banco de Espana	Sveriges Riksbank
Schweizerische Nationalbank	Bank of England	European Central Bank
	Press Communiqué	
	26 September 1999	
	Statement on Gold	

In the interest of clarifying their intentions with respect to their gold holdings, the above institutions make the following statement:

- 1. Gold will remain an important element of global monetary reserves.
- 2. The above institutions will not enter the market as sellers, with the exception of already decided sales.
- 3. The gold sales already decided will be achieved through a concerted programme of sales over the next five years. Annual sales will not exceed approximately 400 tonnes and total sales over this period will not exceed 2,000 tonnes.
- 4. The signatories to this agreement have agreed not to expand their gold leasings and their use of gold futures and options over this period.
- 5. This agreement will be reviewed after five years.

Exhibit 2 The Gold Market in September 1999*

On September 2 it was reported that the IMF was considering a plan to sell gold to member central banks at market prices, thus avoiding sales into the open market; the plan was seen as skeletal at this stage however, and prices failed to react. The borrowing market remained very tight, with wide-spread rumours of reduced central bank lending. The one-month lease rate moved above 4% and the 12-month rate above 3.5%, but this likewise had little impact on prices, which held between \$254 and \$257 until the second Bank of England auction on September 21. This auction was much better received than the first. Gold was fixed at \$255.20 that morning and the sale of 25 tonnes was settled at \$255.75, above the prevailing price. The amount applied for was more than eight times that available, and the scaling factor of 58.5% indicated that few bids were made below market prices. This better-than-expected result triggered a short-covering rally that afternoon and gold raced up to the \$260 level. Further good news appeared on September 23 when the IMF confirmed it would essentially be raising funds for its HIPC initiative through the revaluation of part of its reserves ... reaching the \$271 level on the next day.

These developments were but a prelude to gold's next move. The catalyst was the unexpected announcement on Sunday September 26 that the European Central Bank, the 11 members of the Eurosystem, Sweden, Switzerland and the UK had reached agreement limiting central bank sales of gold to 2,000 tonnes over the next five years, or about 400 tonnes per year, and had also agreed not to expand their gold lending to the market over this period....

Quotations burst higher in response, fixing at \$281.70 on the morning of September 27. Another wave of buying on the next morning carried gold up to a fixing of \$288.25, which meant that the market had recovered all of the ground lost since the May announcement of the UK auctions. Options-related buying was also triggered as prices continued to move higher and expiry date loomed, with grantors of call options forced to scramble for cover, giving the impetus for gold to break above \$300 that afternoon.

By this time a serious shortage of physical metal was emerging. Gold lease rates, already very strong, suddenly rocketed with the one-month rate soaring above 10% on September 29, pushing the market into a rare backwardation (spot prices higher than forward prices). Gold itself maintained its momentum, fixing at \$317.25, and the bid/offer spread in the market widened to \$3 in response to the volatility - one-month volatility widened to 35%. After touching a peak for the quarter of \$325, the market faltered as the rally seemed to become overextended. Lease rates eased as producer buy-backs and speculator short-covering released liquidity into the market, and prices fell back on heavy profit-taking, closing a remarkable quarter at the \$300 level.

^{*} This is excerpted from the World Gold Council, Gold Demand Trends, November 1999.

The McGraw·Hill Companies

21.32 | Derivatives and Risk Management

July Au		gust Septe		mber	Oct	October	
1	262.60	2	255 35	1	254 70	1	307 50
2	263.10	3	255.55		254.70		312 75
5	261.80	4	255.85		254.20	5	325 50
6	257.60	5	255.05	6	254.60	6	318.25
7	257.10	6	255.00	7	255.60	7	323.50
8	257.20	9	256.50	8	256.15	8	323.25
9	257.30	10	256.50	9	256.90	11	317.75
12	256.00	11	257.75	10	256.00	12	320.25
13	254.65	12	260.30	13	256.30	13	321.75
13	255.25	13	260.45	13	256.75	14	321.00
15	253.30	16	259.25	15	257.15	15	315.75
16	255.65	17	259.65	16	255.90	18	311.75
19	253.95	18	261.30	17	255.10	19	309.50
20	252.80	19	258.00	20	255.40	20	307.00
21	254.50	20	258.25	21	258.85	21	302.85
22	253.95	23	255.90	22	263.40	22	302.40
23	255.30	24	253.40	23	264.35	25	302.00
26	253.50	25	252.85	24	270.00	26	296.40
27	254.60	26	253.15	27	281.10	27	290.60
28	254.10	27	253.80	28	301.50	28	296.25
29	253.90	31	254.80	29	307.00	29	299.10
30	255.60	30	299.00				

Exhibit 3 Gold Prices During July to October 1999

Chapter **Twenty Two**

Risk Management in Financial Institutions

Financial institutions need a comprehensive measure of enterprise-wide risk that can be communicated to their investors and regulators and also used by top management for monitoring on a day-to-day basis. This chapter discusses the Value at Risk (VaR) which is used for this purpose. This covers the several methods of computing VaR as well as the complications involved when the portfolio contains significant option positions. The methods used to back test the VaR model are discussed. VaR models which must be supplemented by stress tests that measure the vulnerabilities of the portfolios to extreme movements, are also described. Finally, the chapter lists and analyses the regulatory considerations in the design of risk management systems in financial institutions.

22.1 VALUE AT RISK

As explained in the previous chapter, financial institutions are very different from other businesses. Non-financial businesses operate with relatively low levels of debt and expect to meet most of their liquidity needs from operating cash flows. Their risk management, therefore, focuses on monitoring and hedging their cash flows.

Financial businesses operate with much higher levels of debt and their liquidity needs are rarely met out of operating cash flows. Liquidity in a financial business comes largely by selling liquid assets or raising new debt. The gross financial cash flows generated by these means on an annual basis may be several times the total size of the balance sheet. Uninterrupted access to debt finance is, therefore, critical for a financial business to remain liquid. The liquidity of a financial institution depends on its solvency (the value of its assets exceeding the value of its liabilities) at every point in time. From the perspective of these institutions, therefore, it is critical to monitor the value of their assets and liabilities on a daily (and even on an intra-day basis). It is also critical to measure and manage the risks of changes in these values.

While discussing derivative valuation in this book, we have analysed the risk of each derivative. This risk has been measured in terms of the volatility of the underlying and the Greeks (delta, gamma, vega, and theta) which measure the sensitivities of the derivative to various risk factors. This approach is extremely useful and powerful in describing the risks of a derivative or a portfolio of derivatives that depends only on one or two underlyings. Derivative traders depend on them all the time.

A financial institution on the other hand, typically has a derivative book that is exposed to several hundred underlyings. Listing the deltas, gammas and other greeks with respect to each of these underlyings provides a comprehensive description of the risk profile. However, this risk profile has several disadvantages as a measure of enterprise-wide risk:

- Finance theory requires that enterprise-wide risk be measured at the portfolio level and not at the level of individual positions or individual underlyings.
- A risk profile containing several hundred Greeks is too detailed and too enterprise-specific to be communicated to regulators or to investors at large.
- It is too complex even for senior management within the institution to monitor on a regular basis.

22.2 | Derivatives and Risk Management

It is, therefore, necessary to develop a summary measure of risk which addresses all the above shortcomings. In other words, the risk measure must be easy to understand; it must capture the enterprise risk in a few key numbers; it must be comparable across enterprises so that external stakeholders like regulators and investors can interpret the measure easily.

Value at risk is one such summary measure that has become immensely popular over the last couple of decades. Value at risk attempts to measure the loss that may be incurred on the entire portfolio. Since loss is measured in money terms, losses in various different positions and underlyings can all be aggregated to arrive at a loss on the entire portfolio. The probability distribution of this aggregate loss is a complete description of the risk faced by the enterprise at the total portfolio level but it is too complex for internal management or external stakeholders. Value at risk (VaR) summarizes this distribution in terms of a single number. In order to do this, it is first necessary to choose a particular probability level (95%, 99%, or 99.5%). The 95% VaR, for example, can be defined in the following different equivalent ways¹:

- 1. It is the level of capital that is sufficient to absorb the possible loss 95% of the time.
- 2. It is the level of loss that is exceeded only 5% (100% 95% = 5%) of the time.
- 3. It is the worst of the best 95% of possible outcomes.
- 4. It is the best of the worst 5% of possible outcomes.

The 99% VaR is defined similarly with 99% instead of 95% and 1% instead of 5%. In general the VaR at the x% probability level is defined as the level of capital that is sufficient x% of the time; or the level of loss that is exceeded only 100-x% of the time; or the worst of the best x% of possible outcomes; or the best of the worst 100-x% outcomes. Implicit in the definition is a notion of the horizon over which the loss is measured. For example, an annual VaR number is based on possible losses over a one-year time horizon. A daily VaR number is based on possible losses over a one-day time horizon. Thus every VaR number is based on a VaR level (for example 95% or 99%) and a time horizon (for example year or day). In practice, VaR is most usually computed at the 1-day horizon.

The first two interpretations of VaR given above, make VaR an intuitively appealing and interesting summary measure of risk and account for its popularity among regulators, managers, and others. For example, if a bank regulator is unwilling to let more than 1% of the banks fail per year, then it can achieve this goal by requiring all banks to hold capital equal to 99% VaR measured at an annual horizon. If this is done, then any individual bank will have enough capital to absorb the losses in 99 out of 100 years. Any given bank will thus fail once in 100 years. Over a long period of time, the average² annual failure rate in the entire banking system will then be 1%. Banking regulators do use ideas like this to set minimum capital requirements for all banks.

The controlling shareholder of a bank may also use VaR to measure the risk of the bank. If the bank incurs a loss exceeding the capital, then the controlling shareholder must either recapitalize the bank using its other financial resources or it must be willing to lose control of the bank. Both of these are unpleasant choices for the controlling shareholder. VaR measures the probability that the controlling

¹ These different definitions are equivalent if the loss distribution is continuous. If the distribution is discrete or discontinuous, then these definitions may not all be equivalent. In this chapter, the loss distribution is assumed to be continuous to keep the discussion simple.

² In a particularly bad year, a large percentage of banks may fail and in a good year, no bank might fail. While the long run average failure rate will be 1%, it cannot be guaranteed that only 1% of banks will fail in any given year.

shareholder will be confronted with this difficult choice. If the bank has enough capital to cover its 99.5% VaR for example, it means that there is only a 0.5% chance that the bank will lose its entire capital.

The third and fourth interpretations of VaR highlight some potential difficulties with its concept. For example, the fourth definition says that 95% VaR is essentially the best of the worst 5% of outcomes. This immediately appears unsatisfactory)—why is not the worst of the worst 5% of outcomes or at least the average of the worst 5% of outcomes?

It is easy to argue that worst of the worst 5% is not a meaningful measure of risk because the worst outcome may be unbounded. For example, a financial institution that has sold a futures contract or sold a call option on a stock index faces potentially unlimited losses on the position. There is no theoretical limit to how high the stock index can rise during the life of the contract and the potential losses are, therefore, unbounded. The worst possible outcome is thus– $\infty =$ (minus infinity). This is a meaningless measure of risk for most practical purposes³.

The average of the worst 5% of possible outcomes is however a well-defined and meaningful measure of risk. In the risk literature, this is referred to as expected shortfall (ES), conditional VaR (CVaR), or Tail Conditional Expectation (TCE)⁴. From a theoretical point of view, the average of the worst (ES) is a better measure of risk than the best of the worst (VaR). Consider for example two investment firms that both have a one-day VaR of Rs 10 million at the 99% level. This means that on 99% of the days, the loss will be less than Rs 10 million. The ES measure asks the question as to what happens on the 1% of days when the loss exceeds Rs 10 million. It is possible that in one case, the loss may range from Rs 10 million to Rs 15 million with an average of Rs 15 million. Clearly, the second company is a lot riskier than the first though the both companies have the same VaR. The ES measure (Rs 15 million as compared to Rs 12 million) reveals this picture very well.

Despite its shortcomings, VaR is very popular among financial institutions and their regulators. This chapter is therefore devoted largely to an analysis of how VaR is computed and used. It must be pointed out however that the margining system in most derivative exchanges around the world is closer to an ES measure of risk than to VaR. It has even been asserted⁵ that no derivative exchange in the world uses VaR for margining purposes.

22.2 HISTORICAL SIMULATION

The simplest way to estimate the loss distribution and therefore the Value at Risk is a method known as historical simulation. In this method, the actual historical fluctuation in market prices during some sample period is applied to the current market prices to arrive at different price scenarios for the future.

³ It implies for example that a bank that has sold one call option has the same level of risk as a bank that has sold a thousand call options. The worst possible outcome for both is minus infinity.

⁴ Strictly speaking these different terms are not identical if the loss distribution is discrete or discontinuous. However, as explained in footnote 1 above, the loss distribution is assumed to be continuous throughout this chapter.

⁵ Artzner, P., Delbaen, F, Eber, J.-M., Heath, D. (1999), 'Coherent measures of risk', *Mathematical Finance*, 9(3), Aw203-228. In India, the exchanges and their regulators claim to use VaR for margining purposes. On close examination, however, the margining system in its entirety is closer to ES than to VaR.

22.4 | Derivatives and Risk Management

The current portfolio is revalued using these scenario prices to arrive at scenario portfolio values and scenario gains and losses. The gains and losses from a large number of scenarios gives the loss distribution. The worst 1% or 5% of this distribution is the 99% or 95% VaR.

Table 22.1 Sample historical data on JPY/USD and USD/EUR exchange rates usedas input for historical simulation. Actual simulation would use severalhundred days of data. For illustration, only a few days of data are shown.

DATE	Exchar	ige Rate	Percenta	Percentage Change		
	JPY/USD	USD/EUR	JPY/USD	USD/EUR	•	
03-Jan-05	102.83	1.3476				
04-Jan-05	104.27	1.3295	1.4004%	-1.3431%		
05-Jan-05	103.95	1.3292	-0.3069%	-0.0226%		
06-Jan-05	104.87	1.3187	0.8850%	-0.7899%		
07-Jan-05	104.93	1.3062	0.0572%	-0.9479%		
10-Jan-05	104.32	1.3109	-0.5813%	0.3598%		
ll-Jan-05	103.42	1.3161	-0.8627%	0.3967%		
12-Jan-05	102.26	1.3281	-1.1216%	0.9118%		
13-Jan-05	102.55	1.3207	0.2836%	-0.5572%		
14-Jan-05	102.50	1.3106	-0.0488%	-0.7647%		
18-Jan-05	102.36	1.3043	-0.1366%	-0.4807%		
19-Jan-05	102.52	1.3036	0.1563%	-0.0537%		
20-Jan-05	103.31	1.2959	0.7706%	-0.5907%		
21-Jan-05	102.85	1.3049	-0.4453%	0.6945%		

This method is illustrated in Tables 22.1 and 22.2. Table 22.1 gives the historical exchange rates for two currency pairs — Japanese Yen per US Dollar (JPY/USD) and US Dollar per Euro (USD/EUR) for a sample period from January 2005. In practice, several hundred days of data would be used for the simulation, but to illustrate the method, only a small number of days of data are shown in the table. It is not possible to use the raw exchange rates for the simulation because the current exchange rates are very different from that observed in the sample period. For example, during the sample period the yen is around 102-105¥/\$. If the current exchange rate is 118¥/\$, the rates observed during the sample period are completely unrealistic. To use the historical data, the raw exchange rates are converted into percentage changes as shown in the last two columns of Table 22.1. During the sample period, the percentage changes in yen/dollar range from -1.12% to +1.40%. These percentage changes can be applied to the current exchange rate of 118¥/\$ without any loss of realism.

This is done in Table 22.2 where each day of the sample period is treated as a different scenario defined by the percentage change in the exchange rates on that day. For example, scenario 5 in Table 22.2 corresponds to 10, January 2005 in Table 22.1 and is defined by percentage changes of–0.5813% and 0.3598% in JPY/USD and USD/EUR respectively. These changes are applied to the current exchange rates of 118.00¥/\$ and 1.3500€/\$ shown in the first row of Table 22.2 to give scenario exchange rates of 117.31¥/\$ and 1.3549€/\$.

The scenario exchange rates can be used to value any currency portfolio. In Table 22.2, the portfolio is assumed to consist of ≥ 1.5 billion and ≤ 20.0 million and the investor is assumed to be a US financial

institution that measures all portfolio values in US dollars. The scenario exchange rate of 117.31 (simplies that 1.5 billion is worth \$12.79 million. Similarly the 20.0 million is worth \$27.10 million and therefore the total portfolio is worth \$39.88 million. At current exchange rates the portfolio is worth \$39.71 million. Thus the portfolio experiences a gain of \$0.17 million.

Table 22.2 Historical simulation applied to a portfolio consisting of ¥ 1.5 billion and €20.0 million using historical data from Table 22.1. The current exchange rates are shown on the first row (118.00¥/\$ and 1.3500//\$). The percentage changes from Table 22.1 are applied to the current exchange rates to arrive at the scenario exchange rates. These scenario exchange rates are used to value the portfolio and calculate gains and losses. Actual simulation would use several hundred days of data. For illustration, only a few days of data are shown

Scenario	Percentage Change		Exchange Rate		Por	Portfolio Value		
Number	JPY/USD	USD/EUR	JPY/USD	USD/EUR	¥1.5b	€20.0	Total	Loss
			118.00	1.3500	12.71	27.00	39.71	
1	1.4004%	-1.3431%	119.65	1.3319	12.54	26.64	39.17	-0.54
2	-0.3069%	-0.0226%	117.64	1.3497	12.75	26.99	39.74	0.03
3	0.8850%	-0.7899%	119.04	1.3393	12.60	26.79	39.39	-0.32
4	0.0572%	-0.9479%	118.07	1.3372	12.70	26.74	39.45	-0.26
5	-0.5813%	0.3598%	117.31	1.3549	12.79	27.10	39.88	0.17
6	-0.8627%	0.3967%	116.98	1.3554	12.82	27.11	39.93	0.22
7	-1.1216%	0.9118%	116.68	1.3623	12.86	27.25	40.10	0.39
8	0.2836%	-0.5572%	118.33	1.3425	12.68	26.85	39.53	-0.19
9	-0.0488%	-0.7647%	117.94	1.3397	12.72	26.79	39.51	-0.20
10	-0.1366%	-0.4807%	117.84	1.3435	12.73	26.87	39.60	-0.11
11	0.1563%	-0.0537%	118.18	1.3493	12.69	26.99	39.68	-0.03
12	0.7706%	-0.5907%	118.91	1.3420	12.61	26.84	39.46	-0.26
13	-0.4453%	0.6945%	117.47	1.3594	12.77	27.19	39.96	0.24

Under the given scenarios the portfolio gain/loss ranges from a gain of \$0.39 million (Scenario 7) to a loss of \$0.54 million (Scenario 1). The loss distribution is shown in Figure 22.1 but it must be kept in mind that the sample of 13 scenarios is too small to determine the loss distribution and estimate the VaR with reasonable accuracy. If the same exercise were done with 500 scenarios, the worst 5% would correspond to the worst 25 scenarios and the best of these 25 scenarios (or the worst of the remaining 475 scenarios) would be an estimate of the 95% VaR.

The historical simulation method can handle option positions quite easily. For example, in the simulation of Table 22.2, if the portfolio included a sold option of \in 50.0 million, then the historical data in Table 22.1 would have to be extended to include dollar and eurp interest rates as well as the implied volatility of euro/dollar options. With this additional data, the Black-Scholes model can be used to value the sold option in each scenario. For greater accuracy, historical data would also be needed on the volatility smile. For example, if historical data is available on the at-the-money volatility as well as on straddles and risk reversals, then a volatility smile can be fitted and used to value the sold option. Thus the nature of historical data that is needed for the simulation depends on the sophistication of the pricing model that is used to value the derivatives.

The McGraw·Hill Companies

22.6 | Derivatives and Risk Management



Figure 22.1 Loss distribution from Table 22.2. The distribution of the 13 scenarios is as follows: 1 scenario has a loss between \$0.50 million and \$0.75 million, 3 have a loss between \$0.25 million and \$0.50 million, four have a loss between 0 and \$0.25 million, f our have a gain between 0 and \$0.25 million and one scenario has a gain between \$0.25 million and \$0.50 million

Historical simulation is also easy to use for bonds and fixed income securities. Historical data is needed on interest rates at sufficient number of maturities to interpolate the entire zero yield curve. Any bond can be valued using the zero yield curve.

The historical simulation method can be summarized as follows:

- 1. Let P_0 be the value of the portfolio using current market rates and prices.
- 2. For each of *n* historical scenarios:
 - a) Use the market rates of scenario *i* to calculate the scenario portfolio value P_i
- b) Compute the change in portfolio value as δP = P_i P₀
 3. The set of δP is the profit/loss distribution. The percentiles of this give the VaR. If n = 500, the 5th worst scenario is the 99% VaR, the 25th worst scenario is the 95% VaR, and the 50th worst scenario is the 90% VaR.

The historical simulation method is very simple and requires very few modelling assumptions. Since the actual historical distribution is used, there is no need to assume the normal distribution or any other distribution. This makes the method quite robust. However, the method does have certain disadvantages. It assumes that the future will be like the past in that the distribution of price changes in the future will be the same as in the past. There are two difficulties with this assumption:

1. The expected direction and magnitude of the price change in the future can be different from that found in the past. For example, exchange rate theories assert that the currency with the higher rate of inflation tends to depreciate. During the historical sample period, one currency might have been depreciating because of this factor. At the point of time when the historical simulation

is being done, the inflation rates in the two countries might be very different and the currency which was depreciating in the sample period might be expected to appreciate in future. Similarly, the forward interest rates provide information about the expected movement of interest rates in future and this expectation might be different from what was observed in the historical sample period. This is not as serious a problem as it might appear because over daily horizons, the expected price change can be ignored. As explained in Chapter 14, it is conventional while estimating volatility to assume that the expected price change is zero.

- 2. The method also assumes that the volatility expected in future is the same as the volatility observed in the historical sample period. This is a much more serious problem. As explained in Chapter 14, volatility tends to be clustered and periods of high volatility tend to alternate with periods of low volatility. If the historical sample period happened to be one of low volatility while the current period is one of high volatility, the historical simulation method can produce severely biased estimates of the loss distribution and of the VaR. Using a large historical sample period can reduce the severity of this problem to some extent but does not eliminate it.
- 3. Historical simulation has difficulty in dealing with rare events that did not occur during the sample period. In the context of exchange rates, this is often referred to as the peso problem. The Mexican peso maintained a relatively stable exchange rate against the US dollar for several years. Though most analysts regarded the peso as being severely overvalued, the government was able to maintain this exchange rate for several years through interventions and controls. Ultimately, when the peso was allowed to find its true level, it depreciated very sharply. During the period of stability, the historical simulation would have shown very low risk for a portfolio of peso positions. Most analysts would however have factored in a small probability of a steep discontinuous correction and assessed risk differently. This problem is partly addressed using stress tests as described later in this chapter.
- 4. The method cannot handle relatively new asset classes for which a sufficiently long history of past prices is not available. Some emerging market currencies might for example have been freely traded only for short period of time. Similar problems arise with shares of a company that has only recently been listed.
- 5. A closely related problem is that of structural breaks. If economic policies change dramatically, then the past data might not be representative of future. For example, if a small East European country decides to join the European Union and adopt the Euro as its currency within the next few years, the volatility of its exchange rate might drop very substantially. Historical simulation would however still use the old volatility numbers.
- 6. The historical simulation method can be quite computationally intensive to be computed on a real time basis intra-day. It might be necessary to use fast methods like the delta-normal method discussed below for intra-day purposes while using historical simulation only for less frequent computations.

Despite these disadvantages, the historical simulation method remains a very simple, attractive and robust method for estimating the loss distribution and the VaR and is therefore extremely popular in practice especially when the portfolio is not so large as to make the method computationally expensive.

22.3 DELTA-NORMAL APPROXIMATION

The delta-normal method is based on the normal distribution. If the loss distribution (or equivalently the portfolio value) is assumed to be a normal distribution, then the VaR at any level (for example, 90%, 95%

The McGraw·Hill Companies

22.8 | Derivatives and Risk Management

or 99%) can be read off the normal distribution tables as indicated in Figure 22.2. For example, the 99% VaR corresponds to the worst 1% of outcomes. Figure 22.2 shows that the area under the standard normal distribution to the left of -2.33 is 1 %. This implies that the 99% VaR is at the point 2.33 standard deviations below the mean. Similarly, the 95% VaR (worst 5% outcomes) is 1.64 standard deviations below the mean and the 90% VaR (worst 10% outcomes) is 1.28 standard deviations below the mean. To compute the VaR at any level, it is sufficient to know the mean and standard deviation of the portfolio value.



Figure 22.2 Using the standard normal distribution to estimate Value at Risk. The 99% VaR corresponding to the worst 1% of outcomes is 2.33 standard deviations below the mean, the 95% VaR (worst 5% outcomes) is 1.64 standard deviations below the mean and the 90% VaR (worst 10% outcomes) is 1.28 standard deviations below the mean. To compute the VaR, it is sufficient to know the mean and standard deviation of the portfolio value

The easiest way to obtain a normal distribution for the portfolio value is to assume that the underlyings follow a normal distribution. If the *i*'th underlying has mean return μ_i th and volatility σ_i and the portfolio has a delta Δ_i with respect to the *i*'th underlying, then the change in the portfolio value has a mean of $\sum_i \Delta_i S_i \mu_i$ where S_i is the price of the *i*'th underlying. To find the variance of the portfolio value, it is

necessary to know the covariances σ_{ij} or the correlations ρ_{ij} The formula is $\sum_{ij} \Delta_i \Delta_j S_i S_j \sigma_{ij}$ or

$\sum_{ij} \Delta_i \Delta_j S_i S_j \sigma_i \sigma_j \rho_{ij}$

If the portfolio has only cash, forward or futures positions in the various underlyings, then the above formulas for the mean and variance of the portfolio value are exact. If the portfolio has option positions, then these formulas are approximate as they ignore the effect of gamma, vega and other Greeks. A better estimate of portfolio value for option positions is discussed in the next section. However, the delta approximation has a big advantage over these better approximations—it ensures that the portfolio value follows a normal distribution. This allows the VaR to be estimated using normal distribution tables.

When the portfolio contains bonds, each bond can be decomposed into a portfolio of zero coupon bonds by treating each cash flow as a zero coupon bond maturing on the cash flow date. When the portfolio contains a large number of bonds and other fixed income securities, the number of cash flow dates can become very large and this approach leads to a very large number of zero coupon bonds in the decomposition. It becomes impractical to treat each of these bonds as a different underlying. This is because the delta-normal method requires the volatility of each underlying and also the correlation of each underlying with each other underlying. It is common therefore to use a different approach known as cash flow mapping while implementing the delta-normal method. After decomposing each bond into its component cash flows, the cash flows in turn are mapped into a set of zero coupon bonds of standard maturities.

For example, the standard maturities may be 1 month, 3 months, 6 months, 1 year, 2 years, 3 years, 5 years, 7 years, 10 years and 30 years. A cash flow arising at the end of $4\frac{1}{2}$ years is then mapped into the nearby standard maturities of 3 years and 5 years. The naive method of doing this might be to map it 25% of the cash flow into the 3 year bond and 75% into the 5 year bond so that the average duration of the mapped cash flows is equal to the duration of the original cash flow ($25\% \times 3 + 75\% \times 5 = 0.75 + 3.75 = 4.5$). While estimating VaR, the goal is to match variances and not to match durations. To match variances, it is necessary to consider the volatility of the 3 year and 5 year year bond as well as the correlation between the returns on these two bonds.

Suppose for example that the annualized volatilities are 3.5% and 5% and the correlation is 0.98. By interpolation, the volatility of the 4¼ year bond is 4.625% ($25\% \times 3.5\% + 75\% \times 5\% = 0.875\%$ + 3.750% = 4.625%). If a fraction *w* of the cash flow is mapped into the 3 year bond and a fraction 1 - w into the 5 year bond, then the variance of the mapped cash flow is $w^2 \times 0.035^2 + (1 - w)^2 \times 0.05^2$ + $2 w (1 - w) \times 0.035 \times 0.05 \times 0.98$. The best mapping is given by the choice of *w* that makes this variance equal to $0.04625^2 = 0.002139$. Either by trial and error or by solving the quadratic equation, the optimal *w* is seen to be 0.2408 or 24.08%. This is close to the 25% obtained by the naive method. In practice, the value obtained by matching variances is not far from that obtained by matching durations.

The delta-normal method has two important advantages:

- 1. It is computationally very attractive. Portfolios with several tens of thousands of positions spread across hundreds of underlyings can be handled very fast.
- 2. The data needs for the approach are quite modest. It needs volatilities of each underlying as also the pair wise correlations. Using the methods described in Chapter 14, it is quite straight forward to estimate volatilities on a daily basis for several hundreds of underlyings. Later in this chapter, it will be shown that similar methods can be used for the correlations as well.

Many financial institutions use an approach of this kind for VaR computations to be done on a daily basis. An organization called RiskMetrics provides daily estimates of volatilities and correlations for hundreds of stock indices, exchange rates, interest rates, commodity prices to facilitate the use of this method. Many software tools have also been developed for this purpose.

The main disadvantages of the method are:

- 1. Its reliance on the normal distribution is problematic for VaR at high levels like 99% or 99.5%. If VaR is desired only at 90% or 95%, the normal distribution is often sufficiently close to the true distribution for the delta-normal method to give acceptable results.
- 2. The delta-normal method does not account for option positions accurately. It takes account only of their deltas and not their gammas or other Greeks. This issue is discussed in the next section.

The McGraw·Hill Companies

22.10 | Derivatives and Risk Management

22.4 DELTA-GAMMA APPROXIMATION

Options have payoffs that are asymmetric — equal up and down moves in the underlying do not produce equal gains and losses in the option position. For example in a short option position, the potential losses are unlimited while the potential gain is limited to the option value. This is seen in Figure 22.3 where the loss distribution of a short option position has been depicted. The potentially unlimited losses are reflected in the long left tail. The right tail is quite short indicating that the gains are limited. However, a small gain is the most likely outcome represented by the fact that the mode of the distribution is slightly positive.

Most of the time, the option seller will pocket the premium as the option expires worthless or worth very little. However, this is offset by the small likelihood of large losses. This is another illustration of the insurance nature of options. An insurance company also simply pockets the premium on most policies that it sells but loses heavily on the few policies that lead to claims.



Fig. 22.3 Loss distribution for a short option position is skewed to the left reflecting the asymmetric nature of the option payoffs. Potential losses are unlimited as indicated by the long tail of the distribution. Potential gains are limited as indicated by the short right tail. The mode is positive indicating that the most likely outcome is a slight gain

This skewed distribution is evidently far from being a normal distribution and estimates of the VaR based on the normal distribution are likely to be seriously in error. A proper estimate of the VaR must take account of the skewnesss.

The delta-gamma method is one way of dealing with the problem. The delta-normal method models the change in portfolio values using only the delta. Where there is only one underlying, the delta normal approximation is:

$$\delta P = \Delta S(\delta S) \tag{22.1}$$

where δP is the change in the portfolio value, δS is the change in the underlying, Δ is the delta of the portfolio and S is the current market price of the underlying.

The delta-gamma improves upon this by taking account of the gamma as well:

$$\delta P = \Delta S \left(\delta S\right) + \frac{1}{2} \Gamma S^2 \left(\delta S\right)^2 \tag{22.2}$$

where Γ is the gamma of the portfolio.

The delta-gamma approximation of Eq (22.2) has the following implications:

1. The variance which equals $S^2 \Delta^2 \sigma^2$ under Eq (22.1) has to be corrected by adding a term to capture the effect of the gamma. The corrected variance is

$$S^{2} \Delta^{2} \sigma^{2} + \frac{3}{4} S^{4} \Gamma^{2} \sigma^{4}$$
(22.3)

2. The loss distribution is not symmetric but has a skewness given by

$$\frac{9}{2} S^4 \Delta^2 \Gamma \sigma^4 + \frac{15}{8} S^6 \Gamma^3 \sigma^6$$
(22.4)

Adjusting for the first effect is very simple and consists of using the corrected standard deviation multiplied by the value obtained from normal distribution tables. For example, Figure 22.2 shows that the 99% VaR is given by 2.33 standard deviations. Therefore the corrected standard deviation given by Eq (22.3) has to be multiplied by 2.33 to give the 99% VaR adjusting only for the variance correction. For example, consider an option position for which S = 100, $\Delta = -0.50446$, $\sigma = 1\%$ per day, $\Gamma = 0.178401$. The variance given by Eq (22.1) is $S^2 \Delta^2 \sigma^2 = 100^2 \times (-0.50446)^2 \times 0.01^2 = 0.25448$ implying a standard deviation of 0.50446. The VaR estimate using this standard deviation would be $-2.33 \times 0.50446 = -1.17539$. Eq (22.3) corrects the variance by adding a term $\frac{3}{4} S^4 \Gamma^2 \sigma^4 = \frac{3}{4} \times 100^4 \times (0.178401)^4 \times 0.01^4 = 0.02387$. The variance then becomes 0.25448 + 0.02387 = 0.27835 implying a standard deviation of 0.527589. The VaR estimate using this standard deviation would be $-2.33 \times 0.527589 = -1.22928$.

Adjusting for the skewness is more problematic. One approach⁶ consists of modifying the multiplier (for example -2.33 for 99% VaR) to account for skewness using the formula:

$$w = z + \frac{1}{6} (z^2 - 1) H$$
(22.5)

where w is the corrected percentile, z is the percentile based on the normal distribution and H is the skewness given by Eq (22.4).

For example, consider the earlier option position for which S = 100, $\Delta = -0.50446$, $\sigma = 1\%$ per day, $\Gamma = 0.178401$. Substituting these numbers into Eq (22.4) gives a skewness value of H = -0.214944. Plugging this value into Eq (22.5) along with z = -2.33 (for 99% VaR) gives the corrected percentile as w = -2.49. The 99% VaR for this option should therefore be computed using 2.49 standard deviations instead of 2.33 standard deviations. The corrected standard deviation was computed above as 527589. The VaR adjusted for both the variance correction and the skewness

⁶ This approach is known as the Cornish Fisher approximation.

22.12 | Derivatives and Risk Management

correction is therefore $-2.49 \times 0.527589 = -1.3137$. While the variance correction was only 0.05 (1.18 to 1.23), the skewness correction is larger at 0.09 (from 1.23 to 1.31).

The above discussion was based on just one underlying. When there are multiple underlyings, the delta-normal approximation is given by

$$\delta P = \sum_{i} \Delta_i S_i \delta S_i \tag{22.6}$$

where δP is the change in the portfolio value, δS_t is the change in the *i*'th underlying, Δ_i , is the delta of the portfolio with respect to the *i*'th underlying and S_i is the current market price of the *i*'th underlying.

The delta-gamma approximation with multiple underlyings is given by:

$$\delta P = \sum_{i} \Delta_i S_i \delta S_i + \frac{1}{2} \sum_{i,j} \Gamma_{ij} S_i S_j (\delta S_i)^2$$
(22.7)

where Γ_{ij} is the usual gamma with respect to the *i*'th underlying and Γ_{ij} is a cross-gamma that measures how the delta with respect to the *i*'th underlying changes when there is a movement in the *i*'th underlying.

22.5 MONTE CARLO SIMULATION

The Monte Carlo simulation method is based on the ability of computer software to generate random numbers. These random numbers can be used to generate samples from any given distribution. For example, if the distribution is normal and the means and variances are given, the software can generate a sample of thousands of scenarios consistent with these means and variances.

22.5.1 Monte Carlo with Full Revaluation

Like the historical simulation method, this method is also based on simulation. The difference is that the scenarios are not based on historical data but are generated by computer software based on a specified theoretical distribution. The Monte Carlo method of estimating VaR with full revaluation can be summarized as follows:

- 1. Let P_0 be the value of the portfolio using current market rates and prices.
- 2. For each of *n* Monte Carlo simulated scenarios:
 - a) Use the market rates of scenario i to calculate the scenario portfolio value P_i by revaluing each position in the portfolio
 - b) Compute the change in portfolio value as $\delta P = P_t P_0$
- 3. The set of δP is the profit/loss distribution. The percentiles of this give the VaR. If n 5000, the 50th worst scenario is the 99% VaR, the 250th worst scenario is the 95% VaR, and the 500th worst scenario is the 90% VaR.

Functions for generating random numbers are usually included as part of general purpose statistical, financial and business software. For example, Microsoft Excel includes these functions. These functions produce random numbers which are distributed uniformly from 0 to 1 and can thus be used to generate samples from a uniform distribution. It can be shown that if *x* is distributed uniformly from 0 to 1, then $y = N^{-1}(x)$ follows the standard normal distribution. Moreover $z = \mu + \sigma y$ follows a normal distribution with mean μ and standard deviation σ .

To handle several underlyings, it is necessary to generate a sample from a multivariate normal distribution with specified means, variances and correlations. One can generate say a hundred normal variates and instead of regarding it as a sample of 100 from a univariate normal distribution, regard this

as a sample of 1 from a multivariate normal distribution with a hundred underlyings. The only problem is that these hundred underlyings are then uncorrelated or independent. To generate a sample from a correlated multivariate normal distribution, it is necessary to do some further work.

Suppose it is desired to generate a sample of *n* variables $u_1, u_2, ..., u_n$ with means $\mu_1, \mu_2, ..., \mu_n$, standard deviations $\sigma_1, \sigma_2, ..., \sigma_n$ and covariances σ_{ij} . Suppose also that $y_1, y_2, ..., y_n$ are *n* independent standard normal variables. This means that the y_i have zero means, unit standard deviations and zero covariances with each other. Then the $u_1, u_2, ..., u_n$ can be generated as follows⁷:

- 1. Let $u_1 = \mu_2 + a_1 y_1$. Clearly it is sufficient to set $a_1 = \sigma_1$ to match the mean and variance of u_1 .
- 2. Let $u_2 = \mu_2 + b_1 y_1 + b_2 y_2$. In order for the covariance between u_1 and u_2 to be equal to σ_{12} , it is necessary to ensure that $a_1 b_1 = \sigma_{12}$. This can be done by setting $b_1 = \frac{\sigma_{12}}{a_1}$ where a_1 has the value determined above. To match the variance of u_2 , it is necessary to ensure $b_1^2 + b_2^2 = \sigma_2^2$. This can be done by setting $b_2 = \sqrt{\sigma_2^2 b_1^2}$
- 3. Let $u_3 = \mu_3 + c_1 y_1 + c_2 y_2 + c_3 y_3$. In order for the covariance between u_1 and u_3 to be equal to σ_{13} , it is necessary to ensure that $a_1 c_1 = \sigma_{13}$. This can be done by setting $c_1 = \frac{\sigma_{13}}{a_1}$ where a_1 has the value determined above. In order for the covariance between u_2 and u_3 to be equal to σ_{23} , it is necessary to ensure that $b_1 c_1 + b_2 c_2 = \sigma_{23}$. This can be done by setting $c_2 = \frac{\sigma_{23} b_1 c_1}{b_2}$ where all the variables on the right hand side have already been determined. To match the variance of u_3 , it is necessary to ensure $c_1^2 + c_2^2 + c_3^2 = \sigma_3^2$. This is achieved by setting $c_3 = \sqrt{\sigma_3^2 c_1^2 c_2^2}$.
- 4. Let $u_4 = \mu_4 + d_1 y_1 + d_2 y_2 + d_3 y_3 + d_4 y_4$. Proceeding in the same manner, the values of d_1 , d_2 , ... d_4 can be determined.
- 5. Proceeding in the same manner, all the variables $u_1, ..., u_n$ can be obtained so as to match the given means, variances and covariances.

22.5.2 Partial Simulation (Delta Gamma Approximation)

The Monte Carlo method is computationally quite demanding because thousands of scenarios are generated and in each scenario, the entire portfolio is revalued. One way to speed up the computations is to approximate the portfolio value using the delta gamma approximation of Eq (22.7). The advantage of this approximation is that the computational time for this depends only on the number of underlyings and not on the number of positions. For a portfolio with tens of thousands of positions on a small number of underlyings, the delta-gamma approximation reduces the computational time hugely. However, the VaR estimate is less accurate than with the full revaluation method.

The Monte Carlo method of estimating VaR with delta-gamma approximation can be summarized as follows:

1. Let P_0 be the value of the portfolio using current market rates and prices.

⁷ This method is known as Cholesky decomposition. This is the most popular but not the only method available for this purpose.

- 22.14 | Derivatives and Risk Management
 - 2. For each of *n* Monte Carlo simulated scenarios:
 - (a) Use the market rates of scenario i to calculate the scenario portfolio value P_t using Eq (22.7)
 - (b) Compute the change in portfolio value as $\delta P = P, -P_0$
 - 3. The set of σP is the profit/loss distribution. The percentiles of this give the VaR. If n = 5000, the 50th worst scenario is the 99% VaR, the 250th worst scenario is the 95% VaR, and the 500th worst scenario is the 90% VaR.

22.6 MODELLING AND ESTIMATING CORRELATIONS

Chapter 14 discussed several methods for estimating volatilities. All these methods can be used to estimate covariances or correlations also. This is because the variance of a variable is its covariance with itself. Covariance is thus simply a generalization of the concept of variance.

The most critical point is that all variances and covariances must be estimated using a single consistent method. For example, if the the volatility of two underlyings is estimated using the historical volatility of the last 30 days and their covariance is estimated using historical volatility of the last 50 days, serious problems can arise. For example, the covariance is equal to the product of the standard deviations if the correlation equals unity and is less than the product if the correlation is less than unity. Suppose however, the two underlyings has very low volatility in the last 30 days but much higher volatilities in the 20 days prior to that. The volatilities based on 30 days would be quite low but the covariance based on 50 days would be high. It is quite possible that the estimated covariance exceeds the product of the standard deviations.

An estimate of this kind is not only theoretically unsound but can cause enormous practical difficulties. For example, the variance of a portfolio which is long one asset and short another asset in equal magnitudes is $\sigma_x^2 + \sigma_y^2 - 2\sigma_{xy}$. If by using an unsound method, the estimates are $\sigma_x = \sigma_y = 0.2$ and $\sigma_{xy} = 0.041$ then the estimated variance of the portfolio would be $0.2^2 + 0.2^2 - 2 \times 0.041 = 0.04 + 0.04 - 0.082 = -0.002$ implying that the portfolio has a negative variance and therefore negative risk. This is clearly unacceptable in theory and in practice.

For this reason, it is important to ensure that when estimating portfolio VaR all volatilities and covariances are estimated using the same method with the same parameters. If historical estimates are used, then the same sample period must be used for all volatilities and all covariances. If the exponentially weighted moving average method is used, the same coefficient λ must be used for all underlyings both for volatilities and for covariances. It is a serious mistake to use say $\lambda = 0.90$ for one underlying and say $\lambda = 0.95$ for another underlying.

In practice, the most popular approach is that adopted by RiskMetrics which uses the exponentially weighted moving average method with the same coefficient $\lambda = 0.94$ for all underlyings both for volatilities and covariances. After examining a large number of asset classes around the world, the Risk Metrics team chose 0.94 as the coefficient that gives reasonable results for most underlyings.

22.7 BACK TESTING

Whatever method is used to estimate VaR, the validity of the estimates must be checked by back testing. In back testing, the chosen estimation model is run for several hundred days in the past and the portfolio loss each day is compared with the VaR estimate at the beginning of the day (or equivalently at the end of the previous day). If the VaR model is correct, than the actual loss would exceed the 99% VaR on

approximately 1% of the days. If the actual loss never exceeds the 99% VaR in a large sample, it is apparent that the model is over estimating the VaR. If on the other hand, the 99% VaR is actually exceeded on 10% of the days in a large sample, the model is clearly underestimating the VaR.

The BIS Committee on Banking Supervision for example suggest that if a 99% VaR is back tested with 250 days of data, then the results can be divided into green, yellow and red zones as follows:

- *Green Zone:* The 99% VaR is exceeded less than 5 times and there is little evidence that the model is underestimating VaR. Theoretically, the number should be 1% of 250 or 2.5. However due to random fluctuations, the number may be a little less or a little more. Statistical analysis shows that up to four incidents of loss exceeding VaR can arise by pure chance and therefore there is no cause for concern.
- *Yellow Zone:* The 99% VaR is exceeded more than four times but less than 10 times. There is some evidence that the model is underestimating VaR but the evidence is not conclusive. It is necessary to look at additional data to validate the model. Alternatively, the VaR given by the model may be scaled up by using a multiplicative factor. For example, if a 99% VaR is exceeded on 5 days out of 250, then it is actually a 98% VaR and not a 99% VaR because 5 is 2% of 250. According to the normal distribution, the 99% VaR is at 2.33 standard deviations while the 98% VaR is at 2.05 standard deviations. If the 98% VaR produced by the model is scaled up by a factor
 - of $\frac{2.33}{2.05} = 1.13$, then the result would be a 99% VaR. This scaling factor is of course critically

dependent on the assumption of normality.

• *Red Zone:* The 99% VaR is exceeded 10 or more times implying that the actual level of the VaR is 96% or less. The VaR model is rejected as inappropriate. A better VaR model must be built for risk management purposes.

These three zones are from a regulator's perspective where an underestimation of risk is serious but an overestimation is harmless. From the point of view of the institution itself, overestimation is equally problematic and there should perhaps be a 'black zone' where there are too few days on which the estimated VaR is exceeded. This zone too should lead to revisiting the assumptions of the model to determine whether they are too conservative.

22.8 STRESS TESTING

The Value at Risk techniques described so far in this chapter deal with normal business conditions during which past behaviour of risk factors can be meaningfully used to predict future behaviour. The past behaviour may be used directly as in historical simulation or may be used only to estimate the volatilities and correlations in the delta-normal or Monte Carlo methods.

Stress events are very different. They are characterized by the following features:

- 1. There is a sharp rise in volatilities reflected in extreme movement in market prices.
- 2. The correlation structures observed in the past often breakdown during periods of stress. Sometimes this takes the form of a sharp fall in correlations between assets that are highly correlated under normal business conditions. For example, during the Asian Crisis and Russian default of 1997 and 1998, the correlation between government bond yields and swap yields fell dramatically. This played havoc with various hedging strategies. Stress events can also lead to a sharp rise in correlation between assets that were historically unrelated to each other. For example, during

22.16 | Derivatives and Risk Management

1997 and 1998, correlations between unrelated emerging market currencies increased sharply. This invalidates key assumptions about diversification benefits in large portfolios.

- 3. Illiquidity in key financial markets makes it difficult to exit loss making positions or to implement dynamic hedging strategies.
- 4. As a result of (2) and (3) above, long established hedging relationships, arbitrage conditions and pricing models break down making it very difficult to value complex derivative portfolios.

The risk management literature therefore emphasizes that the VaR models predicated on normal business conditions should be supplemented by stress tests that evaluate the impact of stress events. The BIS Committee on the Global Financial System defines stress testing as 'techniques used by financial firms to gauge their potential vulnerability to exceptional but plausible events.' This definition captures the key elements that make stress testing so difficult in practice - the stress events must be (a) exceptional and (b) plausible. Unlike the scenarios used in VaR estimation, the stress event is an exceptional event that may be observed only once in a decade or so. At the same time, the stress test is useful only if the event is plausible. Totally outlandish scenarios are of little use in practical risk management.

Good stress testing programmes rely on the knowledge of a number of experts to identify potential stress events. The key stress events are based on:

- Understanding the nature of the portfolio and its vulnerability to various risk factors: For example, if the portfolio duration has been carefully managed to immunize the portfolio against changes in the level of interest rates, it is pointless to consider stress events that consist of large parallel shifts of the yield curve. On the other hand, stress events that include changes in the slope or curvature of the yield curve may be more appropriate as the portfolio may be more vulnerable to such events. Similarly, stress events that include large changes in volatilities and correlations may be quite relevant for portfolios with large option positions.
- Understanding the external environment and evaluating the plausibility of different stress events: For example, an appraisal of the economic environment may suggest that exchange rate stress events are more plausible than stock market stress events.
- *Including inputs from contrarian thinkers:* Many organizations fall victim to a form of group think in which there is a consensus on risk factors and vulnerabilities. It may then be useful to consult contrarian thinkers who might highlight plausible stress events that others have ignored.

Scenarios for stress tests are often generated by replicating historical stress events. For example, many stress tests of stock market portfolios include a stress event that consists of a price fall equal in speed and intensity to the crash of 1987 in the US and many other markets. Similarly, the collapse of the European exchange rate mechanism in the early 1990s is widely used as a stress event for currency risk.

Stress tests also use scenarios that have been identified by experts as plausible though they have never occurred in the recent past. For example, a default by an investment grade sovereign has not occurred since the second world war, but a stress test might legitimately include such a scenario.

Based on the above considerations, many financial institutions use two kinds of stress events;

- Routine standardized stress tests that are repeated at periodical intervals and help to monitor the change in the risk profile of the portfolio over time. For example, these tests may be run unchanged every quarter and the results may be compared with those of the previous quarter to assess whether the portfolio has become less risky or more risky.
- Customized stress tests based on current portfolio vulnerabilities as well as the economic environment which would be different on each occasion. These tests may be more plausible and

relevant than the routine events but there is no reference point in the past to which the results can be compared.

22.9 INTERNAL CONTROL SYSTEMS

In Chapter 1, we discussed some well known derivative disasters in banks and financial institutions and concluded that tight internal control systems are necessary to guard against them. This section describes some of these systems.

22.9.1 Separation of functions

Derivative activities in a financial institution involve several functions:

- *Investment or position taking:* This could for example be a decision to take a long position in Italian government bonds. This is based on fundamental or technical analysis of the underlying in relation to the investment objectives of the financial institution itself.
- *Trading:* The actual trades that lead to the achievement of the above position form a separate step and may take several days to implement depending on the liquidity of the market and the size of the position being taken.
- *Settlement:* Cash market trades have to be settled by delivery of cash and the underlying. In derivative trades, mark to market cash flows have to be settled and if the derivative is exercised, the resulting cash market position has to be settled. For example, if a call option is exercised, the bank has to pay out cash and take delivery of the underlying.
- *Custody:* The assets or derivative contracts owned by the bank have to kept in safe custody so that they are not lost or stolen.
- *Valuation and modelling:* There is a need to produce daily or periodic valuations of the entire derivative portfolio for internal control purposes.

When several of these functions are performed by the same individual, frauds are very easy. Ideally each function above must be separated, but a small bank may be forced to combine a few functions with each other because its derivative team is very small. Even then trading and settlement should never be combined. In the Nick Leeson example discussed in Chapter 1, Leeson was able to hide many trades because he had control over both trading and settlement. In a well run bank, this should never happen. Similarly, the valuation of derivatives should be done by a team different from the trading team. In many banks, there is a mid office responsible for valuation and modelling. The mid office should have a strong team of finance and mathematics experts who understand all the derivative instruments traded by the organization.

22.9.2 Trading Limits

Every trader and every trading team should have trading limits that define the maximum amount of positions that they can take. These limits can be set in terms of notional value. For example, a currency trader may be told that her position should not exceed \$10 million on an intra-day basis and \$1 million at end of day. Alternatively, the limits can be set in terms of the delta, gamma, and vega of the portfolio. Most commonly, however, trading limits are set in terms of VaR. Instead of computing the enterprise wide VaR, it is possible to compute the VaR of the positions taken by a single trader or a group of traders. Limits are set on these VaRs at different levels in the organization. This process is in fact one easy way to limit the enterprise wide VaR.

22.18 | Derivatives and Risk Management

It is absolutely important to enforce these trading limits very strictly. There is often a temptation to relax trading limits for traders who are performing well. This is a mistake for two reasons. First, good performance may be the result of good luck and relaxing the trading limit could produce large losses in future. The value of diversification of positions across different traders and across different asset classes must always be kept in mind. Second, good performance may be the result of high risk positions and relaxing the trading limits for high risk traders is very dangerous.

22.9.3 Controlling Model Risk and Liquidity Risk

Banks that trade complex products, end up becoming highly reliant on their own models for pricing and for valuation. Complex products tend to be illiquid and often the mark to market process for these instruments ends up being a mark to model. Controlling model risk is then very important. One way to control this is to have a strong mid office that has its own valuation model that could be different from that used by the trading team for pricing. Another control is the trading limits that ensure that the bank is not overexposed to a single asset class or instrument.

Dynamic trading strategies are crucially dependent on liquidity in the underlying market for implementing the delta hedges. If this liquidity fails, the position can become unhedgeable. Similarly, the VaR model relies on the position being liquidated within the VaR horizon. For example, if a bank uses a 1-day VaR for internal purposes, the assumption is that the position can be liquidated within a day if necessary. This too presumes that there is liquidity in the market. Illiquidity can pose severe risks to those running large derivative books. The LTCM example discussed in Chapter 1 is a good example of the effect of liquidity risk.

22.9.4 Guarding against Herd Behaviour

Sometimes, a bank finds that everyone has similar views on certain economic variables and that this leads to a one-sided market. Such markets are dangerous for two reasons. First, a one-sided market in which there are no constrarian traders can suddenly become illiquid when everybody wants to exit at the same time because there is nobody to take the opposite side of the trade. Second, these markets can produce severe mispricing of derivatives that can also lead to losses when the mispricing is corrected.

22.10 REGULATORY CONSIDERATIONS

Value at risk is used by banking and other regulators to set capital for financial institutions. Until the 1990s, regulators used relatively simple formulas and rules of thumb to determine the capital requirements for various kinds of risks. In the last couple of decades, there has been a recognition that these formulas are too crude and oversimplified to be of much use in modern financial markets. Regulators have come around to relying more and more on the internal models used by the financial institutions themselves. While doing so however, the regulators have imposed a variety of conditions and restrictions which have had a great deal of influence in the practice of risk management.

For example, the BIS requirements for VaR computations for market risk include the following requirements:

- 1. VaR must be computed at the 99% probability level.
- 2. VaR should be estimated for a time horizon of 10 trading days. Most of the examples in this chapter have implicitly used a 1 -day time horizon. Direct estimation of the 10-day VaR is regarded

as impractical and in practice, the 10-day VaR is obtained from the 1-day VaR by using a scaling factor of $\sqrt{10} = 3.16$. In other words, the 10-day VaR is simply 3.16 times the 1-day VaR. This scaling is based on the normal distribution and often overestimates the VaR when the true distribution has fat tails⁸.

- 3. VaR estimates must be produced daily by a risk department that is independent of the business trading departments and must report to senior management.
- 4. Back testing must be done regularly.
- 5. A routine and rigorous process of stress testing must be in place.
- 6. The historical sample period (used for historical simulation or for estimating volatilities and correlations) should be a minimum of one year.
- 7. For each currency, stock market and commodity in which the financial institution has a significant position, the risk factors must include (a) a risk factor for each currency, (b) at least the market index for each stock market and (c) a risk factor for each commodity. In each currency, a minimum of six factors (maturities) must be used to model the yield curve.
- 8. The gamma and vega risk of options must be considered while computing the VaR.

These regulatory requirements dictate the choice of internal risk management models in most financial institutions even where some of these are of dubious value from a theoretical point of view.

Chapter Summary

The Value at Risk (VaR) at the x% probability level is defined as the level of capital that is sufficient x% of the time; or the level of loss that is exceeded only 100 - x% of the time; or the worst of the best x% of possible outcomes; or the best of the worst 100 - x% outcomes. VaR is measured over a specified time horizon—typically one day. Despite its shortcomings (as compared to other measures like Expected Shortfall), VaR is very popular among financial institutions and their regulators.

The historical simulation method applies the actual historical fluctuation (percentage changes) in market prices during some sample period to the current market prices to arrive at different price scenarios for the future. The current portfolio is revalued using these scenario prices to arrive at scenario portfolio values and scenario gains and losses. The worst 1% of this distribution is the 99% VaR.

The delta-normal approach assumes a normal distribution for each underlying and uses the delta of the portfolio with respect to each underlying to derive a normal distribution for the portfolio value. The VaR is then computed using tables of the standard normal distribution. The delta-gamma method refines this further by considering the gammas and cross-gammas of the portfolio.

The Monte Carlo method generates samples conforming to specified means, volatilities and covariances using the random numbers produced by a computer. Revaluation of all positions under these scenarios produces portfolio values under these scenarios. The worst 1% of these is the 99% VaR.

VaR models must be back tested against several hundred days of past data. Ideally, the 99% VaR should be exceeded 1% of the time. If the back test reveals significant departure from this ideal, the VaR might have to be adjusted and in extreme cases, the VaR model itself may have to be modified to be more consistent with the theoretical expectation.

⁸ This is because in practice, the fat tails are more pronounced at daily time horizons and are accentuated at longer time horizons.

22.20 | Derivatives and Risk Management

While VaR measures the risk under normal business conditions, stress tests analyse the risk under exceptional but plausible scenarios in which extreme price movements, sharp rises in volatilities, breakdown of correlation, loss of liquidity might take place. Rigorous stress testing is a critical element of the risk management system for financial institutions.

Suggestions for Further Reading

The deficiencies of value at risk as compared to measures like expected shortfall are highlighted in:

Artzber, P., Delbaen F. Eber, J.M. and Heath, D (1999) "Coherent Measures of risk", *Mathematical Fiance*, 9(3),203-228.

Szego, Giorgio (2004) Risk Measures for the 21st Century, New York, Wiley

The regulatory approach to value at risk is given in:

- Basle Committee on Banking Supervision (1996 updated 2005) "Amendment to the capital to incorporate market risks" <u>http://www.bis.org/publ/dcbsll9.pdf</u>
- Basle Committee on Banking Supervision (1996) "Supervisory framework for the use of 'backktesting' in conjunction wuth the internal models approach to marker risk capital requirements", http://www.bis.org/publ/dcbs22.pdf
- The problems created by fat tailed distributions in calue at risk are explored in:
- Jorion, Philippe (1997) "In Defense of VAR", Derivatives strategy,
- http://www. Derivativativesstrategy. Com/magazine/archive/1997/0497fea2.asp
- Kolman, Joe and Taleb, Nassim (1997) "The World According to Nassim Taleb", Derivatives Strategy,

http://www.derivativativesstrategy.com/magazine/archive/1997/1296qa.asp

One of the earlist and most influential practitioner oriented methodologies for computing value at risk is:

J.P.Morgan Reuters, *RiskMetrics Technical Document*, http://www.riskmetrics.com/rmcovv.html

An excellent overview of the conceptual issues in managing risk in financial institutions is provided by:

Metron, R C. and Perold, A.(1993) "Management of risk capital in financial firms", Chapter 8 in *Hayes, S. L, ed,* (1993) Financial services: perspectives and challenges, Boston, Harvad Business School Press, 215-245. The McGraw·Hill Companies

Index

Accounting 1.2, 1.3, 1.11, 1.12, 5.13, 5.14, 10.9, 10.12, 11.9, 11.10, 15.7, 18.27, 20.1–20.7, 20.9–20.16, 20.18–20.20, 21.6, 21.13, 21.14

- Accrued interest 2.13
- Annuity 19.11, 19.15, 19.16
- Anonymous trading 1.6, 2.3, 7.1
- Arbitrage 1.7–1.9, 1.11, 1.12, 3.1, 3.3–3.7, 3.11–3.13, 8.1, 8.3, 8.4, 8.6,–8.10, 8.13, 8.14, 9.3, 10.12, 12.5, 21.7, 21.9, 22.16,
- Arithmetic mean 16.10
- Asset 1.1, 1.4, 1.6, 1.7, 1.12, 2.2–2.6, 2.13, 3.1–3.12, 4.2, 5.14, 6.1, 6.6, 6.7, 6.8, 6.11, 6.12, 7.2, 8.2–8.5, 8.13, 10.6, 10.7, 10.12, 10.19, 10.20, 11.1, 11.9, 11.10, 12.1, 12.2, 12.3, 12.4, 12.5, 12.6, 12.8, 12.9, 12.11–12.16, 13.1–13.7, 13.9, 14.1, 14.5, 14.7, 14.8, 14.10–15.4, 15.8, 16.1–16.12, 17.11, 18.4, 19.3, 19.14, 19.15, 20.2, 20.3, 20.6–20.10, 20.20, 21.12, 21.15, 21.16, 21.22, 21.29, 21.30, 22.7, 22.14, 22.18
- Assignment 7.2–7.4
- Asymmetry 10.1, 12.6, 15.5
- Available for Sale (AFS) 20.9
- Back testing 21.7, 22.14, 22.19
- Backwardation 3.1, 3.11, 3.12, 21.23, 21.31
- Bank 1.8, 1.9, 1.14, 2.2, 2.4, 2.9, 2.10, 4.1, 4.15, 4.16, 5.1, 5.3, 5.4, 5.6, 5.12–5.14, 77, 5.18, 5.23, 5.24, 6.7, 7.6, 15.9, 17.15, 17.22, 17.23, 17.26–17.28, 17.31, 18.2, 18.20–18.24, 18.26, 18.29, 18.30, 19.1, 19.2, 20.9, 20.10, 21.2, 21.5, 21.7, 21.9, 21.10, 21.12, 21.16, 21.22, 21.23, 21.30, 21.31, 22.2, 22.3, 22.17, 22.18
- Bank for International Settlements (BIS) 2.9, 2.10, 7.5, 22.14, 22.5, 22.17,
- Bankruptcy 1.8, 5.2, 5.6, 17.16, 21.1, 21.3–21.6, 21.8,
- Barrier options 16.4, 16.5, 16.7, 16.8, 16.12,
- Basis risk 4.1, 4.5, 4.7, 4.15, 5.3, 5.4, 20.10, 20.11, 20.12,
- Bear spread, 11.8–11.10
- Benchmark 4.16, 18.6, 19.1, 19.2, 20.8, 20.17, 21.11
- Benefit 3.2, 3.9, 5.1, 5.2, 6.10, 7.9, 11.3, 11.4, 19.1, 21.13, 21.15, 21.29
- Benefits 1.1, 1.8, 1.11, 3.3, 5.1, 5.2, 5.8, 12.7, 20.2, 20.5, 20.15, 21.3, 21.6, 22.16
- Binomial 1.12, 9.1, 9.3–9.5, 9.8–9.10, 9.12–9.15, 10.16, 10.20, 15.9, 16.9, 16.12, 17.6, 17.7, 17.8, 17.11, 17.12

Binomial trees 9.3, 9.8, 9.13, 15.9, 17.7,

Bivariate normal 16.11,

- Black 1.3, 1.12, 6.8, 8.13, 9.13, 9.15, 10.1, 10.4, 10.5, 10.6, 10.7, 10.11–10.21, 12.1, 12.3, 12.11, 12.12, 12.14, 12.15, 12.16, 13.9, 13.10, 14.1, 14.3, 14.6, 14.8, 15.1–15.3, 15.7–15.9, 16.1, 16.6, 16.8, 16.9–16.12, 17.2, 17.4, 17.5, 17.6, 17.12, 19.1, 19.3, 19.10, 19.15, 19.16, 20.12, 20.13, 22.5,
- Bond 1.3, 1.6, 1.10, 1.11, 2.5, 2.6, 2.7, 2.12–2.14, 3.4– 3.6, 3.11, 4.2–4.5, 4.8, 4.9, 4.15, 4.16, 6.9, 7.8, 8.6, 8.8, 8.9, 8.11, 8.14, 12.1, 12.4, 13.4, 15.5, 16.13, 17.5–17.12, 17.14, 17.15, 17.17, 17.22–17.26, 18.1, 18.4, 18.5, 18.6–18.12, 18.14, 18.16–18.18, 18.20, 18.21, 18.24, 19.1–19.5, 19.7–19.12, 19.14, 19.16, 19.17, 20.4, 20.8, 20.10, 20.11, 20.14, 20.17, 22.6, 22.9, 22.15
- Bootstrapping 18.12, 18.13
- Bull spread 11.8–11.10, 15.5, 16.3,
- Butterfly 13.4–13.12,
- Calendar spreads 13.9, 13.11, 13.12
- Call Option 1.5, 5.5, 6.1, 6.3, 6.4, 6.6–6.13, 7.3, 7.5, 8.8, 8.10, 8.11, 8.14, 9.2, 9.3, 9.5–9.7, 9.9, 9.10, 9.15, 10.4–10.9, 10.13, 10.14, 10.19, 10.21, 11.1, 11.5, 11.6, 11.8, 11.10, 12.1, 12.2, 12.4, 12.7, 12.8–12.10, 12.17, 13.1, 13.10, 15.1, 15.6, 16.1, 16.2, 16.4, 16.6, 16.7, 16.11–16.13, 17.1, 17.2, 17.4, 17.6, 17.12, 17.14, 17.15, 17.17, 19.4, 19.7, 19.10, 19.16, 20.17, 21.13, 22.3, 22.10, 22.17
- Callability 19.1, 19.7, 19.9, 19.16
- Cap 19.1–19.6, 19.10, 19.13, 19.15–19.17, 20.4, 20.17–20.19, 21.21
- Caplet 19.1–19.3, 19.5, 19.6, 19.10, 19.14–19.16, 20.17, 20.18
- Cash settlement 2.1, 2.4, 2.6, 2.7, 2.8, 2.11–2.15, 7.3, 7.9, 20.3
- Certainty equivalent 8.14
- Cheapest to Deliver (CTD) 4.9, 4.15,
- Chicago Board of Trade (CBOT) 2.10, 2.14,
- Chicago Mercantile Exchange (CME) 2.6, 2.14, 4.11, 7.9
- Collateral 2.15, 3.4, 5.4, 21.2
- Commodities 1.3, 1.4, 1.6, 1.7, 1.13, 2.1, 2.2, 2.9, 2.10, 3.1, 3.10, 3.12, 7.1, 7.6, 7.7, 10.1, 10.7, 10.13, 10.20
- Commodity derivatives 1.2, 1.7, 2.5, 2.7, 2.11, 2.14, 18.1

I.2 | Index

Commodity Futures Trading Commission (CFTC) 2.11 Conditional expectation 10.4, 10.5, 10.17, 22.3 Conditional Value at Risk (CVaR), 22.3 Contango 3.11, 3.12, 21.12–21.14, 21.23, Contingent 15.9, 17.5, Continuous compounding 3.7, 3.8, Continuous trading 8.14, 16.5, Contract design 2.4, 2.7, 2.14, 7.1, 7.9 Contract specification 2.7, 2.15, Convenience yield 3.9–3.13, 4.14, 4.15, 10.13, 10.20 Core competence 5.2, 21.4, 21.7 Correlation 4.5, 4.14, 4.16, 19.13, 19.17, 20.2, 20.14, 22.9, 22.14, 22.15, 22.20 Cost of Carry 1.11, 3.1, 3.6–3.8, 3.12, 4.14, 8.3, 8.4, 10.12, Coupon 2.12, 2.13, 3.4–3.6, 3.9, 3.11, 4.2, 4.8, 4.9, 8.8, 16.13, 17.6, 17.7, 17.9, 17.10, 17.11, 17.14, 17.15, 17.17, 18.1–18.3, 18.6–18.8, 18.10–18.12, 18.15-18.19, 19.2-19.5, 19.7-19.11, 19.14-19.17, 20.8, 22.9 Covariance 4.5, 22.13, 22.14 Covered call 11.9, 11.10, Credit risk 1.6, 1.7, 2.2, 2.3, 2.14, 5.6, 7.1, 17.6, 17.13, 20.10, 20.14, 20.16, 21.15 Creditworthiness 1.7, 2.2, 18.8, 18.9, 20.10, 20.11, 21.2, 21.15 Crisis 3.6, 5.9, 18.22, 18.23, 21.11, 21.22, 22.15 Currency 1.3, 1.4, 1.6, 1.8, 1.12, 2.5, 2.6, 2.9, 2.10, 2.12, 3.11, 4.7, 4.10, 4.11, 5.2, 5.6, 5.7, 5.9, 5.10, 5.12, 5.14, 5.16, 5.19, 5.21, 7.2, 7.6, 7.8, 7.9, 10.10, 10.19, 10.20, 11.5, 11.10, 12.14, 12.15, 12.16, 15.9,

- 16.10, 17.14, 17.15, 17.17, 17.23, 17.28, 18.1–18.4, 18.15, 18.17–18.20, 18.22, 18.23, 18.29–18.31, 20.3,
- 20.4, 20.22, 21.1, 21.2, 21.22, 22.4, 22.6, 22.7, 22.16, 22.17, 22.19
- Curvature 15.7, 22.16,
- Customized contracts 5.3
- Default 1.6, 2.2, 2.3, 4.2, 5.21, 5.23, 5.26, 5.27, 8.14, 17.10, 17.11, 17.24, 18.7, 21.8, 21.15, 21.23, 22.15, 22.16
- Deliverable 2.6, 2.7, 2.13,
- Delivery 1.1, 2.1–2.7, 2.12–2.14, 3.6, 3.9, 3.10, 3.12, 5.3, 5.12, 7.1, 7.2, 7.3, 7.7, 7.9, 20.3, 21.13, 21.15, 21.29, 22.17,
- Delta 7.5, 9.3, 9.7, 9.9, 9.10, 9.15, 12.1–12.8, 12.11, 12.12, 12.14–12.17, 13.1–13.3, 13.6, 13.7, 13.10–13.12, 15.6, 15.8, 15.10–16.6, 16.11, 20.12, 20.13, 22.7–22.13, 22.15, 22.17–22.19,
- Deregulation 1.4
- Derivative 1.1–1.4, 1.6–2.1, 2.8, 2.14, 2.15, 3.1, 3.3,

4.5, 5.3, 5.7, 5.13, 7.5, 7.9, 8.1, 8.2, 8.7, 8.8, 8.10, 8.11, 8.13, 9.8, 16.1, 17.7, 20.1, 20.2, 20.3, 20.4, 20.6, 20.7, 20.10, 20.11, 20.14, 20.15, 20.21, 20.22, 21.2, 21.7, 21.8, 21.9, 22.1, 22.3, 22.16, 22.17, 22.18 Derivative disasters 1.1, 1.2, 1.8–1.11, 1.13, 21.8, 22.17 Barings 1.8, 1.9, 1.13, 13.12 Hamanaka 1.9, 1.14 LTCM 1.11, 22.18 Metallgesselschaft 1.9, 1.13, 1.14, 4.15, 4.16, 21.8 Proctor and Gamble 1.9, 1.10, 21.8, Sumitomo 1.9, 1.13, 1.14, Dilution 17.3–17.5, 17.12, 17.24 Diminishing marginal utility 8.2–8.5, 8.13 Disclosures 21.2 Discontinuity 16.2, 16.3, 16.5, 16.6 Discount 2.7, 2.12, 3.1, 4.2, 4.8, 4.9, 5.9, 8.1, 8.2, 8.4, 8.7, 8.13, 8.14, 9.1, 10.5, 17.6, 17.7, 17.10, 17.12, 18.8, 18.9, 19.4, 19.11, 19.16, 20.8 Discounting 3.1, 4.11, 8.1, 8.2, 8.4, 8.5, 8.7, 8.9, 8.13, 9.1, 9.2, 9.6, 9.7, 10.4, 10.7, 12.4, 12.14, 17.6, 17.7, 17.9, 17.10, 18.7, 18.9, 18.18, 19.3, 19.10, 20.12, Diversification 5.3, 22.16, 22.18, Dividend 3.2–3.5, 3.8, 3.9, 3.12, 3.13, 5.15, 6.7, 6.8, 6.11, 10.1, 10.5–10.16, 10.18–10.21, 12.1, 12.3, 12.4, 12.8, 12.9, 12.12, 12.14–12.17, 14.10, 15.9, 16.1, 16.2, 16.6, 16.8, 16.11, 16.13, 17.7, 17.8, 17.15, 17.28, 18.28 Documentation 5.13, 5.14, 20.1, 20.14, 20.15 Doubling strategy 1.8, 1.9, 1.14, Drift 16.6, 16.8, 16.9, Duration 4.2, 4.8, 4.9, 4.15, 4.16, 5.31, 18.4, 18.8, 18.24, 22.9, 22.16, Dynamic hedging 12.7, 12.16, 16.12, 22.16 Early exercise 9.13, 10.14, 10.15, 12.9 Enterprise wide risk 22.1 Equity 2.1, 2.5, 2.9–2.11, 2.14, 3.1, 3.6, 4.10, 5.12, 5.15, 5.16, 5.18, 5.26, 7.6-7.9, 15.4, 16.11, 17.12, 17.14, 17.16, 17.17, 17.29, 18.1, 18.22, 20.19, 20.21, 21.6, 21.7, 21.8, 21.11, 21.19 Eurex 2.12, 2.14, 7.7, 7.8 Euro 2.12, 4.10–4.14, 7.2, 7.8, 17.14, 17.16, 17.24, 21.11, 22.4, 22.5, 22.7, Eurodollar 1.3, 2.5, 2.13, 2.14, 4.15, 7.8 Euronext 2.8, 2.11, 7.7, 7.8 Exercise price 6.1, 6.7–6.13, 7.2, 9.2, 9.8, 9.9, 9.12, 10.4-10.8, 10.11-10.16, 10.19, 10.21, 12.1, 12.3, 12.14, 12.15, 12.17, 13.1, 13.2, 14.1, 15.1, 15.9–16.2,

16.11, 17.1, 17.2, 17.5, 17.6, 17.7, 17.24 Exotic options 1.12, 16.1, 16.3, 16.5, 16.11, 16.12

Index | I.3

Asian options 1.12, 16.9, 16.10 Barrier options 16.5, 16.6, 16.8, 16.9, 16.13, Binary option 8.14, Chooser options 16.11, Compound options, 16.11,

- Knock in options, 16.9,
- Knock out options 16.5, 16.9
- Expectation 3.2, 8.5, 8.6, 8.13, 9.10, 10.1, 10.3–10.5, 10.13, 10.17, 11.1, 12.4, 12.11, 16.2, 16.3, 17.4, 19.3, 19.13–19.16, 22.3, 22.7, 22.19,
- Expected shortfall 22.3, 22.19, 22.20
- Exponentially weighted moving average 14.4–14.6, 14.8, 14.10, 22.14
- Exposure 1.3, 2.4, 4.1, 4.3, 4.5–4.7, 4.11–4.15, 5.1, 5.3, 5.4, 5.6, 5.7, 5.14, 5.20, 5.22, 7.4, 9.3, 11.1, 12.1, 13.1, 13.6, 15.8, 16.13, 18.4, 18.17, 18.23, 20.1, 20.2, 20.4, 20.9, 20.14, 21.1, 21.2, 21.7, 21.8, 21.12, 21.29,
- Fair value 3.12, 3.13, 5.13, 8.14, 11.10, 13.4, 17.4, 17.6, 17.7, 20.1–20.15, 20.18, 20.19, 20.21
- Far month 13.9, 13.10, 13.11, 13.12, 15.8,
- Fat tails 22.19,
- Financial distress 5.6, 21.1–21.8,
- Financial institutions 1.2–1.4, 1.8, 1.11, 1.12, 5.4, 7.8, 18.23, 18.29, 20.3, 21.6, 21.7, 21.12, 22.1, 22.3, 22.9, 22.16–22.20
- Financial intermediaries 2.2,
- Floor, 19.1, 19.2, 19.4, 19.5, 19.16, 20.4, 21.13
- Floorlet 19.4,
- Floortion 1.6,
- Forecast 2.4, 4.6, 14.6–14.8, 15.5, 15.10, 20.9, 21.4, 21.27
- Forward contracts 1.2, 1.4, 1.6, 1.7, 1.13, 2.1, 2.2, 2.4, 2.14, 2.15, 3.1, 3.3, 4.1, 4.10, 4.14, 4.16, 5.1, 5.3, 5.4, 5.5, 5.13, 6.2, 6.5, 18.1, 18.3, 18.4, 18.16, 18.19, 18.20, 19.3, 20.3, 20.5, 21.12, 21.14, 21.15, 21.26, 21.27,
- Forward price 2.2, 2.4, 3.1–3.12, 4.10, 4.12–4.14, 4.16, 6.2, 6.11, 8.4, 12.2–12.4, 13.1, 13.3, 15.6, 19.15, 20.3, 21.12, 21.14
- Forward Rate Agreement (FRA) 18.16,
- Forward rates 18.10, 18.13, 18.14, 18.16, 18.17, 18.19, 18.20, 19.6, 19.11, 19.13, 19.16, 19.17, 20.18,
- Friedman Milton 1.7, 2.2
- Futures 1.2–1.4, 1.6–1.8, 1.10, 1.11, 1.13–3.1, 3.3, 3.10–3.13, 4.1–4.3, 4.5–5.1, 5.4, 5.5, 5.7, 7.1–7.3– 7.9, 8.3, 8.13, 10.1, 10.7, 10.11–10.13, 10.19–10.21, 11.1, 11.2, 11.10, 12.3–12.5, 12.14, 12.16, 12.17, 13.2, 13.4, 13.12, 14.10, 15.8, 15.10, 18.4, 19.3, 20.3, 20.6, 20.7, 21.1, 21.2, 21.30, 22.3, 22.8

- Futures markets 1.2,-1.4, 1.11, 2.1, 2.6, 2.8-2.11, 2.14, 2.15, 4.15, 5.4, 7.2, 7.4, 8.13, 11.1, 21.1
- Gamma 9.15, 12.5–12.9, 12.11–12.13, 12.15–12.17, 13.1, 13.6–13.9, 13.11, 13.12, 15.8–15.10, 16.3, 16.4, 16.11, 20.12, 20.13, 22.1, 22.8, 22.10–22.13, 22.17, 22.19
 - Rho 12.14–12.17,
 - Theta 12.8–12.12, 12.15–12.17, 13.6, 13.9–13.11, 20.12, 20.13, 22.1
 - Vega 9.15, 12.12–12.14, 12.16, 12.17, 13.1, 13.6– 13.9, 13.11, 13.12, 15.8–15.10, 16.5, 16.6, 16.11, 20.12, 20.13, 22.1, 22.8, 22.17, 22.19
- Garch 14.5-14.10, 15.5, 15.9,
- Geometric mean 10.2, 10.3, 16.10
- Greek 12.1, 12.8, 12.12
 - Delta 7.5, 9.3, 9.7, 9.9, 9.10, 9.15, 12.1–12.8, 12.11, 12.12, 12.14–12.17, 13.1–13.3, 13.6, 13.7, 13.10–13.12, 15.6, 15.8, 15.10–16.6, 16.11, 20.12, 20.13, 22.7–22.13, 22.15, 22.17– 22.19
- Hamanaka 1.9, 1.14
- Hedge 1.3, 1.4, 1.7, 1.8, 1.10–1.12, 2.12, 4.1–5.8, 5.10, 5.13, 5.14, 7.4, 7.6, 11.10, 12.7, 13.1, 15.4, 15.8, 16.1, 16.3–16.5, 16.7, 16.8, 16.10–16.13, 18.3, 18.4, 18.17, 18.19, 18.24, 19.1, 19.17, 20.1–20.16, 20.18, 20.19, 20.21–21.2, 21.4, 21.7, 21.8, 21.12–21.16, 21.27–21.30
- Hedge effectiveness 4.5, 4.12, 4.14, 4.15, 5.4, 20.1, 20.14, 20.15,
- Hedged 1.1, 1.10, 4.1–4.7, 4.14, 5.2, 5.7, 5.13, 5.14, 7.4, 11.4, 11.5, 11.7, 11.9, 11.10, 12.1, 12.5, 12.8, 12.11, 15.8, 16.8, 18.24, 20.1, 20.5–20.12, 20.14, 20.21, 20.22, 21.2, 21.6, 21.16, 21.23, 21.26, 21.27, 21.29
- Hedger 5.1, 5.3–5.5, 11.1, 20.10
- Held for Trading (HFT) 20.10
- Held to Maturity (HTM) 20.10
- Illiquidity 18.6, 22.16, 22.18
- Increment 2.6, 7.2
- Initial margins 2.3, 7.1
- Innovation 1.3, 1.4, 1.14
- Insurance 1.14, 3.9, 5.2, 5.3, 5.6, 5.24, 5.33, 6.1, 6.6, 6.7, 6.8, 6.12, 10.13–10.16, 20.13, 21.2, 21.7, 22.10 Intercontinental Exchange (ICE) 2.11, 7.7
- Interest rate 1.3, 1.4, 1.6, 1.7, 1.12, 2.5, 2.9–2.14, 3.1, 3.3–3.6, 3.8–3.13, 4.2, 4.8, 4.9, 4.11, 4.14–4.16, 5.5, 6.7, 6.8, 6.12, 6.13, 7.5, 7.6, 7.8, 8.14, 9.14, 10.10–

I.4 | Index

10.12, 10.14, 10.19, 10.21, 11.10, 12.1, 12.4, 12.11-12.14, 12.16, 12.17, 14.1, 14.8, 15.9, 16.6, 16.8, 16.13, 17.5, 17.11, 17.12, 17.15, 18.1-18.9, 18.14-18.20, 18.22, 18.23, 19.1, 19.3, 19.4, 19.8, 19.9, 19.14, 19.15, 19.17, 20.2-20.4, 20.8-20.10, 20.14-20.21, 21.12 International Accounting Standards 20.1, 20.2, 20.15 Intrinsic value 6.9–6.11, 7.4, 9.8, 9.10, 9.15, 10.16, 10.20, 12.12, 13.9, 13.10, 16.5, 16.6, 16.12, 20.12-20.14, 20.17-20.19, Inventory 3.9, 17.19, 20.6, 20.7 Inverse floater 20.8, Investment 1.7–1.9, 3.1–3.4, 3.6, 3.13, 4.2, 4.7, 4.13, 5.1, 5.17, 5.18, 6.7, 6.9, 6.11, 8.1–8.5, 8.14, 9.10, 13.4, 13.6, 17.14, 17.16, 17.23, 17.27, 17.28, 18.28, 19.1, 20.3, 20.10, 20.16, 21.3, 21.5, 21.14, 21.22, 21.23, 22.3, 22.16, 22.17, Investor 1.1, 2.3, 2.4, 2.11, 3.4, 3.10, 4.7, 4.9-4.11, 4.13, 8.5, 11.1–11.4, 11.8, 11.9, 13.1, 13.2, 13.4, 13.6, 13.11, 16.3, 16.4, 16.7, 16.9, 17.23, 18.22, 19.2, 21.1, 21.12, 21.30, 22.4 Invoice price 2.6, 2.13 LEAPS 7.7, Leeson 1.8, 1.9, 1.13, 1.14, 13.12, 22.17 Leverage 1.9, 1.11, 1.13, 20.3, 21.9 Libor 2.13, 2.14, 4.15, 16.13, 18.6, 18.8, 18.18, 18.19, 18.23, 18.24, 18.30–19.2, 19.4, 19.15, 20.17 Libor Market Model (LMM) 19.1, 19.2, 19.4, 19.15, Liffe 2.11, 7.7, 7.8 Liquidity 1.1, 1.2, 1.4, 1.9–1.11, 1.13, 1.14, 2.14, 4.10, 4.14, 5.7, 7.7, 18.4, 20.16, 21.6, 21.8, 21.29–21.31, 22.1, 22.17, 22.18, 22.20, Logarithm 3.8, 9.13–9.15, 10.1–10.3, 10.18, 19.15, 20.14 Logarithmic return 10.1–10.3, 10.5, 10.18, 14.1, 14.4, 16.1, 16.2 Lognormal distribution 19.15 LTCM 1.11, 22.18 Manipulation 1.2, 2.7, 16.9, 16.10, 21.22 Margins 2.3, 2.8, 2.14, 2.15, 4.10, 5.3, 5.4, 5.12, 7.1, 7.4, 7.9, 17.24, 18.23 Mark to market 2.3, 2.8, 2.12, 2.14, 2.15, 4.10–4.13, 4.15, 5.1, 5.3, 5.4, 5.7, 5.10, 5.14, 7.1, 7.4, 20.7, 21.7,

21.8, 22.17, 22.18 Maturity 1.4–1.6, 1.10, 2.2–2.6, 3.1, 3.3, 3.5–3.7, 3.10–3.12, 4.1, 4.2, 4.7–4.15, 5.4, 5.5, 6.1–6.13, 7.1, 7.2, 7.7, 8.3, 8.4, 8.10, 8.13, 8.14, 9.8, 9.10, 9.14, 9.15, 10.4–10.6, 10.11–10.16, 10.19, 10.20, 11.1, 11.2, 11.10, 12.1–12.10, 12.12–12.15, 13.1, 13.4,

13.6, 13.10, 13.12, 14.1, 14.3, 14.6, 14.8, 14.10, 15.1, 15.5, 15.6, 15.8, 15.9, 16.1–16.3, 16.5, 16.7–16.13, 17.2, 17.3, 17.5–17.7, 17.9–17.12, 17.14, 17.15, 17.17, 17.23, 17.24, 17.27, 18.2, 18.3–18.9, 18.11, 18.12, 18.14, 18.15, 18.18, 18.19, 18.24, 18.30, 18.31, 19.3, 19.5, 19.7–19.10, 19.15, 20.10, 20.12, 20.13, 20.18, 21.7, 21.8, 21.15, 21.16, 21.26 Mean reversion 14.6, 14.7, 15.5, 15.9, Merton Robert Metallgesselschaft 1.9, 1.13, 1.14, 4.15, 4.16, 21.8 MIBOR 18.2–18.8, 18.19, 19.2–19.4, 19.14, 19.16, 20.8 Miller Merton 1.14 Mismatch, 4.15, 5.12, 20.13, Mispricing 22.18 Mode 22.10 Model 1.10, 1.12, 3.1, 3.6, 3.7, 4.14, 5.12, 6.8, 7.4, 8.2, 8.3, 8.4, 8.6, 8.8-8.10, 8.13, 9.1, 9.3, 9.5-9.7, 9.10, 9.13, 9.15, 10.1, 10.2, 10.6, 10.7, 10.9, 10.11, 10.12, 10.16, 10.18, 10.20, 12.1, 12.3, 12.12, 12.16, 14.6-14.10, 15.5, 15.7–15.9, 16.6, 16.11, 17.1, 17.3–17.5, 17.7, 17.12, 17.14, 19.1–19.4, 19.10, 19.11, 19.15, 19.17, 21.10, 22.1, 22.5, 22.14, 22.15, 22.18, 22.19 Monte Carlo 22.12-22.15, 22.19 Multi Commodity Exchange (MCX) 2.5, 2.11 National Commodity Derivative Exchange (NCDEX) 2.15 Near month 2.11, 2.12, 13.9–13.12, 15.8, Net settlement 20.3 New York Board of Trade (NYBOT) 2.11, 7.7 New York Metal Exchange (Nymex) 2.11, 7.7 Normal distribution 9.13, 9.15, 10.1–10.4, 10.17, 12.4, 14.8, 15.1, 15.2, 15.8–15.10, 16.6, 16.10–16.12, 17.4, 19.3, 19.4, 19.15, 22.6–22.13, 22.15, 22.19 Notional value 2.9, 2.13, 2.14, 4.8, 4.9, 5.13, 7.5, 7.6, 7.9, 13.12, 18.1, 20.3, 20.16, 22.17 Novation 2.3, 2.4, 4.10, 5.4, 7.1 One Chicago 2.11 American options 6.1, 6.7, 7.4, 7.7, 9.10, 10.1, 10.13, 10.16, 10.20, 16.12, 17.11, 19.9 Basket of options 19.10, 19.13, Call option 1.5, 5.5, 6.1, 6.3, 6.4, 6.6–6.13, 7.3, 7.5, 8.8, 8.10, 8.11, 8.14, 9.2, 9.3, 9.5–9.9, 9.10,

7.5, 8.8, 8.10, 8.11, 8.14, 9.2, 9.3, 9.3–9.9, 9.10,
9.15, 10.4–10.9, 10.13, 10.14, 10.19, 10.21,
11.1, 11.5, 11.6, 11.8, 11.10, 12.1, 12.2, 12.4,
12.7–12.10, 12.17, 13.1, 13.10, 15.1, 15.6, 16.1,
16.2, 16.4, 16.6, 16.7, 16.11–17.1, 17.2, 17.4,
17.6, 17.12, 17.14, 17.15, 17.17, 19.4, 19.7,
19.10, 19.16, 20.17, 21.13, 22.3, 22.10, 22.17

Index | I.5

```
Determinants of option price 1.11, 6.7,
```

- Exercise price 6.1, 6.7–6.13, 7.2, 9.2, 9.8, 9.9, 9.12, 10.4–10.8, 10.11–10.16, 10.19, 10.21, 12.1, 12.3, 12.14, 12.15, 12.17, 13.1, 13.2, 14.1, 15.1, 15.9–16.2, 16.11, 17.1, 17.2, 17.5, 17.6, 17.7, 17.24,
- Exotic options 1.12, 16.1, 16.3, 16.5, 16.11, 16.12,
- Expiration 2.6, 7.2, 7.3
- Option markets 7.1, 7.6, 7.7, 11.10, 12.17, 13.12, 19.5, 19.13, 19.15
- Option premium 6.4–6.7, 6.8, 7.2, 7.5, 7.9, 11.1–11.6, 11.10, 12.8, 13.1, 13.2, 13.6, 16.6, 20.3, 20.13, 21.13
- Option Strategies 11.1, 11.6, 13.1, 13.6–13.11
- Put call parity 6.11, 6.13, 9.15, 10.5, 11.9, 12.1, 12.4, 13.4, 13.11, 16.1, 16.5, 16.11, 19.5, 19.13,
- Put option 6.1, 6.5–6.8, 6.10–6.13, 7.4, 8.14, 9.8,– 9.12, 9.15, 10.5, 10.6, 10.8, 10.15, 10.16, 10.20, 10.21, 11.3, 11.4, 11.6, 11.9, 12.2–12.5, 12.8–12.10, 12.12–12.14, 12.17, 13.1, 16.1, 16.2, 16.10–16.12, 17.7, 17.14, 17.15–17.17, 17.24, 19.4, 19.7, 19.16, 20.12, 20.13, 21.13
 - Strike price 1.5, 5.5, 6.1, 6.3, 6.4, 6.13, 7.2, 7.3, 7.8, 8.8, 8.10, 11.1, 11.2, 11.5, 11.6, 11.9, 12.2, 12.3, 12.6, 12.11, 12.13, 13.6, 13.7, 13.10– 13.12, 14.8, 15.2–15.4, 15.6, 15.7, 16.12, 16.13, 17.5, 21.13
 - Option strategies 11.1, 11.6, 13.1, 13.6–13.9, 13.11
 - Bear spread 11.8-11.10
 - Bull spread 11.8-11.10, 15.5, 16.3,
 - Butterfly 13.4–13.12,
 - Calendar spreads 13.9, 13.11, 13.12
 - Covered call 11.9, 11.10
 - Range forward 11.6, 11.7, 11.9, 11.10, 21.13,
 - Risk reversal 15.5,
 - Straddle 13.1–13.4, 13.6–13.9, 13.11,
 - Strangle 13.3, 13.4, 13.6–13.9, 13.11, 13.12, 15.5, Optional 8.2–8.4, 8.6, 8.8, 8.13, 9.8, 9.13, 10.17,
 - 17.10, 17.11, 19.14, 19.15, Optionality 1.4, 1.6
 - Optionality 1.4, 1.6
- Payoff diagram 1.4, 1.5, 6.5, 6.11, 13.5, 13.6
- Profit diagram 6.4–6.6, 11.11, 13.1–13.5, 13.9–13.12, Peso 21.11, 22.7
- Physical settlement 2.1, 2.6, 2.7, 2.14, 7.2, 7.3, 7.9
- Portfolio 1.7, 3.4, 4.7, 4.9, 4.10, 4.13, 4.15, 4.16, 5.18, 7.1, 8.3, 8.4–8.8, 8.10, 8.11–8.13, 9.3, 9.7, 9.10, 9.15, 10.9, 11.11, 15.8, 15.9, 16.4, 17.16, 18.4, 18.8, 18.21, 18.22, 19.14, 20.2, 20.9, 20.16, 20.21, 21.1, 21.27, 21.29, 22.1, 22.2, 22.4–22.14, 22.16, 22.17, 22.19

Precious metals 1.4, 1.6, 2.11, 7.7

- Present value 2.13, 3.1–3.3, 3.7, 3.10–3.12, 4.2, 4.3, 4.5, 4.8–4.13, 4.15, 6.9–6.11, 10.7, 10.19, 13.4, 13.5, 17.11, 18.7, 18.9–18.12, 18.17, 18.19, 19.11, 19.14, 19.16, 20.8, 20.12, 20.15, 20.17, 20.18
- Principal 1.11, 2.2, 2.8, 2.9, 2.12, 3.4, 3.7–3.9, 5.3,
 5.11, 6.1, 6.9, 6.11, 7.5, 7.7, 7.9, 17.16, 17.17, 17.25,
 17.26, 17.28, 18.1, 18.2, 18.5–18.7, 18.11, 18.17,
 18.19, 18.20, 18.30, 18.31, 19.12, 20.1, 20.4, 20.9,
 20.10, 20.14–20.19, 20.22, 21.1, 21.8, 21.22,
- Probability 3.11, 6.7, 8.2–8.7, 8.9, 8.14, 9.1, 9.5, 9.8, 9.9, 9.11–9.14, 10.2–10.5, 12.14, 13.5, 15.1, 15.3– 15.5, 15.7, 15.9, 16.1, 16.2, 16.5, 16.7, 16.8, 17.4, 17.10, 17.11, 17.24, 22.2, 22.7–22.19
- Proctor and Gamble 1.9, 1.10, 21.8
- Put call parity 6.11, 6.13, 9.15, 10.5, 11.9, 12.1, 12.4, 13.4, 13.11, 16.1, 16.5, 16.11, 19.5, 19.13
- Put Option 6.1, 6.5–6.8, 6.10–6.13, 7.4, 8.14, 9.8–9.12, 9.15, 10.5, 10.6, 10.8, 10.15, 10.16, 10.20, 10.21, 11.3, 11.4, 11.6, 11.9, 12.2–12.5, 12.8–12.10, 12.12– 12.14, 12.17, 13.1, 16.1, 16.2, 16.10–16.12, 17.7, 17.14–17.17, 17.24, 19.4, 19.7, 19.16, 20.12, 20.13, 21.13
- r-square 4.4-4.15, 20.14
- Range forward 11.6, 11.7, 11.9, 11.10, 21.13
- Redemption 2.12, 4.8, 4.9, 4.13, 17.5–17.8, 17.14, 17.15, 17.17, 17.25–17.27, 18.3, 18.7, 18.10, 18.12, 18.19, 18.27–18.29, 19.7, 19.8, 19.14
- Regulator 2.11, 21.30, 22.2, 22.15,
- Re-insurance 5.3
- Replication 8.10, 8.12, 9.7, 9.8, 9.10, 9.15, 16.4
- Reset 18.4, 18.7, 18.8, 18.18, 18.20, 19.3, 19.9, 20.8,
- Residual 2.6, 3.10, 3.12, 4.5, 5.26, 21.1, 21.2, 21.4,
- Revaluation 20.10, 21.31, 22.12, 22.13, 22.19
- Reverse cash and carry 3.1, 3.5–3.7, 3.12, 3.13
- Rho 12.14-12.17
- Risk 1.1, 1.3, 1.4, 1.6–1.14, 2.2, 2.3, 2.6, 2.12, 2.14, 3.1–3.6, 3.10–3.12, 4.1, 4.15, 5.1–5.8, 5.10, 5.12–5.14, 5.22, 6.6, 6.7, 6.11, 6.13, 7.1, 7.4, 7.6, 8.1–9.3, 9.5– 10.11, 10.13–10.16, 10.18–10.21, 11.1, 11.2, 11.5, 11.9–11.11, 12.1–12.5, 12.7, 12.11, 12.12, 12.14– 12.17, 13.1, 13.5, 13.6, 13.8, 13.9, 13.11–14.1, 14.8, 14.10–15.5, 15.7–15.9, 16.1–16.3, 16.6, 16.8–16.10, 16.12, 16.13, 17.3, 17.4, 17.6–17.8, 17.10–17.13, 17.15, 17.23, 18.3, 18.4, 18.7–18.10, 18.13, 18.17, 18.19, 18.20, 19.1, 19.3, 19.4, 19.6, 19.14–19.16, 20.2, 20.3, 20.6–20.12, 20.14, 20.15, 20.16, 20.19, 21.1– 21.9, 21.12–21.16, 21.22, 21.25, 21.26, 22.1–22.3, 22.5, 22.7, 22.8, 22.14–22.16, 22.18–22.20

I.6 | Index

Enterprise wide risk 22.1, Exposure 1.3, 2.4, 4.1, 4.3, 4.5–4.7, 4.11, 4.12, 4.14, 4.15, 5.1, 5.3, 5.4, 5.6, 5.7, 5.14, 5.20, 5.22, 7.4, 9.3, 11.1, 12.1, 13.1, 13.6, 15.8, 16.13, 18.4, 18.17, 18.23, 20.1, 20.2, 20.4, 20.9, 20.14, 21.1, 21.2, 21.7, 21.8, 21.12, 21.29, Risk Adjusted Discount Rate (RADR) 8.1, Risk aversion 8.2–8.5, 8.7, 8.13 Risk management 1.3, 1.4, 1.11–1.13, 4.1, 5.5–5.7, 5.10, 5.13, 5.14, 5.22, 7.4, 8.1, 12.1, 18.20, 19.6, 20.9, 20.14, 20.16, 21.1–21.9, 21.25, 21.26, 22.1, 22.15, 22.16, 22.18, 22.19, 22.20 Risk Adjusted Discount Rate (RADR) 8.1, Risk aversion 8.2–8.5, 8.7, 8.13 Risk management 1.3, 1.4, 1.11–1.13, 4.1, 5.5–5.7, 5.10, 5.13, 5.14, 5.22, 7.4, 8.1, 12.1, 18.20, 19.6, 20.9, 20.14, 20.16, 21.1-21.9, 21.25, 21.26, 22.1, 22.15, 22.16, 22.18-22.20 Risk neutral distribution 10.2-10.5, 10.18, 13.6, 15.1-15.5, 15.7, 15.8, 16.6, 16.9 Risk reversal 15.5, RiskMetrics 22.9, 22.14, 22.20 Rogue traders 1.9, 1.11, 1.13 Sample 1.14.1, 1.14.3–1.14.5, 14.8, 22.3–22.7, 22.12– 22.15, 22.19 Scenario 4.12, 8.14, 15.8, 17.6, 18.7, 21.16, 21.29, 22.4-22.6, 22.12, 22.13, 22.14, 22.16, 22.19 Seasonality 10.9, 15.5 Security 1.1, 1.13, 5.21, 5.22, 5.28, 8.2, 8.3, 8.6, 21.2 Sensitivity 4.8, 4.10, 5.22, 9.15, 12.16, 20.8, 21.28, 21.29 Settlement 1.11, 2.1, 2.4, 2.6, 2.7–2.9, 2.11–2.15, 7.1– 7.5, 7.9, 19.5, 19.6, 20.3, 22.17 Shareholder 21.1, 21.8, 21.10, 21.11, 21.19, 21.30, 22.2, 22.3 Short selling 3.6, 3.12, 11.3, 21.22 Simulation 16.12, 22.3–22.7, 22.12, 22.13, 22.15, 22.19 Skewness 12.4, 15.8, 22.11, 22.12 Speculation 5.2, 5.7, 17.23, 18.4, 21.4, 21.7, 21.9, Spot 1.13, 2.1, 2.2–2.4, 2.7, 3.1–3.7, 3.9–3.13, 4.12, – 4.15, 5.9, 5.13, 6.2, 7.9, 8.3, 8.4, 10.11–10.13, 10.20, 10.21, 11.1, 11.6, 12.3-12.5, 12.16, 13.1, 15.8, 17.17, 18.32, 19.5, 20.11–20.13, 21.12–21.15, 21.17, 21.23, 21.26-21.28 Spread 1.4, 3.7, 4.2, 4.3, 5.3, 11.6, 11.8–11.11, 13.4,

Spread 1.4, 3.7, 4.2, 4.3, 5.3, 11.6, 11.8–11.11, 13.4, 13.7, 13.9, 13.10, 13.11, 13.12, 15.5, 16.3, 16.4, 16.5, 17.12, 18.30, 18.31, 20.21, 21.1, 21.10, 21.11–21.13, 21.31, 22.9, Square root scaling 14.1, 16.10

Stakeholder 21.3

Standardization 2.4, 2.6, 4.10, 5.4, 7.1

- States of the world 8.3, 8.6, 8.13,
- Statistical 3.8, 4.3, 4.5, 4.10, 14.6, -14.8, 15.7, 20.14, 22.12, 22.15,
- Stochastic 8.14, 14.9, 15.7, 15.9
- Stock 1.1, 1.2, 1.4, 1.6–1.8, 1.13, 2.1, 2.5–2.8, 2.11,
 2.12, 2.14, 2.15, 3.1–3.6, 3.9, 3.12, 4.7, 4.8, 4.10,
 4.15, 5.12, 5.14–5.16, 5.20, 5.24, 6.2, 6.8–6.13, 7.2,
 7.4–7.7, 7.9, 8.6–8.13, 9.1–10.16, 10.18, 10.19,
 10.21, 11.2–11.4, 11.8–11.11, 12.1–12.12, 12.14,
 12.15, 12.17, 13.1, 13.2, 15.9, 16.1, 16.2, 16.6, 16.13,
 17.1, 17.4–17.16, 17.25, 18.22, 18.23, 20.21, 21.2,
 21.10, 21.11, 21.21, 21.22, 21.27, 21.28, 22.3, 22.9,
 22.16, 22.19
- Storage 1.9, 1.13, 2.2, 3.2–3.5, 3.7, 3.8, 3.9, 3.10–3.13, 4.14, 20.12
- Straddle 13.1–13.4, 13.6–13.9, 13.11,
- Strangle 13.3, 13.4, 13.6–13.9, 13.11, 13.12, 15.5,
- Stress test 22.16,
- Sumitomo 1.9, 1.13, 1.14
- Supervision 22.15, 22.20
- Swap 1.6, 1.7, 1.13, 5.29, 7.8, 18.1–18.6, 18.8–18.12, 18.16–18.21, 18.24, 18.30–19.1, 19.5–19.13, 19.15–19.17, 20.2, 20.3, 20.8–20.10, 20.21, 22.15
- Swap Market Model (SMM) 19.1, 19.11, 19.15
- Swaps 1.4, 1.6, 1.7, 1.12, 4.14, 5.5, 18.1–18.4, 18.6, 18.9–18.11, 18.15, 18.17–18.21, 19.1, 19.17, 20.21, 20.22
- Swaption 19.1, 19.7–19.13, 19.15–19.17, 20.21, 20.22, Symmetry 16.6,
- Taxes 5.15, 18.28, 20.20, 21.3
- Theta 12.8–12.12, 12.15–12.17, 13.6, 13.9, 13.10, 13.11, 20.12, 20.13, 22.1
- Time decay 12.9, 12.16, 13.9
- Transportation 4.6, 4.7, 5.5, 5.6, 17.19, 17.22
- Treasury 1.3, 2.6, 2.12, 4.2–4.5, 4.9, 7.8, 17.15, 18.8, 18.32, 21.7, 21.8, 21.12, 21.16
- Uncertainty 1.3, 4.6, 4.15, 5.1, 5.6, 5.9, 5.14, 8.1, 8.4– 8.6, 10.2, 12.6, 15.5, 17.11, 18.6, 18.24, 20.8, 20.18, 21.3
- Unconditional 15.7, 15.8, 17.7,
- Uncorrelated 4.14, 22.13
- Underlying 1.1, 1.2, 1.4–1.6, 1.12, 1.13, 2.4–2.7, 2.10, 2.13, 2.14, 4.10, 4.14, 4.15, 5.22, 6.2–6.6, 6.8–6.10, 6.12, 6.13, 7.1–7.5, 7.8, 7.9, 8.2–8.4, 8.13, 11.1, 11.10, 12.1, 12.5, 12.6, 12.8, 12.9, 12.11–12.13,

- 12.15, 12.16, 13.1–13.12, 14.8, 14.10, 15.5, 15.6, 15.8, 16.1–16.3, 16.5, 16.6, 16.11, 19.3, 19.5, 19.7, 19.10, 19.12, 19.13, 19.14, 19.16, 20.2, 20.3, 20.12, 21.1, 21.7, 22.1, 22.8, 22.9–22.12, 22.14, 22.17– 22.19
- Unhedgeable 1.13, 22.18
- Unwinding 2.4, 7.1
- Upfront premium 5.5
- Utility theory 8.1, 8.4
- Valuation 1.12, 3.1, 3.3, 4.10, 8.1–8.8, 8.10, 8.13, 9.1– 9.3, 9.5–9.7, 9.11, 10.1, 10.3, 10.4, 10.6, 10.11– 10.13, 10.16, 10.20, 12.4, 12.11, 12.12, 13.5, 14.9, 15.9, 16.1, 16.6, 16.8–16.12, 17.1, 17.3–17.8, 17.10– 17.15, 18.1, 18.3, 18.6, 18.8, 18.10, 18.20, 19.2, 19.4, 19.10, 19.11, 19.13–19.17, 20.10, 20.11, 20.19, 21.8, 22.1, 22.17, 22.18
- Value 1.1, 1.4, 1.6, 1.12, 1.13, 2.5, 2.8, 2.9, 2.11–2.14, 3.1–3.3, 3.6–3.13, 4.2, 4.3, 4.5, 4.6, 4.8–4.13, 4.15, 5.1, 5.2, 5.9, 5.13, 5.22, 6.1, 6.2, 6.6–6.12, 7.4–7.7, 7.9, 8.1, 8.2, 8.5, 8.7, 8.8, 8.10–9.15, 10.3–10.16, 10.19,–10.21, 11.6, 11.10, 12.1, 12.8, 12.9, 12.11, 12.12, 12.14, 12.16, 13.1, 13.4, 13.5, 13.9, 13.10, 13.11–14.1, 14.4, 14.6–14.8–15.3, 15.5, 15.8, 16.1– 16.3, 16.5–16.13, 17.1–17.14, 17.20–17.23, 17.27– 17.29, 18.1, 18.3, 18.6–18.12, 18.17–18.20, 18.26, 19.1–19.3, 19.5, 19.6, 19.8–19.16, 20.1–20.21, 21.1,– 21.8, 21.11, 21.15, 21.22–21.29, 22.1–22.20
- Value additivity 18.10
- Value at Risk (VaR) 1.12, 22.1, 22.2, 22.19 Conditional Value at Risk (CVaR) 22.3 RiskMetrics 22.9, 22.14, 22.20, Tail Conditional Expectation (TCE) 22.3
- Variance 4.5–4.7, 4.15, 8.3, 8.13, 9.13, 9.14, 10.3,

10.4–10.18, 14.1, 14.3, 14.4–14.8, 14.10–15.3, 16.10, 22.8, 22.9, 22.11–22.13, 22.14

- Vega 9.15, 12.12, 12.13, 12.14, 12.16, 12.17, 13.1, 13.6–13.9, 13.11, 13.12, 15.8–15.10, 16.5, 16.6, 16.11, 20.12, 20.13, 22.1, 22.8, 22.17, 22.19
 - Exponentially weighted moving average 14.4–14.6, 14.8, 14.10, 22.14
 - Square root scaling 14.1, 16.10
- Volatility clustering 14.5
- Volatility surfaces 15.6, 15.8
- Volatility clustering 14.5
- Volatility smile, 1.12, 13.6, 15.1, 15.3–15.5, 15.8–15.10, 16.9, 19.5, 19.13, 22.5
- Risk reversal 15.5
- Volatility skews 15.4, 15.5
- Volume 1.1, 1.4, 2.9, 5.29, 5.31, 7.5, 7.8, 7.9, 17.29, 18.23
- Warrant 17.1–17.7, 17.9, 17.10, 17.12, 17.13,
- Weighted average 4.8, 4.9, 10.8, 14.4, 14.5, 14.8, 18.11, 18.13, 18.14, 18.16, 18.17, 19.5, 19.13, 19.15–19.17
- Yield 2.12, 2.13, 3.7–3.13, 4.2–4.5, 4.8, 4.14, 4.15, 5.22, 10.3, 10.6–10.10, 10.12–10.14, 10.20, 10.21, 12.3, 12.9, 12.14, 12.15, 12.17, 14.10, 15.9, 16.2, 16.6, 16.8, 16.11, 16.13, 17.6–17.8, 17.11, 17.12, 17.15, 18.1, 18.8, 18.9, 18.11, 18.12, 18.13, 18.14, 18.15–18.17, 18.19, 18.20, 18.32, 19.4–19.6, 19.11, 19.12, 20.16, 22.6, 22.16, 22.19
 Yield to Maturity (YTM) 18.8, 18.9
- Zero coupon bonds 18.10, 18.16, 19.14, 19.15, 22.9 Zero yields 18.11–18.13, 18.17, 18.19, 18.20, 19.5, 19.11