

Risk Management and Financial Derivatives

*A Guide to the
Mathematics*

Edited by

SATYAJIT DAS

McGraw-Hill

New York San Francisco Washington, D.C. Auckland Bogotá
Caracas Lisbon London Madrid Mexico City Milan
Montreal New Delhi San Juan Singapore
Sydney Tokyo Toronto

First published by The Law Book
Company Limited t/as LBC Information
Services, Sydney, Australia, 1997.

© Satyajit Das 1997. Copyright in this edition is held by the Editor.
Copyright in individual chapters is held by the relevant author(s).

This publication is copyright. Other than for the purposes of and subject to the conditions prescribed under the Copyright Act, no part of it may in any form or by any means (electronic, mechanical, microcopying, photocopying, recording or otherwise) be reproduced, stored in a retrieval system or transmitted without prior written permission. Inquiries should be addressed to the publishers.

Library of Congress Cataloging-in-Publication Data

Risk management and financial derivatives : a guide to the mathematics
/ edited by Satyajit Das.

p. cm.

Includes index.

ISBN 0-07-015378-7

1. Derivative securities—Mathematical models. 2. Risk management—Mathematical models. 3. Investment analysis—Mathematical models. I. Das, Satyajit.

HG6024.A3R577 1998

332.64'5—dc21

97-46822

CIP

McGraw-Hill



A Division of The McGraw-Hill Companies

Copyright © 1998 by The McGraw-Hill Companies, Inc. All rights reserved.
Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior written permission of the publisher.

3 4 5 6 7 8 9 DOC/DOC 9 0 2 1 0 9

ISBN 0-07-015378-7

The sponsoring editor for this book was Stephen Isaacs, the editing supervisor was John M. Morriss, and the production supervisor was Suzanne W. B. Rapcavage.

Printed and bound by R. R. Donnelley & Sons Company.

McGraw-Hill books are available at special quantity discounts to use as premiums and sales promotions, or for use in corporate training programs. For more information, please write to the Director of Special Sales, McGraw-Hill, 11 West 19th Street, New York, NY 10011. Or contact your local bookstore.

This book is printed on recycled, acid-free paper containing a minimum of 50% recycled de-inked fiber.



Preface

1. BACKGROUND AND OBJECTIVES OF THE BOOK

Modern financial management entails an appreciation of a number of key mathematical concepts. This is particularly relevant to risk management and risk management products, such as financial derivatives.

The knowledge of the mathematical concepts and the inherent assumptions underlying these mathematical concepts has tended to be concentrated among a select group of quantitative analysts employed by individual organisations (the so-called *rocket scientists* or *quants*). However, increasingly, the central role played by these products in capital markets is forcing a broader range of personnel to be aware of and utilise these concepts. This may be in a purely supervisory capacity or in using the concepts and products in day-to-day activities. This trend will continue and increase over time.

This increasing emphasis on mathematical finance creates problems for individuals who may not be comfortable with the basic concepts underlying the actual techniques utilised. This is usually caused by either a lack of exposure to the techniques or the absence of use of the concepts since graduation from university or other institutions. Hence, there is a very strong and increasing demand for material which explains the mathematical basis of risk management and financial derivatives in a *non-technical* manner to allow non-specialists to gain an appreciation of the concepts that are utilised.

Risk Management and Financial Derivatives: A Guide To The Mathematics is directed to fill this gap in the literature.

The book consists of a collection of papers from leading market practitioners covering:

- (1) the basic mathematics underlying risk management and financial derivatives products; and
- (2) application of the basic techniques in a number of common settings to promote understanding of the actual use of the concepts. Applications covered include the most common applications of mathematics to finance, such as:
 - yield curve modelling and bond/fixed income pricing;
 - pricing derivatives, both forwards and options;
 - investment management applications; and
 - risk management, in particular, value at risk and portfolio stress simulations based on monte carlo techniques.

The style of the book is as practical as possible. The text avoids mathematical notation to the maximum degree feasible so that an intuitive grasp of the concepts can be gained. The book also includes numerous

detailed worked examples, wherever possible, to help the reader understand the concepts and see how the practical calculations are undertaken.

The target audience for this book includes:

- financial institutions, particularly commercial and investment banks, as well as brokers, active in trading activities;
- liability and investment managers who utilise or are looking at utilising trading risk management techniques in the management of trading risk;
- service industries, consultants, IT firms, accountants et cetera, active in advising traders or users of these techniques; and
- regulatory agencies.

The book can also be used as the basis for practical in-house training programs, as well as in post-graduate programs such as MBA or Applied Finance courses in financial markets, either as the primary text or supplementary reading.

2. CONTENT AND STRUCTURE

The book is structured around several themes, which correspond to the parts of the book.

Part 1—Introduction

This consists of a single chapter (Chapter 1) designed to outline the fundamental basis of the application of mathematics to financial markets, in particular the essential risk reward basis of value and the use of risk neutrality and arbitrage as the basis for all valuation or pricing.

Part 2—Interest Rates and Yield Curves

This section and the three which follow it are focused on specific application areas of derivatives in risk management and derivative products. In this part, the focus is on fixed interest markets made: Chapter 2 covers interest rates, the pricing of bonds and measures on interest rate risk, utilising duration and convexity; and Chapter 3 covers yield curve modelling, with coverage of the derivation of zero-coupon rates and the generation of yield curves.

Part 3—Derivative Pricing

Part 3 covers the application of mathematical techniques in the pricing of financial derivatives. Chapter 4 covers the pricing of forward and futures contracts. Chapter 5 examines the pricing of option contracts. Chapters 6 and 7 focus on more advanced issues in option pricing—specifically, the pricing of interest rate options utilising term structure models and the valuation of non-standard or exotic options. Chapter 8 focuses on the estimation of volatility, which is a central issue in the valuation of options. Chapter 9 looks at more complex approaches to estimation of volatility using ARCH/GARCH approaches to modelling volatility changes. Chapters 10 and 11

examine the risk of options: Chapter 10 examines the measurement of option price sensitivity using the Greek alphabet of risk (delta, gamma, theta, vega and rho); and Chapter 11 focuses on the use of option price sensitivities to replicate option profiles (the practice of delta hedging).

Part 4—Investment Management

Part 4 focuses on investment management applications of mathematics. Chapter 12 examines the concept of efficient portfolios and diversification within investment portfolios and the optimisation of portfolio structures. Chapter 13 examines the immunisation of bond portfolio using duration-based hedging techniques. Chapter 14 covers the concept and practice of portfolio insurance, utilising basic option theory to guarantee minimum values of portfolios. Chapter 15 looks at the creation, structuring and management of indexed portfolios.

Part 5—Risk Management

Part 5 covers the mathematics of risk measurement and management. Chapter 16 focuses on the use of value of risk techniques to model trading risk. Chapter 17 examines the use of stress testing, primarily using monte carlo simulation methods to measure risk not captured by traditional models of risk. Chapter 18 extends the framework of risk management to cover credit risk.

Part 6—Mathematical Techniques

Part 6 sets out the mathematical techniques underlying the applications described in previous sections. In Chapter 19 the basic mathematics underlying the financial applications is described, together with material which allows the reader to further her or his knowledge of the techniques and so come to a more complete understanding of the subject matter. The chapter also includes cross-references to built-in mathematical functions found in spreadsheet packages, to allow the reader to program some of these applications.

The book is designed either to be read through from start to finish or as a reference source where individual sections are read as required.

4. ACKNOWLEDGMENTS

I would like to thank all the contributors to the book. They, all busy practitioners, add a practical dimension and real world application focus to the work.

I would like to thank the Publishers—LBC Information Services (Fiona Dixon, Carolyn Uyeda and Julie Burke); Irwin Publishing (Kevin Commins and Stephen Isaacs); and MacMillan Press (Jane Powell and Stephen Rutt). I would also like to thank Kim Paino, who edited the book.

I would like to thank my parents—Sukumar and Aparna Das—for their continued support and encouragement. In particular, I would like to thank my friend Jade Novakovic without whose help, support, patience, tolerance and encouragement this would never have been completed. This book is dedicated to these three people.

SATYAJIT DAS

Sydney
November 1997

Table of Contents

<i>Preface</i>	v
<i>About the Authors</i>	xi
<i>Selected Bibliography</i>	xvii

Part 1

Introduction

1. Risk-reward Relationships—Foundations of Derivatives	3
<i>by Lance Smith</i>	

Part 2

Interest Rates and Yield Curves

2. Interest Rates, Bond Pricing, Duration and Convexity	19
<i>by Roger Cohen</i>	
3. Interest Rate and Yield Curve Modelling	47
<i>by Satyajit Das (with a contribution from Roger Cohen)</i>	

Part 3

Derivative Pricing

4. Pricing Forwards and Futures Contracts	97
<i>by John Martin</i>	
5. Pricing Options.....	221
<i>by Satyajit Das</i>	
6. Interest Rate Option Pricing Models	275
<i>by Jim Rowlands</i>	
7. Pricing Models for Complex/Exotic Options	288
<i>by Dr Garry de Jager</i>	
8. Estimating Volatility.....	307
<i>by Satyajit Das</i>	
9. Estimating Volatility and Correlation Using ARCH/GARCH Models	337
<i>by Carol Alexander</i>	
10. Measuring Option Price Sensitivity—The “Greek Alphabet” of Risk.....	355
<i>by Satyajit Das</i>	
11. Option Replication Utilising Delta Hedging	390
<i>by Satyajit Das</i>	

Part 4**Investment Management**

- | | |
|--|-----|
| 12. Portfolio Optimisation | 431 |
| <i>by Geoffrey Brianton</i> | |
| 13. Risk Management for Bond Portfolios..... | 470 |
| <i>by Roger Cohen</i> | |
| 14. Portfolio Insurance..... | 482 |
| <i>by Stuart Roe</i> | |
| 15. Indexation of Portfolios..... | 500 |
| <i>by Frances Cowell</i> | |

Part 5**Risk Management**

- | | |
|---|-----|
| 16. Value at Risk Models..... | 547 |
| <i>by Satyajit Das and John Martin</i> | |
| 17. Portfolio Simulation: Stress Testing Techniques | 685 |
| <i>by Lance Smith</i> | |
| 18. Credit Risk Measurement..... | 697 |
| <i>by Alan Bustany</i> | |

Part 6**Mathematical Techniques**

- | | |
|-----------------------------------|-----|
| 19. Mathematical Techniques | 717 |
| <i>by Thomas R Gillespie</i> | |
| <i>Index</i> | 779 |

About the Authors

Dr Carol Alexander

After working in the mathematics department of the University of Amsterdam, and later as a bond analyst for Phillips and Drew (UBS), Carol Alexander joined the mathematics faculty at the University of Sussex in 1985. Since 1990 she has developed an international reputation for the time-series analysis of financial markets, specialising in risk measurement and investment analysis. In 1996 she moved to a part-time post at the university, when she became the Academic Director of Algorithmics Inc (www.algorithmics.com) and Editor in Chief of NetExposure, the electronic journal of financial risk (www.netexposure.co.uk).

During the past few years she has been consulting in risk management and time-series analysis for banks, corporates and other financial institutions. As a result, most of her academic research work now concentrates on applied financial econometrics, specialising in volatility and correlation analysis. Carol has developed a large number of general and in-house training courses covering the general areas of risk management and investment analysis for financial institutions.

With Algorithmics Inc, she is developing VAR models for large-scale risk management systems and new methods for historical simulation using pattern recognition. Her pattern recognition algorithm, developed with Dr Ian Giblin (Department of Mathematics, University of Pisa) won the first international non-linear financial forecasting competition in 1996. She also works with Dr Peter Williams (Department of Cognitive Science, University of Sussex) on using neural networks to estimate mixtures of normal distributions to model, "fat-tailed" distributions and term structures of kurtosis. Her research on emerging markets includes the analysis of equity and currency derivatives in the Asia-Pacific region, and the hedging of equities with new cointegration software.

She speaks at many international conferences and on mathematical techniques for risk and investment management and has written numerous articles in both academic and professional journals. Carol's books include the edited *Handbook for Risk Management and Analysis* (1st ed, April 1996; 2nd ed in two volumes, forthcoming February 1998, Wiley). More details of professional and academic publications are given on www.maths.sussex.ac.uk/Staff?COA.html.

Geoffrey Brianton

Geoffrey Brianton has an Honours degree in economics and statistics. He is an independent consultant in the area of investment risk management. For the last ten years he has worked as a Fund Manager in London, Sydney and Melbourne. During that time he has developed and managed a wide variety of quantitative funds, including indexed, capital protected and arbitrage

funds. With a background in econometrics and operations research, he has a particular interest in the use of optimisation techniques, both in regard to the problems of portfolio construction and some of the broader problems found in finance such as optional hedging strategies.

Alan Bustany

Alan Bustany is a Principal Consultant in the Financial Services Industry Practice of Price Waterhouse Urwick. He is a Wrangler from Trinity College, Cambridge, and was part of the British team at the 1977 International Mathematical Olympiad in Belgrade. He is an Associate Fellow of the Institute of Mathematics and its Applications, and of the Securities Institute of Australia. Alan has used his mathematical background in a range of technology-related fields, including relational database theory; artificial intelligence; formal methods of software engineering; knowledge-based systems; and financial derivatives.

Mr Bustany has more than 15 years' experience in the computing industry in consulting, commercial, and product development roles. His current focus is strategic consulting in financial markets, risk management, and credit measurement. His clients include Bankers Trust, National Australia Bank, Macquarie Bank and Westpac.

Roger Cohen

Roger Cohen is currently employed at Deutsche Morgan Grenfell, in the Equity Derivatives and Trading Group. His main focus is on the development of structured equity products, pricing and risk management. Roger commenced his career in the financial markets in the Quantitative Applications Division at Macquarie Bank in 1992. Since then he has worked at Natwest Markets in the Global Markets area. He has presented papers at finance conferences in Australia and abroad. Prior to entering the financial markets, Roger was a lecturer in the Faculty of Engineering at the University of Sydney. His main research focus was in the computational modelling of fluid motion.

Frances Cowell

Prior to graduating from university, Frances worked in the biomedical library at the New South Wales Institute of Technology, and then as a Rehabilitation Counsellor for disabled people. After obtaining a Bachelor of Arts degree in psychology and statistics from the University of New South Wales, she worked in the wholesale liquor industry as a Marketing Manager, before completing a Master of Business Administration at the Australian Graduate School of Management.

Frances entered the investment management industry in January 1983 as the Research Analyst for Aetna Life & Casualty, then a major life office. Faced with analysing a portfolio of 300 stocks, she employed a novel approach to stock analysis, using a desktop computer and spreadsheet to simultaneously analyse expected risk-return profiles for a large number of stocks, thereby identifying apparently mispriced issues for closer analysis.

This was combined with industry-wide analyses to estimate industry growth rates and of competitive advantage to identify best performing stocks within industry groups.

In 1984, with the growth of the share price index futures market in Australia, Frances was attracted by the opportunities for SPI arbitrage and joined Australian Bank to exploit these. There she also established and managed an early portfolio protection operation which enabled the bank's borrowers to cap their interest rate risk. This portfolio protection program drew on a variant of Black-Scholes option pricing technology (delta hedging) to hedge the bank's sold option positions.

Later, Frances combined stock index arbitrage principles with delta hedging to exploit underpriced SPI futures contracts with minimal risk. This operation, which was unique at the time, was conducted within a major Australian institutional broking house.

In 1991 Frances began working on indexed portfolios for a major Australian superannuation fund. This work covered domestic and global equities and fixed interest. She managed the derivatives enhanced domestic equities index portfolio, which grew to over A\$1 billion. From there she joined County NatWest to take over the management of index portfolios. At County, Frances further developed customising capabilities for indexation clients, with particular focus on adding value by managing after-tax returns to index portfolios.

Satyajit Das

Satyajit Das is an international specialist in the area of financial derivatives, risk management, capital markets, and treasury management.

He presents seminars on financial derivatives/risk management and treasury management/corporate finance in Europe, North America, Asia and Australia. He acts as a consultant to financial institutions and corporations on derivative instruments, risk management, and treasury/financial management issues.

Between 1988 and 1994, Mr Das was the Treasurer of the TNT Group, an Australian-based international transport and logistics company, with responsibility for the Global Treasury function, including liquidity management, corporate finance, capital markets, and financial risk management. He was also involved in the financial restructuring of the TNT Group in the early 1990s. During 1994, Mr Das acted as a consultant to the TNT Group in the areas of financial strategy and policy, capital allocation/management and strategic risk management.

Between 1977 and 1987, he worked in banking with the Commonwealth Bank of Australia, Citicorp Investment Bank and Merrill Lynch Capital Markets, specialising in fundraising for Australian and New Zealand borrowers in domestic and international capital markets and risk management, involving the use of derivative products, including swaps, futures and options.

In 1987, Mr Das was a Visiting Fellow at the Centre for Studies in Money, Banking and Finance, Macquarie University.

Mr Das is the author of *Swap Financing* (IFR Publishing Ltd/The Law Book Company Ltd, 1989); *Swaps and Financial Derivatives: The Global Reference to Products, Pricing, Applications and Markets* (IFR Publishing Ltd/The Law Book Company Ltd/Irwin Professional Publishing, 1994); *Exotic Options* (IFR Publishing Ltd/LBC Information Services, 1996); and *Structured Notes and Derivative Embedded Securities* (Euromoney Publications, 1996). He is also the editor of *The Global Swaps Market* (IFR Publishing Ltd, 1991). He has published on financial derivatives, corporate finance, treasury and risk management issues in professional and applied finance journals.

Mr Das holds Bachelors' degrees in Commerce (accounting, finance and systems) and Law from the University of New South Wales, and a Masters degree in Business Administration from the Australian Graduate School of Management.

Dr Garry de Jager

Garry de Jager holds a PhD in option pricing, a Masters of Business Administration and a Bachelor of Science in mathematics. He is currently the Senior Manager, International Capital Markets Research, for Chase Manhattan Bank, Sydney. His primary interests are modelling derivatives for commodities, foreign exchange and interest rate products. In his role supporting the trading desks, his special emphases are exotic options, hedging strategies, simulation analysis, volatility and correlation analysis, and provision of models for system development. He is a frequent traveller to Chase sites around the world, as well as a regular speaker at seminars on these topics.

Garry has wide commercial experience, having previously been the Director of Research and Development for the specialist financial option company Optech International; Marketing Representative for IBM; and Production Manager in the printing industry. Prior to joining Chase Manhattan he taught finance and Information Systems at the Queensland University of Technology.

Thomas R Gillespie

Thomas Gillespie graduated from the University of Sydney in 1986 with a Bachelor of Science, and in 1988 with a Bachelor of Arts majoring in mathematics and statistics. After graduation, Thomas joined Bankers Trust Australia and specialised in foreign exchange and fixed interest options. In 1989 he moved to James Capel Australia and there specialised in equity derivatives and arbitrage. This led to work with the HSBC group in Sydney, Tokyo and London on a number of projects, including Japanese derivatives arbitrage and automated trading systems. In 1994 Thomas decided to take up part-time studies again, earning a Graduate Diploma in Science in 1994, and is now pursuing a PhD at the University of Sydney. Thomas' current research interests are fitting alternative stochastic processes to financial time series, and the development of new estimation methods. Thomas is currently working with County NatWest Australia, specialising in equity derivatives and structured products.

John Martin

John Martin has worked, taught and published extensively in the areas of treasury, derivatives and financial risk management. He was closely involved in the development of the derivatives industry in both Australia and New Zealand in roles varying from market trader, risk manager, regulator, and educator. John's area of interest is in financial risk management, and he has written numerous articles on this topic, including his recently published book, *Derivatives Maths* (IFR Publishing Ltd, 1996).

Currently, John is a Divisional Director of Australian-based treasury advisory firm Oakvale Capital Ltd. He is responsible for the provision of specialist financial risk management advice to clients across a wide range of industries, including electricity, retailing, agriculture, and financial services.

Prior to joining Oakvale, John was the Risk Manager of the Sydney Futures Exchange from 1991 to 1994, and was then a Director of Financial Risk Management Consulting Pty Ltd. He has also held positions with the Reserve Bank of Australia, and with TNT as Manager, Treasury Planning.

John is regarded throughout the Asia/Pacific region as a leading author and speaker in the areas of treasury and financial risk management, and has lectured extensively to professional groups in the Asia/Pacific region.

John holds a Bachelor's degree in Economics with Honours from the University of Sydney.

Steuart Roe

Steuart Roe is a specialist in equities and finance for County NatWest Securities Australia Ltd, a member of the NatWest Markets Group. He specialises in derivative product development, implementation, and distribution.

Prior to joining County NatWest Securities Australia Ltd, he was an Associate Director, Structured Investments for County NatWest Investment Management Ltd. In this role he was responsible for the design and implementation of tailored investment solutions for individual clients. In particular, he was responsible for the management of in excess of A\$2 billion in assets, the bulk of which was in portfolio insurance.

Steuart holds a Bachelor of Science in mathematics and statistics from the University of Melbourne, and a Master of Applied Finance from Macquarie University.

Dr Tim Rowlands

After completing a PhD in theoretical chemistry at Cambridge University in the United Kingdom in 1989, Tim returned to Australia and commenced work for Macquarie Bank in a quantitative analysis role. His work involved supporting the Fixed Interest Division with emphasis on developing and implementing new products, including being part of a team implementing a term structure option pricing model. Tim then moved to the State Bank of New South Wales in 1992, and, from 1993, led the Quantitative Group with work focusing on interest rate options, retail product enhancements (including options incorporating prepayment risk) and risk management methodologies.

In 1995 Tim became Head of Treasury Risk Management, with responsibilities across the full spectrum of market and credit risk, policy and compliance. In mid-1996 Tim moved to Westpac to take responsibility for methodology in the trading risk area. This role has included reviewing and refining the bank's VAR calculations and the implementation of historical simulation capability for risk measurement. In addition, Tim has been involved in research into applying VAR techniques to credit and operational risk to improve capital allocation and risk adjusted performance measurement within the business.

Lance Smith

Lance Smith has a PhD in mathematics from Duke University and was an Assistant Professor of Mathematics at Columbia University prior to joining Salomon Brothers in 1986. At Salomon he worked on the equity proprietary trading desk and was responsible for the development of all the desk's financial pricing and hedging models used for assessing and limiting the risk in a derivatives book. He also developed and implemented several proprietary trading strategies, as well as hedged the more exotic derivative OTC transactions undertaken by the desk. In 1993 he and some of his long-term colleagues founded Imagine Software Inc. Along with developing state of the art risk management technology, he continues to explore new methodologies in pricing derivative securities and quantifying risk in the financial markets.

Selected Bibliography

Chapter 3

Adams, Kenneth J, "Fitting Yield Curves To Market Data" (1994) 1(2) *The Pacific ALM Journal* 27.

Adams, Kenneth J and Van Deventer, Donald R, "Fitting Yield Curves And Forward Rate Curves With Maximum Smoothness" (1994) (June) *The Journal of Fixed Income* 52.

Cox, David, "Yield Curves And How To Build Them" (1995) 4 *Capital Market Strategies* 29.

Flavell, Richard, "Discount Function Construction" *IFR Swaps Weekly*, 27 July 1994, p 4.

—— "Fitting Interest Rate Curves" (1995) *Matlab Financial Computing Brief*.

McEnally, Richard W and Jordan, James V, "The Term Structure of Interest Rates" in Fabozzi, Frank J and Fabozzi, T Dessa (eds), *The Handbook of Fixed Income Securities* (4th ed, Irwin Professional Publishing, Chicago, 1995), pp 779-830.

Smith, David, "By The Bootstraps" (1990) 3(6) *Risk* 40.

—— "Zero Coupon Yields" (1995) 5 *Capital Market Strategies* 51.

Chapter 4

Das, S, *Swaps and Financial Derivatives* (2nd ed, The Law Book Company Ltd, 1994).

Francis, Toy and Whittaker (eds), *The Handbook of Equity Derivatives* (1995).

Hull, J C, *Options, Futures and Other Derivative Securities* (Prentice-Hall, 1993).

Martin, J S, *Derivative Maths* (IFR Publishing Ltd, 1996).

Chapter 5

Black, Fischer, "The Pricing of Commodity Contracts" (1976) 3 (March) *Journal of Financial Economics* 167.

Black, Fischer, "The Holes in Black-Scholes" (1988) 1(3) *Risk* 30.

Black, Fischer, "How We Came Up With the Option Formula" (1989) (Winter) *Journal of Portfolio Management* 4.

Black, Fischer, "How to Use the Holes in Black-Scholes" (1989) 1(4) *Continental Bank Journal of Applied Corporate Finance* 67.

Black, Fischer, "Living Up To The Model" (1990) 3(3) *Risk* 11.

- Black, Fischer, Derman, E, and Toy, W, "A One Factor Model of Interest Rates and Its Application to Treasury Bond Options" (1990) (Jan-Feb) *Financial Analysts Journal* 33.
- Black, Fischer, and Scholes, Myron, "The Pricing of Options and Corporate Liabilities" (1973) 81 *Journal of Political Economy* 637.
- Cookson, Richard, "Moving in The Right Direction" (1993) 6(10) *Risk* 22.
- Cox, J and Rubinstein, M, *Options Markets* (Prentice-Hall, Inc, Englewood Cliffs, NJ, 1985).
- Derman, Emanuel, "Model Risk" (1996) 9(5) *Risk* 34.
- Garman, Mark B and Kohlhagen, Steven W, "Foreign Currency Option Values" (1983) 2 *Journal of International Money and Finance* 231.
- Grabbe, J Orlin, "The Pricing of Call and Put Options on Foreign Exchange" (1983) 2 *Journal of International Money and Finance* 239.
- Hull, John C, *Options Futures and Other Derivatives Securities* (2nd ed, Prentice-Hall, Inc, Englewood Cliffs, NJ, 1993).
- Hull, John C, *Introduction To Futures and Options Markets* (2nd ed, Prentice-Hall, Inc, Englewood Cliffs, NJ, 1995).
- Jarrow, Robert A and Rudd, Andrew, *Option Pricing* (Richard D Irwin, Homewood, Illinois, 1983).
- Leong, Kenneth, "In the Eye Of The Beholder" (1990) 3(7) *Risk* 38.
- Leong, Kenneth, "The Emperor's New Clothes" (1990) 3(8) *Risk* 11.
- Leong, Kenneth, "Exorcising The Demon" (1990) 3(9) *Risk* 29.
- Leong, Kenneth, "Price Versus Value" (1991) 4(10) *Risk* 22.
- Leong, Kenneth, "Model Choice" (1992) 5(11) *Risk* 60.
- Mason, R C, "Building On Black Scholes" (1988) 1(11) *Risk* 13.
- Meisner, James F and Labuszewski, John W, "Modifying the Black-Scholes Option Pricing Model For Alternative Underlying Instruments" (1984) (Nov-Dec) *Financial Analysts Journal* 23.
- Ritchen, Peter, *Options—Theory, Strategy and Applications* (Scott, Foresman and Company, Glenview, Illinois, 1987).
- Rowley, Ian, "Pricing Options Using The Black-Scholes Model" (1987) (May) *Euromoney Corporate Finance* 108.
- Rowley, Ian, "Option Pricing Models: How Good Is Black-Scholes?" (1987) (June) *Euromoney Corporate Finance* 30.
- Smithson, Charles W, "Wonderful Life" (1991) 4(9) *Risk* 37.
- Smithson, Charles W with Shang Song, "Extended Family (1) and (2)" (1995) 8(10) *Risk* 19; (1995) 8(11) *Risk* 52.
- Swiss Bank Corporation, *Understanding Derivatives* (Prospects Special, Basel, Switzerland, 1994).
- Tompkins, Robert, *Options Explained* (MacMillan Press, England, 1994).

Chapter 6

- Black, Scholes, (1973) 81 *Journal of Political Economy* 637.
 Black, (1976) 3 *Journal of Financial Economics* 167.
 Black, Derman, Toy, (1990) *Financial Analysts Journal* 33.
 Brace, Gatarek, Musiela, (1995) (Aug) *Risk* 34.
 Brennan, Schwartz, (1983) 4 *Finance* 119.
 Cox, Ingersoll, Ross, (1985) 53 *Econometrica* 363.
 Cox, Ross, Rubinstein, (1979) 7 *Journal of Financial Economics* 229.
 Heath, Jarrow, Morton, (1992) 60 *Econometrica* 77.
 Ho, Lee, (1986) 41 *Journal of Finance* 1,011.
 Ho, (1995) (Summer) *Journal of Derivatives*.
 Hull, *Options, Futures and Other Derivatives* (3rd ed, Prentice-Hall, 1997).
 Hull, White, (1990) 3 *Review of Financial Studies* 573.
 Jamshidian, (1989) 44 *Journal of Finance* 205.
 Jarrow, *Modelling Fixed Income Securities and Interest Rate Options* (1996).
 Longstaff, Schwartz, (1992) 47 *Journal of Finance* 1,259.
 Rebonato, *Interest-Rate Option Models* (Wiley & Sons, 1996).
 Rendleman, Bartter, (1979) 34 *Journal of Finance* 1,092.
 Smithson, (1995) (Oct-Nov) *Risk*.
 Vasicek, (1977) 5 *Journal of Financial Economics* 177.

Chapter 7

- Black, F, "The Pricing of Commodity Contracts" (1976) 3(1 and 2) *Journal of Financial Economics* 167.
 Black, F, Derman, E and Toy, W, "A One-Factor Model of Interest Rates and Its Application to Treasury Bond Options" (1990) (Jan-Feb) *Financial Analysts Journal* 33.
 Black, F and Scholes, M, "The Pricing of Options and Corporate Liabilities" (1973) 81 *Journal of Political Economy* 637.
 Conze, A and Viswanathan, "Path Dependent Options: The Case of Lookback Options" (1991) 56(3) *Journal of Finance* 1,893.
 Cox, J and Ross, S, "The Valuation of Options for Alternative Stochastic Processes" (1976) (Jan-Mar) *Journal of Financial Economics* 145.
 Cox, J C, Ross, S and Rubinstein, M, "Option Pricing: A Simplified Approach" (1979) 7 *Journal of Financial Economics* 229.
 Geske, R, "A Note on an Analytical Valuation Formula for Unprotected American Call Options on Stocks with Known Dividends" (1979) 7 *Journal of Financial Economics* 375.
 Goldman, M B, Sosin, H B and Gatto, M A, "Path Dependent Options: 'Buy at the Low, Sell at the High'" (1979) 34(5) *Journal of Finance* 1,111.
 Grabbe, J O, "The Pricing of Call and Put Options on Foreign Exchange" (1983) 2 *Journal of International Money and Finance* 239.

- Heath, D C, Jarrow, R A and Morton, A, "Contingent Claim Valuation with a Random Evolution of Interest Rates" (1990) 9(1) *The Review of Futures Markets* 54.
- Hull, J C and White, A, "Efficient Procedures For Valuing European And American Path-dependent Options" (1993) 1(1) *Journal of Derivatives* 21.
- Merton, R C, "Option Pricing When Underlying Stock Returns Are Discontinuous" (1976) 3 *Journal of Financial Economics* 125.
- Roll, R, "An Analytical Valuation Formula for Unprotected American Call Options on Share with Known Dividends" (1977) 5(2) *Journal of Financial Economics* 251.
- Rubinstein, M and Reiner, E, "Barrier Options" in *Exotic Options*, unpublished manuscript by M Rubinstein (1991).
- Whaley, R, "On the Valuation of American Call Options on Stocks with Known Dividends" (1981) 9 *Journal of Financial Economics* 207.

Chapter 8

- Carrado, Charles and Miller, Thomas, "Volatility Without Tears" (1996) 9(7) *Risk* 49.
- Chance, Don, "Leap Into the Unknown" (1993) 6(5) *Risk* 60.
- Clarke, Roger C, "Estimating And Using Volatility: Part 1" (1994) (Fall) *Derivatives Quarterly* 40.
- Cox, J and Rubinstein, M, *Options Markets* (Prentice-Hall, Inc, Englewood Cliffs, NJ, 1985).
- Engle, Robert F, "Statistical Models For Financial Volatility" (1993) (Jan-Feb) *Financial Analysts Journal* 72.
- Hargreaves, Guy D, "Volatility in Interest Rate Options", paper presented to a conference, "Options on Interest Rates", IIR Conferences, 9-10 March 1992, Sydney, Australia.
- Hull, John C, *Options Futures and Other Derivatives Securities* (2nd ed, Prentice-Hall, Inc, Englewood Cliffs, NJ, 1993).
- Hull, John C, *Introduction to Futures and Options Markets* (2nd ed, Prentice-Hall, Inc, Englewood Cliffs, NJ, 1995).
- Leong, Kenneth, "Estimates, Guesstimates and Rules Of Thumb" (1991) 4(2) *Risk* 15.
- Leong, Kenneth, "Mean Streets" (1991) 4(5) *Risk* 45.
- Murphy, Gareth, "When Options Price Theory Meets The Volatility Smile" (1994) (Mar) *Euromoney* 66.
- Tompkins, Robert, *Options Explained* (MacMillan Press, England, 1994).
- Tompkins, Robert, "Answer in the Cards" (1995) 8(6) *Risk* 55.

Chapter 9

- Alexander, C O, "Volatility and correlation forecasting" in C O Alexander (ed), *Handbook of Risk Management and Analysis* (Wiley, 1996).
- Alexander, C O, "Evaluating the Use of RiskMetrics™ as a Risk Measurement Tool for Your Operation: What are its Advantages and Limitations" (1996) 2 *Derivatives Use, Trading and Regulation* 277.
- Alexander, C O, "Splicing methods for Value-at-Risk" *Derivatives Week*, Learning Curve, June 1997.
- Alexander, C O and Leigh, C, "On the Covariance Matrices used in VAR Models" (1997) 4 (Spring) *Journal of Derivatives* 50-62.
- Alexander, C O and Chibumba, A M, "Orthogonal factor GARCH", forthcoming, *Journal of Empirical Finance*, 1997.
- Alexander, C O and Williams, P M, "Term Structure Forecasts of Foreign Exchange Volatility and Kurtosis: A Comparison of Neural Network and GARCH Methods", forthcoming, *Risk*, 1997.
- Becker, S, "A Survey of Risk Measurement, Theory and Practice" in C O Alexander (ed), *Handbook of Risk Management and Analysis* (Wiley, 1996).
- Bollerslev, T, "Generalized Autoregressive Conditional Heteroskedasticity" (1986) 31 *Journal of Econometrics* 307.
- Broadie, M and Glasserman, P, "Estimating Security Price Derivatives Using Simulation" (1996) 42(2) *Management Science*.
- Duan, J, "Cracking the Smile" (1996) 9(12) *Risk* 55-59.
- Engle, R F, "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation" (1982) 50(4) *Econometrica* 987.
- Engle, R F and Mezrich, J, "Grappling with GARCH" (1995) 8(9) *Risk* 112.
- Engle, R F and Rosenberg, J, "GARCH Gamma" (1995) 2 *Journal of Derivatives* 47.
- Figlewski, S, "Forecasting Volatility Using Historical Data", New York University Salomon Center (Leonard N Stern School of Business), Working Paper, Series No S-94-13, 1994.
- Fitzgerald, M D, "Trading volatility" in C O Alexander (ed), *Handbook of Risk Management and Analysis* (Wiley, 1996).
- Hull, J C, *Options, Futures, and Other Derivative Securities* (Prentice-Hall, 1993).
- Taylor, S J, "Modeling Stochastic Volatility: A Review and Comparative Study" (1994) 4(2) *Mathematical Finance* 183.

Chapter 10

- Brown, Stephen J, "Estimating Volatility" in Figlewski, Stephen, Silber, William L and Subrahmanyam, Marti G, *Financial Options* (Business One Irwin, Homewood, Illinois, 1990), pp 516-537.
- Chance, Don M, "Translating The Greek; The Real Meaning Of Call Option Derivatives" (1994) (July-Aug) *Financial Analyst's Journal* 43.

- Cox, J and Rubinstein, M, *Options Markets* (Prentice-Hall, Inc, Englewood Cliffs, NJ, 1985).
- Hull, John C, *Options Futures and Other Derivatives Securities* (2nd ed, Prentice-Hall, Inc, Englewood Cliffs, NJ, 1993).
- Hull, John C, *Introduction To Futures and Options Markets* (2nd ed, Prentice-Hall, Inc, Englewood Cliffs, NJ, 1995).
- Tompkins, Robert, *Options Explained* (MacMillan Press, England, 1994).

Chapter 11

- Asay, Charles and Edelsburg, Charles, "Can a Dynamic Strategy Replicate the Returns of an Option" (1986) 6(1) *Journal of Futures Markets* 63.
- Silber, William L, "Marketmaking In Options: Principles and Implications" in Figlewski, Stephen, Silber, William L, and Subrahmanyam, Marti G, *Financial Options* (Business One Irwin, Homewood, Illinois, 1990), pp 485-516.
- "The Dangers Of Neutrality" (1991) 4(10) *Risk* 48.
- Frye, Jon, "Greek Alphabet Soup: A Recipe For Success" (1988) 1(4) *Risk* 26.
- Hull, John C, *Options Futures and Other Derivatives Securities* (2nd ed, Prentice-Hall, Inc, Englewood Cliffs, NJ, 1993).
- Hull, John C, *Introduction to Futures and Options Markets* (2nd ed, Prentice-Hall, Inc, Englewood Cliffs, NJ, 1995).
- Leong, Kenneth, "Solving Mystery" (1990-1991) 4(1) (Dec-Jan) *Risk* 68.
- Rubinstein, Mark and Leland, Hayne E, "Replicating Options With Positions In Stock And Cash" (1991) (July-Aug) *Financial Analyst's Journal* 63.
- Rutherford, Janette, Sher, Answer, and Fitzgerald, Desmond, "Building Blocks" (1989) 2(7) *Risk* 42.
- Smith, A L H, *Trading Financial Options* (Butterworths, London, 1986).
- Westminster Equity, "Building Confidence" (1990) 3(4) *Risk* 33.

Chapter 12

- Arnott, Robert D and Wagner, Wayne H, "The Measurement and Control of Trading Costs" (1990) (Nov-Dec) *Financial Analysts Journal*.
- Black, Fischer, "Universal Hedging: Optimising Currency Risk and Reward In International Equity Portfolios" (1989) (July-Aug) *Financial Analysts Journal*.
- Black, Fischer and Litternam, Robert, "Global Portfolio Optimisation" (1992) (Sept-Oct) *Financial Analysts Journal*.
- Jeffrey, Robert H and Arnott, Robert D, "Is your Alpha Big Enough to Cover Its Taxes?" (1992) (Sept) *Journal of Portfolio Management*.
- Jobson, D and Korkie, B, "Putting Markowitz Theory to Work" (1981) (Summer) *Journal of Portfolio Management*.

Jobson, J D and Korkie, B, "Improved Estimation for Markowitz Portfolios Using James-Stein Type Estimators", proceedings of the American Statistical Association, Business and Economic Statistic Section, 1979.

Markowitz, Harry, *Mean-Variance Analysis in Portfolio Choice and Capital Markets* (Basil Blackwell Inc, Oxford, 1990).

Markowitz, Harry, "Portfolio Selection: Efficient Diversification of Investments", *Cowles Foundation Monograph No 16* (John Wiley and Sons, New York, 1959).

Perold, A F, "Large Scale Portfolio Optimisation" (1984) 30(10) *Management Science*.

Wolfe, P, "The Simplex Method for Quadratic Programming" (1959) 27(3) *Econometrica*.

Chapter 14

Black, F and Scholes, M, "The pricing of options on corporate liabilities" (1973) 81 (May-June) *Journal of Political Economy* 637.

Hull, J, *Options, Futures, and Other Derivative Securities* (2nd ed, Prentice-Hall International Inc).

Perold, A F and Sharpe, W F, "Dynamic Strategies for Asset Allocation" (1988) (Jan-Feb) *Financial Analyst Journal*.

Chapter 15

Chiang, A C, *Fundamental Methods of Mathematical Economics* (McGraw Hill, New York, 1974).

Kritzman, Mark, *The Portable Financial Analyst* (Probus, Chicago, Ill, 1995).

Lakonishok, Josef, Schleifer, Andre and Vishny, Robert W, *Study of the US Equity Money Manager Performance*, Brookings Institute Study, 1992.

Markowitz, H M, *Portfolio Selection: Efficient Diversification of Investment*, Cowles Foundation Monograph 16 (Yale University Press, New Haven, Ct, 1959).

Rosenberg, B, *Extra Market Components of Covariance in Security Returns* (1974) (Mar) *Journal of Financial and Quantitative Analysis* 263.

Rudd, Andrew, *Optimal Selection of Passive Portfolios* (1980) (Spring) *Financial Management* 57.

Rudd, Andrew and Clasing, Henry K Jr, *Modern Portfolio Theory* (2nd ed, Andrew Rudd, Orinda, CA, 1988).

Sharpe, William F, "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk" (1970) 19(3) *Journal of Finance* 425.

Sharpe, William F, *Portfolio Theory and Capital Markets* (McGraw Hill, New York, 1970).

Chapter 16

- Batlin, Carl, "Regulatory Approaches to Risk Management" (1996) 6 (June) *Financial Derivatives & Risk Management* 69.
- Beckstrom, Rod A, Lewis, Don and Roberts, Chris, "VAR: Pushing Risk Management to the Statistical Limit" (1994) 3 *Capital Market Strategies* 9.
- Beckstrom, Rod A and Campbell, Alyce R, "Value At Risk: Theoretical Foundations" (1996) 7 *Capital Market Strategies* 44.
- Beckstrom, Rod A and Campbell, Alyce R (eds), *An Introduction to Value at Risk* (C.ATS Software Inc, Palo Alto, 1995).
- Beder, Tanya Styblo, "VAR: Seductive But Dangerous" (1995) (Sept-Oct) *Financial Analysts Journal* 12.
- Bock, Jerome T, "VAR Approaches to Market Risk Management" (1996) 6 (June) *Financial Derivatives & Risk Management* 15.
- Cocks, Graham, "Measuring Market Risk Within a Bank" (1996) 1(2) *Journal of Applied Finance and Investment* 63.
- Morgan, J P, *RiskMetric™—Technical Document* (3rd ed, J P Morgan, New York, 1995).
- Morgan, J P, *RiskMetric™—Technical Document* (4th ed, J P Morgan, New York, 1996).
- Lawrence, Dr Colin, Robinson, Gary and Stiles, Matthew, "Incorporating Liquidity Into The Risk Measurement Framework" (1996) 6 (June) *Financial Derivatives & Risk Management* 24.
- Lawrence, Dr Colin and Robinson, Gary, "Value at Risk: Addressing Liquidity and Volatility Risks" (1996) 7 *Capital Market Strategies* 24.
- Leong, Kenneth, "The Right Approach" (1996) (June) *Risk Value-At-Risk Supplement* 9.
- Longerstaey, Jacques, "VAR, RiskMetrics™ And Market Risk Methodology" (1995) 6 *Capital Market Strategies* 9.
- Longerstaey, Jacques and Zangari, Peter, "Commoditising the VAR Framework: Incorporating Gamma Risk" (1996) 6 (June) *Financial Derivatives & Risk Management* 49.
- Matten, Chris, "The Capital Allocation Challenge for the Banks" (1995) 4-5 *Swiss Bank Corporation Prospects* 2.
- Matten, Chris, "Risk Adjusted Performance Measurement" (1996) 6 (June) *Financial Derivatives & Risk Management* 37.
- Ramaswami, Murali, "Why Capital at Risk is the Recommended Measure" (1996) 6 (June) *Financial Derivatives & Risk Management* 53.
- Salzberg, Mike, "Firmwide Risk Management Systems" (1996) 6 (June) *Financial Derivatives & Risk Management* 62.
- Sharma, Jitendra, "Practical Issues of Value-At-Risk" in (1996) *Corporate Finance Risk Management & Derivatives Yearbook* 8.
- Shimko, David, "What is VAR?" (1995) 8(12) *Risk* 25.
- Shimko, David, "VAR for Corporates" (1996) 9(6) *Risk* 28.
- Shimko, David, "Investor's Return On VAR" (1996) 9(7) *Risk* 27.

- Smith, Lance C, "Portfolio Simulation: Stress-Testing Techniques" (1996) 6 (June) *Financial Derivatives & Risk Management* 31.
- Smithson, Charles and Minton, Lyle, "Value at Risk" (1996) 9(1) *Risk* 25.
- Smithson, Charles and Minton, Lyle, "Value at Risk (2)" (1996) 9(2) *Risk* 38.
- Spinner, Karen, "Adapting Value at Risk" (1996) (April) *Derivatives Strategy* 14.
- Stambaugh, Fred and Cohen, Robert, "Value at Risk: Its Measurement and Uses" (1995) 4 (Dec) *Financial Derivatives & Risk Management* 45.
- Taylor, Charles R and MacDonald, William A, "The Future of Market Risk Management" (1996) 6 (June) *Financial Derivatives & Risk Management* 10.
- Turner, Chris, "VAR as an Industrial Tool" (1996) 9(3) *Risk* 38.
- Wilson, Duncan, "Marriage of Ideals" (1995) 8(7) *Risk* 7.
- Wilson, Duncan, "VAR in Operation" (1995) 8(12) *Risk* 24.

Chapter 19

- Abramowitz, M and Stegun, I, *Handbook of Mathematical Functions* (9th ed, Dover, 1972).
- Bhattacharya, R and Waymire, E, *Stochastic Processes with Applications* (Wiley, 1990).
- Black, F and Scholes, M, "The Pricing of Options and Corporate Liabilities" (1973) 81 (May-June) *Journal of Political Economy* 637.
- Feller, W, *An Introduction to Probability Theory and Its Applications* (Volume I, 3rd ed, Wiley, 1968).
- Feller, W, *An Introduction to Probability Theory and Its Applications* (Volume II, 2nd ed, Wiley, 1971).
- Garman, M and Klass, M, "On the Estimation of Security Price Volatilities from Historical Data" (1980) 53 (1) *Journal of Business*.
- Hogg, R and Craig, A, *Introduction to Mathematical Statistics* (5th ed, Prentice-Hall, 1995).
- Huber, P, *Robust Statistics* (Wiley, 1981).
- Hull, J, *Options Futures and Other Derivative Securities* (2nd ed, Prentice-Hall International, 1993).
- Hull, J and White, A, "The Pricing of Options on Assets with Stochastic Volatilities" (1987) 42 (June) *Journal of Finance* 281.
- Jarrow, R and Rudd, A, *Option Pricing* (Irwin, 1983).
- Jarrow, R and Rudd, A, "Approximate Option Valuation for Arbitrary Stochastic Processes" (1982) 10 *Journal of Financial Economics* 347.
- Jarrow, R and Rudd, A, "Tests of an Approximate Option Valuation Formula" in Brenner, M (ed), *Option Pricing* (Lexington Mass, DC Heath, 1983, pp 81-100).
- Knuth, D, *The Art of Computer Programming: Seminumerical Algorithms* (2nd ed, Vol 2, Addison-Wesley, 1981).

- Kreyszig, E, *Advanced Engineering Mathematics* (John Wiley and Sons, 1993).
- Launer, R and Wilkinson, G, *Robustness in Statistics* (Academic Press, 1979).
- Lawrence, C, Zhou, J and Tits, A, *Users Guide for CFSQP Version 2.4: A C Code for Solving (Large Scale) Constrained Nonlinear (Minimax) Optimisation Problems, Generating Iterates Satisfying All Inequality Constraints* (Electrical Engineering Department and Institute for Systems Research, University of Maryland, College Park, MD 20742 USA, 1996).
- Malliaris, A and Brock, W, *Stochastic Methods in Economics and Finance* (North-Holland, 1982).
- McKean, H, *Stochastic Integrals* (New York, Academic Press, 1969).
- Mendenhall, W, Scheaffer, R and Wackerly, D, *Mathematical Statistics with Applications* (4th ed, Duxbury Press, 1990).
- Microsoft Corporation, *Microsoft Excel User's Guide* (1993).
- Moro, B, "The Full Monte" (1995) 8(2) *Risk*.
- Neter, J and Wasserman, W, *Applied Linear Statistical Models* (Irwin, 1974).
- Neter, J, Kutner, M, Nachtsheim, C and Wasserman, W, *Applied Linear Statistical Models* (4th ed, Irwin, 1995).
- Papageorgiou, A and Traub, J, "Beating Monte Carlo" (1996) 9(6) *Risk*.
- Parkinson, M, "The Extreme Value Method for Estimating the Variance of the Rate of Return" (1980) 53(1) *Journal of Business*.
- Press, W, Flannery, B, Teukolsky, S and Vetterling, W, *Numerical Recipes in c* (2nd ed, Cambridge University Press, 1992).
- Ross, S, *An Introduction to Probability Models* (4th ed, Academic Press, 1989).
- Schuss, Z, *Theory and Application of Stochastic Differential Equations* (Wiley, 1980).
- Taha, H, *Operations Research: An Introduction* (5th ed, Macmillan Publishing, 1992).
- Tezuka, S, *Uniform Random Numbers: Theory and Practice* (Kluner Academic Publishers, 1995).

**Risk Management
and Financial Derivatives:**

A Guide to the Mathematics

Part 1

Introduction

WWW.TRADING-SOFTWARE-COLLECTION.COM
ANDREYBBRV@GMAIL.COM SKYPE: ANDREYBBRV



Chapter 1

Risk-reward Relationships—Foundations of Derivatives

by Lance Smith

1. INTRODUCTION

Classical portfolio theory examines risk/reward tradeoffs from a “mean-variance” framework. In this model, the “risk” of an individual security is encapsulated by the variance (or, equivalently, *standard deviation*) of its returns. The higher the variance, the more uncertain the return, and therefore the greater the risk. Diversification by assembling a portfolio of securities enables investors to decrease their variance while maintaining their expected profitability target.

The analysis of the risk in derivatives, as well as the pricing methodology, differs from the classical framework in many respects. In fact, in the case of equities, the only point they have in common is the assumption that the risks of an underlying security are characterised by the mean and variance of its returns. At this point, the analysis diverges.

First of all, a derivative security (such as an option) has an asymmetrical return pattern, so that although the risks of the *underlying* security may be summarised by a mean variance model, this description is clearly insufficient to adequately capture the risks in the *option*. We will see that the value of the derivative security is determined chiefly by the price of the underlying and its variance. However, in order to understand this connection, we must first introduce two concepts: *hedging* and *arbitrage*.

2. HEDGING AND ARBITRAGE

The process of reducing or eliminating a particular risk in a portfolio through a trade, or a series of trades, (or contractual agreements) is called *hedging*. The corresponding trades or contractual agreements are referred to as *hedges*. In a typical investment portfolio, there is little or no hedging; rather, the investor simply tries to achieve the greatest upside potential, given an acceptable level of risk. This is the framework of modern portfolio theory.

The pricing of derivatives goes to the other extreme: here, the requirement is to construct a portfolio (the hedging portfolio) that eliminates all of the risks introduced by the derivative security being analysed. In particular, the hedging portfolio is required to replicate a return pattern identical to that of the derivative security, so that, from the point of view of an investor, the two alternatives—replicating portfolio and derivative security—are indistinguishable.

This latter point introduces the notion of *arbitrage*. If the replicating portfolio and the derivative security produce the same return pattern, then they should have the same value. If they currently have different values in the marketplace, there is then an opportunity for *arbitrage*, that is, one can sell the higher-valued representation and purchase the lower-valued one, securing a risk free profit.

In the real world, things seldom work out as neatly, and there may be other reasons for an apparent arbitrage opportunity. Most pricing models ignore these secondary issues and focus on the primary risks. It is important, therefore, to understand the underlying assumptions and the implications for the pricing model. In the next section, we will illustrate these points by carefully inspecting two examples.

3. REPLICATING PORTFOLIOS

This section will attempt to provide an intuitive understanding of the “building blocks” of derivatives from a trading and risk management perspective, as opposed to an abstract mathematical one. The basic point of view is that in order to determine the price of a derivative security, one needs to understand how to *hedge* the security, and that the theoretical value is then determined by calculating the cost of the hedge. In this context, *hedging* will refer to a trade, or a series of trades in an appropriate *underlying security* in such a way as to offset the corresponding risk in the derivative security. In practice, there may be alternative choices of hedging security, but these can be compared to this basic case.

By examining the hedging process in some detail, we can arrive at a better appreciation of the risks that are not being adequately hedged; that is, we can better understand “where the model breaks down”. The trader’s job is then to determine a price for these unhedged risks given the trader’s current portfolio. The risk manager’s job is to understand and quantify the total unhedged risk in the firm’s position and ensure that it is maintained within acceptable limits, given the firm’s risk/reward profile.

We will illustrate these points by carefully examining two basic examples.

The first, that of a forward contract on a stock, requires only a *static* hedge. The second, an option on a stock, requires a *dynamic* hedge. In both cases we will see that the price of the security is determined by two considerations:

1. **cost of the hedge (or value of the replicating portfolio); and**
2. **compensation for unhedged risk.**

3.1 Example 1: A one year forward contract on a stock

A customer wishes to purchase 100,000 shares of a particular stock from you, one year from today. He wants to determine the price today at which he will purchase the stock in one year.

We will determine the price, the *forward price* of the stock, by examining the cost of the hedge. We seek a trade or series of trades that will exactly

offset the risk inherent in the forward contract. We know that in one year we need to have in hand 100,000 shares of stock. One way to achieve this is to purchase the shares of stock today and set them aside for delivery in one year. What will this cost? Let us assume that the current (spot) price of the stock is \$100, so that we can purchase the stock at \$100 per share today. What must we charge in one year in order to break even? The \$100 per share that I have spent on my hedge could have been invested elsewhere, such as in money market securities which earn interest. This is a relatively "risk free" transaction. We wish to place a hedge so that the forward contract is equally risk free. Put another way, as a trader, I will be charged interest on the \$100 per share that I have tied up in my hedge. If the rate at which I borrow money for one year is 10% (simple annual), then I will incur \$10 per share of interest charges during the life of the forward contract, so I need to receive \$110 per share at maturity in order to break even. So, at first pass, the forward price of the stock should be \$110.

However, if the stock pays a dividend of, say, \$3 over the next year, then I will receive \$3 per share in my hedging portfolio. This reduces the cost of my hedge and the customer will expect to be rebated by adjusting the forward price to \$107. That is:

Total cost of hedge = \$107 per share.

At this point we should stop and investigate if there are any risks that we have ignored. One risk is that the \$3 of dividends is not guaranteed. We can only *forecast* \$3, recognising that the company of the stock could adjust this amount either up or down. For this reason we might modify our price to provide a cushion against this event. An alternative is to simply agree to *pass through* the dividends at the maturity of the contract. In this way neither ourselves nor the customer bear the dividend risk, which is basically unhedgeable. In this case the forward price would remain at \$110 per share, and the customer receives any interim dividends.

We have also ignored the effect of changes in interest rates. If we finance our stock hedge overnight, then we are at risk to changes in interest rate levels because our forward price was determined assuming a fixed rate of 10%. This may be overcome by taking term financing for one year at 10%, instead of overnight.

Another risk which has been ignored by our neat analysis is *counterparty* risk; that is, the risk that the counterparty may default on moneys owed to us. Suppose that over the next three months the stock plummets to \$25 per share. Let us assume we have a pass-through forward so that the price of the original forward contract has been set at \$110. At this point our counterparty will owe us about \$85 more per share above the current market price, in nine months (the remaining term of the forward contract). On 100,000 shares this comes out to be \$8,500,000. This is essentially the (unrealised) amount of money that we have lost on our hedge, so that if the counterparty defaults we are really out this amount. For this reason there may be collateralisation requirements built in to the contract. We note that futures contracts are exchange-traded forward contracts with a daily collateralisation requirement (that is, variation margin).

Finally, another risk that has been ignored is that of *liquidity*, or the practical limitation of implementing a hedge in the marketplace. In order to

hedge our position in this example, we must purchase 100,000 shares of stock. Unless the forward price is calculated off of our average cost of acquiring the stock hedge, we will also be at risk in that the purchase of this much stock may impact on the price.

The main point we are making here is that the theoretical determination of the forward price is as we initially calculated. However, by thinking through what is required to actually hedge the position, we arrive at a better understanding of the risks involved.

Exercise:


Suppose the customer wishes to *sell* a forward contract on the stock. Suppose that in addition the stock can be borrowed or loaned for an annual fee of 2.0%. What should the forward price be?

We next turn to an example of a European style (no early exercise) equity option. Our goal is to understand the hedging process for such a security and its impact upon the theoretical valuation of the option. We will proceed as intuitively as possible.


3.2 Example 2: European style call option on a stock

A customer wishes to purchase a three month call option on 100 shares of stock. The current stock price is \$100, as is the exercise (strike) price of the option.

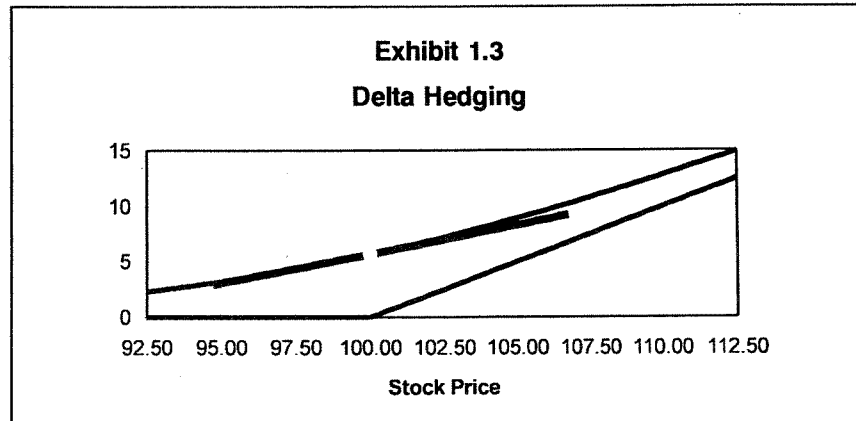
We will begin the analysis with the *expiration diagram* of this security (Exhibit 1.1). That is, a graph of the value of this security at expiration. Clearly, to the extent that the stock price exceeds the exercise price, the excess will be the value of the option, and if the stock price falls below the exercise price, the option will expire worthless. Therefore, the expiration diagram is simply a “hockey stick” as presented below.



We currently have three months to go on this security. What should the graph of the call option value look like when the three months are taken into account? A number of things are easy to see. First of all, for very low stock prices, say around 80, the option has very little value as the possibility of the stock rising 20 points is rather remote (again, we are speaking *intuitively* here). For stock prices around 100, the option has value in that if the stock price drops, there is nothing to lose, but if it rises, there is much to gain, in fact, point for point with the stock price. Finally, for high stock prices, say 120, the option is of course, worth at least 20 points but not a whole lot more because the option has as much to lose as it has to gain depending upon whether the stock falls or rises; put another way, the one-sidedness of the expiration diagram is not as evident with the stock at 120. These intuitive samples indicate that the graph should look something like the following:



Note that the graph begins to parallel the hockey stick at both extremes. We see that, on the downside, the graph should approach the expiration diagram. However, on the upside, it will approach a value that is a fixed amount above the expiration diagram. We now will set about to hedge our position. Obviously, as the stock price increases, so does the value of the call option. As we are short the option we will suffer a loss unless we have positioned a hedge that compensates for this loss. One way is to purchase some shares of stock. Then, as the stock price increases, our hedge will make money to offset the losses on the option position. The flip side is of course that if the stock price *drops* our hedge will *lose* money, but our short option position will *gain*. The first question is: How much stock should we purchase? If we make a brief return to first semester calculus we will recall the notion of *tangent line*. The amount of stock we should purchase will correspond to the *slope* of the tangent line. For high stock prices this slope approaches 100%, and for low stock prices it approaches 0%.



Let us suppose that the current slope, corresponding to the current stock price of 100, is 50%. We should then purchase 50 shares of stock, since the option is actually on 100 shares, not just 1. Suppose next that the stock rallies to 102, where the slope is now 60%. We are underhedged! In order to get hedged again we should purchase another 10 shares, paying 102 per share. Now suppose the stock falls back to 100. Now we are *overhedged* and should sell off 10 shares at 100. Of course we are now back to where we began, except that we have lost 2 points on 10 shares for a \$20 loss. If the stock continues to see-saw in this fashion we will continue to “buy high and sell low” in the process of actively hedging our position. This is an odd way to make money, and in fact we appear to be losing. However, the other side of the coin is that, as time passes, the value of the option should “decay”. That is, the graph eventually should approach the hockey stick. To the extent that the option decays, we will make back some of the money we are losing by hedging. It is a race between the two effects: trading losses versus time decay. In other words, if the stock is *volatile*, we will be whipsawed more and lose the race, and therefore lose money. If the stock is very quiet, we will win the race and make money. The famous Black-Scholes option pricing formula basically calculates what these hedging costs will be over the life of the option and then asserts that the (present value) of these costs should equal its price. That is, just as in the calculation for the forward contract, *the value should equal the cost of the hedge*. This time it is more complicated, of course, because the hedging process (termed “delta-hedging”) is *dynamic*. We summarise this discussion below.

The theoretical price of an option is equal to the cost of hedging the option.

With this point in mind, we can now examine what *information* is required to price an option. Clearly, the terms of the option (exercise price, expiration date, et cetera) are required. We also need to know various market parameters such as the stock price, dividend information and interest rates, as these are all relevant to calculating the cost of the hedge, just as for the stock

forward. However, because of the dynamic nature of the hedging process there is one more parameter which we need: the *volatility* of the stock (or its square, the *variance*). This is a statistical measurement, calculated as an (annualised) standard deviation of the stock price returns that quantifies the extent of the “whipsawing” we can expect when dynamically hedging the option position. The more volatile the stock, the more expensive our hedge, so that the price of the option should increase with the volatility of the stock. Typical values range from 10% to 50%. The lower end corresponds to a stock index, while a blue chip stock is typically in the 20%-35% range.

We summarise the required information in a table:

<p>Inputs for Option Valuation</p> <p>Contractual Terms</p> <p>Exercise Price</p> <p>Expiration Date</p> <p>Style (American or European)</p> <p>Market Information</p> <p>Stock Price</p> <p>Interest Rates</p> <p>Dividend Information</p> <p>Statistical Information</p> <p>Volatility</p>
--

At this point it is important to note that all of the inputs are known *except for the volatility* (and to some extent the dividends). That is, the cost of hedging will be influenced by the volatility of the stock *during the life of the option*. This is not really known, but must be predicted. We can certainly measure the past volatility of the stock, but there is no guarantee of the future volatility. This has been compared to “skiing down a slope backwards”, watching the trees go by. Because of the inherent intractability of the volatility parameter, there is always a degree of uncertainty in what the hedging costs will be. In short, if we are hedging only by dynamically adjusting our stock hedge, this parameter risk is unhedged. In a typical situation a trader will have a portfolio of options, both long and short of varying maturities, so that the risk of the unknown volatility is netted across the entire portfolio. Then the risk manager needs to be concerned with issues such as three month volatility versus six month volatility; that is, the “term structure” of volatility. We will not explore this aspect here; it is a natural extension of the ideas presented so far.

Exercise:

For very high stock prices, we have stated that the option value approaches a fixed amount above the expiration diagram. Calculate this value by calculating the cost of the hedge (note that the slope is 100%).

4. GLOSSARY OF TERMS

Before we can continue with our discussion, it is important that we define some frequently used technical terms.

A Glossary of Greeks

<i>Delta</i>	Greek symbol δ . This is simply the slope of the tangent line. This risk is <i>hedgeable</i> with the underlying stock.
<i>Gamma</i>	Greek symbol γ . This measures how rapidly the slope changes. For a trader, it is used to anticipate how much rehedging will be required for a given move in the stock price. This risk is <i>unhedgeable</i> , except with other option-like securities.
<i>Theta</i>	Greek symbol θ . This measures the time decay. In a sense (referring to the "race" mentioned earlier), theta is the flipside of gamma. This risk is also <i>unhedgeable</i> , except with other option-like securities.
<i>Sigma</i>	Greek symbol σ . Measures the <i>volatility</i> of the stock (expressed as an annualised standard deviation of returns). Also <i>unhedgeable</i> , except with other option-like securities.
<i>Kappa</i>	Greek symbol κ (also called "vega"). This measures the sensitivity to the volatility assumption, σ . Again, <i>unhedgeable</i> except with other option-like securities.
<i>Rho</i>	Greek symbol ρ (also called "dv01"). Measures the sensitivity to interest rates for (usually) a one "basis point" change (that is, .01%). This is <i>hedgeable</i> by trading, for example, an appropriate bond.

5. RISK NEUTRALITY

Up to now, we have examined derivative pricing from the perspective of hedging. As it turns out, there is an alternative method of calculation that is actually a byproduct of the hedging approach, termed the "Risk Neutrality Hypothesis". It states that the price of the derivative security can be

calculated by making simplifying assumptions about the underlying process, and then computing its “expected value” under the simplified process, discounting as if it were a known cashflow. We will first illustrate this principle in the case of an ordinary stock option, and then indicate why it is true with a brief foray into the binomial world.

This time we begin again with the hockey stick expiration diagram. We next calculate the “expected value” of this by simply taking each point on the diagram and multiplying by a probability, and then adding them all up. This sounds tedious; fortunately, computers are very good at this. The resulting number is then present valued to today. The probability distribution is obtained by taking the original process for the stock price (that is, *lognormal*—we have carefully avoided actually writing it down) and replacing the *expected return* of the stock—a very subjective number—with the “risk free” rate (actually, our financing rate). We have retained the stock volatility, σ , which is in general a less subjective number than the expected return.

In order to understand why the risk neutral calculation is equivalent to calculating the hedging costs, we will consider a simple one step binomial “tree”. That is, suppose that over the next time period the stock, currently at a price of S , can either go up to a price S^+ , or go down to a price S^- . We will assume that $S^+ = uS$ and $S^- = dS$ with $0 < d < 1 < u$ (that is, u stands for “up” and d stands for “down”). *We do not assume that we know the probability of either event.* This is tantamount to not knowing the expected return on the stock. Next, we attempt to hedge an option on this stock over the next time period, assuming that at the price S^+ it will equal C^+ and at S^- it will equal C^- (for example, if the option expires in the next time period). We wish to construct a hedge of stock and cash (accruing at a rate of $r\%$) which will “hedge” this option. To that end, we assume that we hold a position of m shares of stock and n units of cash (m will be a fractional share in this calculation). We certainly are not concerned about roundlots right now. Thus, the value of our hedging portfolio is currently

$$mS + nB$$

We will assume that over the next time period that B will grow to FB ; F is a “future value” factor; and $F-1$ is the interest earned per unit of cash during this time period. Our task is to determine m and n so that we are hedged in the two events $S = S^+$ and $S = S^-$. That is, we must determine m and n so that:

$$mS^+ + nFB = C^+$$

and

$$mS^- + nFB = C^-$$

These are two equations in two unknowns which are easily solved. For example, we find that m is just

$$m = (C^+ - C^-) / (S^+ - S^-)$$

which corresponds exactly to the tangent line slope originally discussed. In a similar fashion we can solve for n , and substitute back in to find that, after rearranging terms

$$mS + nB = (pC^+ + qC^-) / F$$

where

$$p = (F - d)/(u - d)$$

and

$$q = (u - F)/(u - d).$$

We note that

$$p + q = 1$$

and

$$0 < p < 1$$

$$0 < q < 1$$

as long as $F < u$. This last assumption simply asserts that the stock must have a chance of outperforming the cash instrument in any time period. (Otherwise, why would anyone ever buy the stock!) These equations imply that p and q can be interpreted as *probabilities* (termed *arbitrage probabilities*) and the value of the hedging portfolio today, which equals the option price today, is simply equal to the expected value of its price over the next time period, using the arbitrage probabilities), discounted by the factor F . This is nothing more than a binomial description of the risk neutrality principle.

The power of the risk neutrality methodology is that once it has been demonstrated that a particular derivative security can be hedged—usually by considering a similar one step binomial tree (or in the case of more complicated multifactor securities, a multinomial tree)—an elaborate mathematical machinery can then be called into play to actually perform the calculation. The main *drawback* of this methodology is that one can be easily seduced into the risk neutral world where none of us actually live, and forget some of the model assumptions that brought us there. This is why, from the point of view of risk management, it is important to understand how the models can break down, and how to best *stress test* them.

6. APPLICATIONS OF THE RISK NEUTRALITY PRINCIPLE

The risk neutrality calculation methodology can be applied to a wide variety of derivative securities. The key point is to verify the hedgeability of the relevant risks; the problem is then reduced to an “expected value” calculation, which is usually quite tractable or at least amenable to a wide variety of mathematical techniques.

6.1 Example: A three month “look back” option

This is a security that will payout in three months the difference between the *maximum* stock price reached over the next three months, S^* , and today's stock price, K . That is, the payment will be:

$$S^* - K.$$

K is currently known, while S^* is not. Is this hedgeable? We resort to binomial logic. With one period to go to expiration, we will know the current maximum S^* , and therefore, the maximum at expiration in either case $S = S^+$

(where the maximum will be the greater of S^+ and S^*), or $S = S^-$ (where the maximum will remain S^*). This means that we can then construct our hedge just as before.

This reasoning can be extended to earlier periods as well, establishing that the security is hedgeable. However, in order to *calculate* the theoretical price, we can now invoke the Risk Neutrality Hypothesis and employ alternative means. Now any mathematician can apply advanced techniques (partial differential equations, Green's functions, stopping times, Monte Carlo as a last resort, et cetera) to solve the problem. Some of these techniques have been around for over 100 years and have become practical with the advent of computer technology. For American style (early exercise) securities, this calculation is performed over a small time interval, after which the security can be tested for early exercise, just as in a binomial option pricing model.

Once the model is in hand, it should be *stress tested* in order to reveal hidden risks that are *unhedgeable* (except with other options or option-like securities).

7. ARBITRAGE—A CLOSER LOOK

As we have noted, the basic premise of all pricing models is that the value of the derivative security should equal that of a replicating portfolio; conversely, if the values differ, than arbitrageurs can play one against the other and obtain a riskless profit. In an efficient market this arbitrage activity should force convergence of market prices to theoretical prices.

All of this is approximately true, and more true for some derivative securities than for others, but there are often legitimate reasons for a "mispricing", usually due to features of the marketplace that have not been adequately incorporated into the pricing model. Obvious features are transaction costs such as brokerage fees and stamp duties. Some others are listed below:

7.1 Counterparty risk (for OTC transactions)

This has been illustrated in the example of the forward contract.

7.2 Cash flow risk

The pricing models usually assume that money can be borrowed whenever it is needed. This may not be really true. If we return to the example of the forward contract, recall the scenario where the stock has fallen to \$25 a share. If we have purchased the stock on margin, this would trigger a massive margin call; if we have insufficient funds, our hedge will be liquidated.

7.3 Parameter risk

As discussed in the example of the European style call option, the cost of the hedge will depend largely upon the experienced volatility during the hedging. This parameter can only be estimated.

7.4 Horizon risk

This is a more subtle risk, and has to do with the *rate of convergence* between the market price and the theoretical price; that is, the rate at which the trade increases in profitability. In the case of a three month call option, it will take at most three months. But for, say, a convertible bond, it could take as long as 15 years. This lengthy horizon can easily disincentivize arbitrageurs from stepping in (they tend to have notoriously short time horizons).

8. CONCLUSION

This chapter has discussed the pricing of equity derivatives, but the same techniques apply to other financial derivatives as well. The general procedure is to:

- (i) identify the primary risks and a mathematical model for their evolution through time;
- (ii) identify underlying “hedging instruments” for the risks;
- (iii) verify the hedgeability (this may be done with an appropriate binomial tree, or *multinomial* tree in the case of multiple risks);
- (iv) calculate the value by invoking the Risk Neutrality Hypothesis; and
- (v) consult a mathematician to actually perform the calculation. Also ask for a list of the model parameters. These will give rise to corresponding sensitivities which should be stress tested in order to develop a better understanding of the unhedged risks.

Part 2

Interest Rates and Yield Curves

Chapter 2

Interest Rates, Bond Pricing, Duration and Convexity

by Roger Cohen

1. WHAT IS AN INTEREST RATE?

The term interest rate is commonly used to describe the growth or earning potential associated with an amount of money. Common occurrences of interest rates include the advertised return on money deposited in a bank account, home loan mortgage rates, financing costs and so forth. The types of rates, and the context in which they are used is often confusing. In this chapter, various representations of interest rates will be introduced and explained. The context in which interest rates are used in the financial markets will be explained.

An interest rate refers to the rate of growth or decay of an asset over time. It is a measure of the value of the asset at the present relative to its value in the future. Although the asset is usually cash, it need not be. Interest rates allow us to quantify questions such as “*is it worth more now or in the future?*” or “*what will this be worth in ten years time?*”. They are a measure of the earning (or expense) associated with deferred consumption. For example, an individual may have a sum of money that is not needed at the present time. Another individual may need money (perhaps to buy food, or to build a house). The two individuals can enter into an agreement where the latter gets the use of the money at the present time. The full amount plus an additional sum will be repaid at a future date. This additional sum is the interest paid by the borrower to the lender. By deferring consumption of the money, the lender is rewarded with extra cash or interest. This is the governing principle of most financial transactions. The financial markets just formalise the mechanics of such transactions. They provide an efficient framework for transferring capital. The underlying principal is that a lender receives interest for the use of capital by a borrower. This applies where the lender is an individual, an organisation or even an entire country. The converse applies to the borrower.

Although varying in complexity, interest rate transactions involve a borrower and a lender. A bank deposit, for example, is effectively a loan by an individual to a bank. The underlying rationale is that the bank is able to use this money to earn income greater than the expense associated with the interest paid to the depositor. An example of this is that home loan mortgage rates (where the bank is the lender and an individual is the borrower) are invariably higher than deposit rates (where the bank is now the borrower). The appropriate rates of interest are dependent on the parties involved in a transaction, the risk of the transaction (that is, whether the borrower will be

able to repay the debt), and the amount, structure and timing of the transaction. This will not be discussed further here.

Interest rates are also used to quantify the growth or decay of commodities other than money. This includes gold and other precious metals, and certain financial instruments such as stocks. Interest rates need not be positive. If there are storage costs, or time wastage associated with holding a commodity, then its interest rate may be negative. This means that a lender of such a quantity will pay a fee to the borrower.

1.1 Representations of interest rates

There are many representations of interest rates. Usually (but not always) rates are a positive amount expressed as a percentage. They are measures of the growth or decay of an asset. For example, an advertised rate of 10% paid annually, means that \$100 will be worth \$110 after one year. Generally a rate is expressed as a percentage, and a basis. The percentage gives the amount of growth that is expected, the basis tells the period over which this growth is compounded. If the 10% of the above example were paid semi-annually (that is, twice a year), then after six months \$100 would be worth \$105.¹ If the interest is compound, then after a further six months, the \$105 is increased by 5% giving \$110.25 after one year. From this we can see that 10% semi-annually is then equivalent to 10.25% as an annual rate. Similarly, 10% quarterly would mean a growth of 2.5% is applied each quarter. Simple rates are applied over a period without compounding. An annual simple rate of 10% would increase \$100 to \$110, \$120 and \$130 over 1, 2 and 3 years respectively. If the 10% were compounded annually, then \$100 would increase to \$110, \$121 and \$133.10 respectively.

1.2 Interest rate arithmetic

There is nothing special about how an interest rate is expressed. Rates can be changed from one basis to any other. The premise for doing this is to realise that, at the end of a period, the final amount must be the same when any rate basis is used. To convert from one basis to another,

$$\left(1 + \frac{r_1}{b_1}\right)^{b_1} = \left(1 + \frac{r_2}{b_2}\right)^{b_2}$$

where r_1 and r_2 are interest rates with bases b_1 and b_2 .

Example:

What is the quarterly equivalent of 7.5% semi-annual?

$$\left(1 + \frac{r}{4}\right)^4 = \left(1 + \frac{0.075}{2}\right)^2$$

This solves to give $r = 0.07431$ or 7.431% as a quarterly rate.

1. By convention, non-annual rates are divided by the number of periods in a year. Thus 10% semi-annually means interest is 5% every half year; 10% quarterly would be 2.5% per quarter and so forth.

An interest rate can be converted from any basis to any other.² It is convention to express the periodicity of rates in years. There are various conventions for determining the number of days in a year. These will be discussed below.

Commonly used interest rate bases include

1. *Simple*: This is where the rate is expressed exactly over the period required. For example, 8% over 45 days would mean that \$1 grows to $\$1(1 + 0.08 \cdot 45/365) = \1.0099 after 45 days.
2. *Daily*: These rates are compounded daily. Cash or overnight rates set by most central banks are daily effective. A daily rate of 10% means that after n days, \$100 increases to $\$100(1+0.1/365)^n$.
3. *Monthly, quarterly and semi-annual*: Rates are compounded 12, four and two times per year.

1.3 Year basis

In financial calculations, the number of days per year used when calculating interest periods is not necessarily the actual number of days in the period. Convention uses either a 365 day or a 360 day year. Months may use the actual number of days or be considered as having a fixed number of days—usually 30. Arithmetic for these conventions can be found in DAS, *Swaps and Financial Derivatives* (2nd ed, 1994), pp 174-178.

1.4 Continuous rates

Another interest rate basis is that of continuous or continuously compounding rates. These never actually appear explicitly in financial instruments or transactions. They are a representation most often used internally in financial calculations. They provide a simple means for performing interest rate arithmetic. Within many financial models, rates are converted to continuous rates. They are converted back to their original basis after manipulation.

In the section above, we showed various compounding periods ranging from one year down to one day. If the compounding period is decreased further, then in the limit (where the period becomes infinitesimal), we have continuously compounding rates. If we had a rate r compounded continuously, then after one year, we would have the following growth:

$$\lim_{n \rightarrow \infty} \left(1 + \frac{r}{n}\right)^n = e^r$$

An amount $\$A$ would grow to $\$Ae^r$ after one year. For the general case, using rate r for time t (in years), the growth would be e^{rt} . The advantage of continuously compounding rates will become apparent when we introduce

2. It is interesting to note that when rates are advertised, they are commonly expressed in a basis that makes them look attractive to an investor. Deposit rates are often expressed as annual effective even when they are compounded (10.25% annual effective looks more attractive to the uninformed than 10.00% paid semi-annually). For borrowing rates, the converse is often the case. Mortgage rates are often compounded monthly or even daily (a rate of 10% daily is the same as 10.516% annual effective).

forward rates later in this chapter. It is also essential in many models used for option pricing.

Exhibit 2.1				
Value of \$1 in 1 Year, at 10% Interest Using Various Bases				
Basis	Periods Per Year	Rate	Formula	Value
Annual Effective	1	10%	$(1 + 0.1)$	\$1.10000
Semi Annual	2	10%	$(1 + 0.1/2)^2$	\$1.10250
Quarterly	4	10%	$(1 + 0.1/4)^4$	\$1.10381
Monthly	12	10%	$(1 + 0.1/12)^{12}$	\$1.10471
Daily	365	10%	$(1 + 0.1/365)^{365}$	\$1.10516
Daily (360)	360	10%	$(1 + 0.1/360)^{360}$	\$1.10516
Continuous	∞	10%	$e^{0.1}$	\$1.10517

10% continuously compounded is the same as 10.517% annual effective.

It is a common perception that continuously compounded rates are complex. This is not true. The reality is that working with continuously compounding rates is simpler than with discrete rates. This is due to the exponential growth and decay used when valuing with continuous rates. The perception of difficulty arises as rates need to be converted to continuous before they can be used, and then often back to simple rates after manipulation. To convert a rate to continuous we use the relationship

$$e^{rt} = (1 + R)^T$$

where r is the continuous rate applied over time t years, and R is the periodic rate applied over T periods. The relationship between t and T is $t = T \cdot f$ where f is the number of periods per year over which R is effective.

Example: continuous and discrete rates

What is the continuous and annual equivalent of a semi-annual rate of 7.5%?

When converting a rate from one basis to any other, we will still get the same return. After one year, the semi annual rate R will gross one dollar up to $\$(1 + R/2)^2$ dollars. A continuously compounded rate r will gross a dollar up to $\$1 \cdot e^r$ dollars. For equivalence, these must be equal. Similarly, the annual rate r' will also gross up to the same dollar value.

$$\left(1 + \frac{R}{2}\right)^2 = e^r = (1 + r')$$


Solving when $R = 7.5\%$ gives an annual effective rate of 7.6406% and a continuous rate of 7.3628%.

$$\left(1 + \frac{0.075}{2}\right)^2 = e^{0.073628} = (1 + 0.076406)$$

The continuously compounded rate can be used in calculations. It is easier to work with exponentials and logarithms rather than the discrete representations.

1.5 Valuing cashflows

From the discussion of interest rates above, we are now in a position to value any cashflow to or from the present day. If we are valuing future cashflows to the present, this is often referred to as *discounting* or obtaining the *net present value* or *NPV*. The reverse—finding the value in the future of an amount of cash at present—is referred to as *grossing up*.



Any cashflow can be valued by discounting at an appropriate rate

$$Val = \frac{a}{(1+r)^t} = e^{-Rt} = a.DF(t)$$

The concepts of discounting and grossing up form the basis of most financial transactions. Investors are trying to maximise their return or the NPV of their assets, while borrowers are seeking the minimum cost for their borrowings. This is all done within a framework where the risks involved and the structure of the transactions are considered.

1.6 Forward rates

As well as being able to express interest rates in many different bases, it is also important to specify the exact period over which the interest rates apply. Rates in the above sections are considered to apply from the present to a date in the future, or vice versa. These are referred to as *spot rates*. The spot date is usually the present date, or the date out of which a transaction begins.

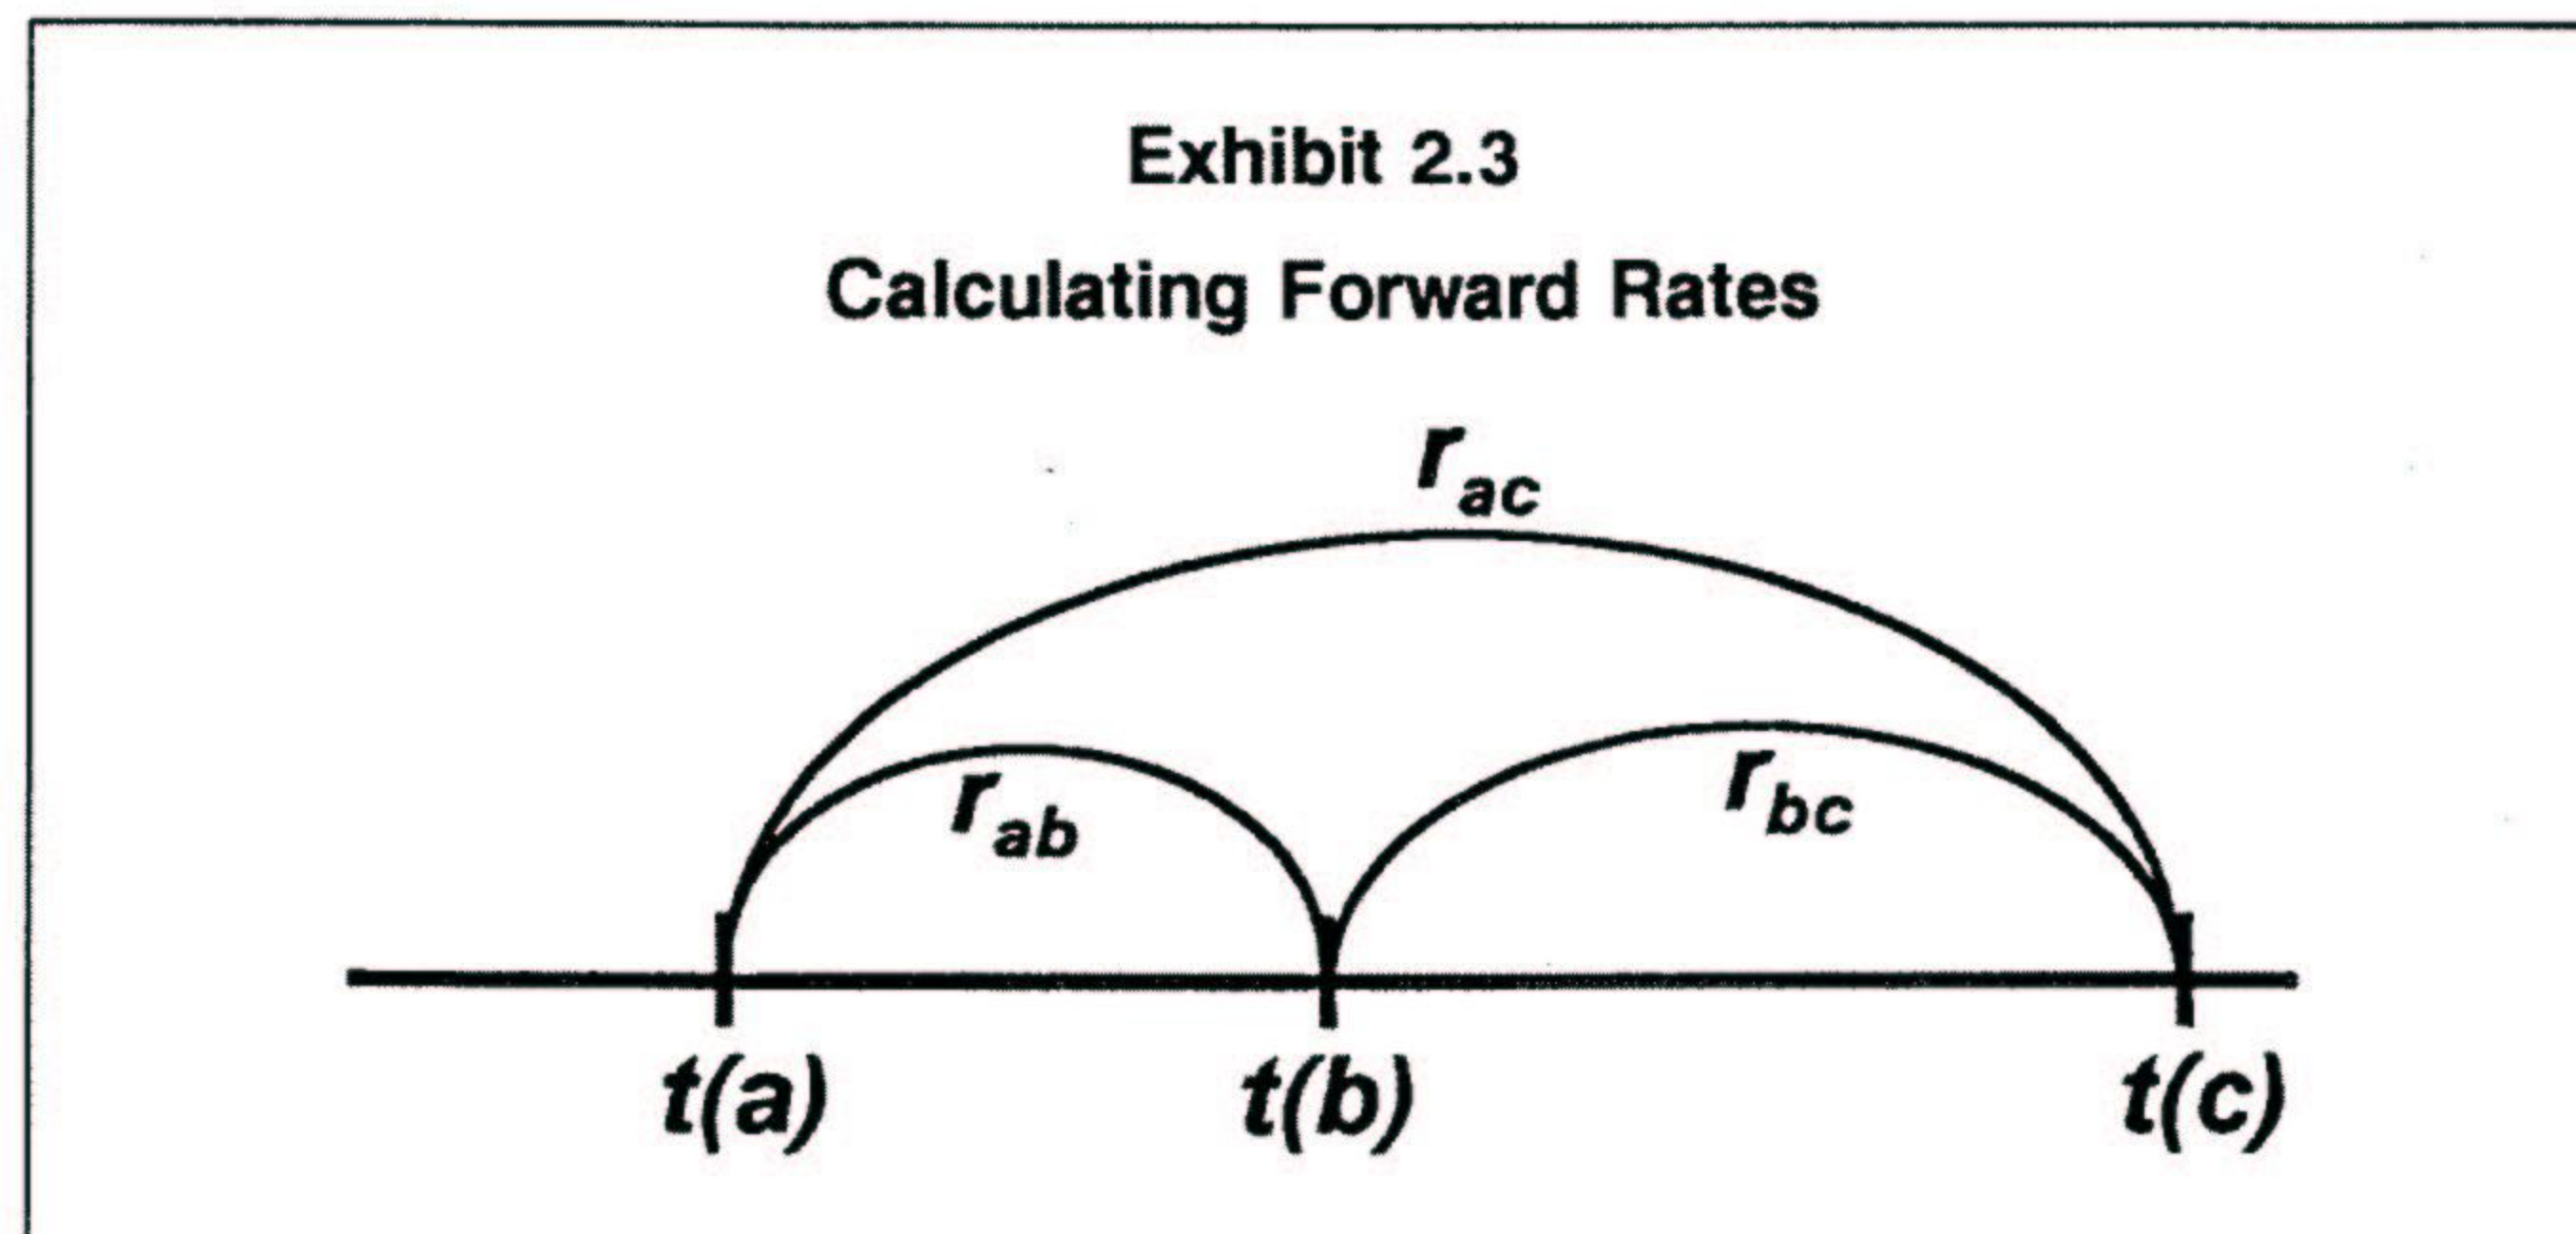
This need not always be the case. Where rates apply over a period that does not involve the present (or spot date), we refer to these rates as *forward rates*. An example—the three month rate in three months time is referred to as the three month forward. There is a deterministic relationship between spot and forward rates. This relationship stems from the principle that a cashflow should have the same NPV no matter how it is discounted. If this

does not hold, then there is a basis for increasing the NPV by revaluing the cashflow differently. Such an occurrence is called an *arbitrage* (see the section below). In the financial markets, professional arbitrageurs constantly exploit such occurrences. Because of this, they do not often exist, and if they do then it is only for very short periods.

If it is assumed (usually this is the case) that there is no arbitrage, then over any period, the value of a cashflow will be invariant whether it is discounted by a single rate or a series of forward rates. The general principle

$$\frac{1}{(1+r_{ac})^{t(ac)}} = \frac{1}{(1+r_{ab})^{t(ab)}} \times \frac{1}{(1+r_{bc})^{t(bc)}}$$

where r^{ij} is the rate applicable over period $t(ij)$.




In the market, there may be slight differences in value depending on the path for discounting or grossing up. Usually this will be less than the margin lost if the difference were to be exploited.

Example: forward rates

What is the forward rate from three months to six months given the following spot rates?

	Spot Rate
1 month	7.49
3 Month	7.75
6 Month	7.82



To solve for the forward rate from three to six months, we note that under the no arbitrage principal, one dollar in six months must have the same NPV if it is discounted by the six month spot rate, or by the forward rate from three to six months then the three month spot rate. If these are annual effective rates, then

$$\frac{1}{(1+r_{06})^{\frac{6}{12}}} = \frac{1}{(1+r_{03})^{\frac{3}{12}}} \times \frac{1}{(1+r_{36})^{\frac{3}{12}}}$$

The subscripts on the rates above refer to the start and end of the period over which the rates apply. Spot rates all have a subscript starting with zero (the spot date). Solving this gives $r_{36} = 7.8900\%$

Continuously compounded make calculations of forward rates extremely simple. In the example above, we can do the calculation using continuous rates.

	Spot Rate	Continuous
1 month	7.49	7.2228
3 Month	7.75	7.4644
6 Month	7.82	7.5293

The forward rate from three to six months in the continuous representation is given by

$$e^{-r_{06} \frac{6}{12}} = e^{-r_{03} \frac{3}{12}} \times e^{-r_{36} \frac{3}{12}}$$

This simplifies to

$$e^{-r_{36} \frac{3}{12}} = e^{-(r_{06} \frac{6}{12} - r_{03} \frac{3}{12})}$$

or

$$r_{36} = \frac{12}{3} \left(r_{06} \frac{6}{12} - r_{03} \frac{3}{12} \right)$$

The forward rate is a simple arithmetic expression, which gives the continuously compounded forward rate from three to six months as 7.5942%. Converted back to an annual effective rate we get 7.8900%, which is exactly the same as when calculated using the annual effective rates directly. Once the conversion to the continuously compounding domain is made, calculations are generally simpler.

Futures contracts are common manifestations of forward rates. In most markets participants have access to a strip of bill futures. These are usually three month instruments which start at various dates in the future. Forward rate agreements or FRAs are instruments which provide a guaranteed forward rate.

Digression: an arbitrage

The calculations from above are based on the principle of *no arbitrage*. To illustrate what happens when there is an arbitrage opportunity, we use the spot rates of the example above, but instead of solving for the forward rate of 7.89%, assume that there exists some instrument which will pay a rate of, say, 8.5% for the forward period from three months to six months.

	Rate
1 month	7.49
3 month	7.75
6 month	7.82
3-6m fwd	8.5

Using these rates, consider the value of a cashflow (say \$1,000,000) in six months time. If it were discounted by the six month rate of 7.82%, its NPV would be \$963,053. If the NPV were calculated by discounting from six to three months at the forward rate of 8.5%, then from three months to spot at the three month spot rate 7.75%, the NPV is now \$961,697 (or \$1,356 less than using the six month spot rate). If these rates were all available to a market participant, and the risk associating with borrowing or lending at them were equivalent, then an investor can borrow money for six months—the first three at the three month spot rate, the latter at the forward. This money would be lent at the six month spot rate. At the end of six months the investor would be \$1,356 better off, with no net outlay. This represents an arbitrage opportunity. In reality if such an opportunity were to occur, it would quickly be exploited. This would cause the rates to be adjusted until the arbitrage disappears.

Theoretically, there should be no difference to the NPV of a future cashflow, no matter what path the discounting follows. Over one year 365 one day rolls should be the same as one year roll and so forth. This forms the basis of arbitrage free pricing theory. It is a foundation for the pricing of many complex instruments, including options.

1.7 Discount factors

Whenever we use an interest rate, we need to specify the period over which the rate applies, and the basis in which the rate is expressed. This applies to both spot and forward rates. Often this leads to undue complexity. Another way to express rates is as *discount factors*. A discount factor or DF is just the value of one dollar when discounted over the required period. As an example, if the NPV of a dollar at some future date is 94 cents, then the discount factor at this date is 0.94. Discount factors have the advantage that they do not depend on any specific basis. Using discount factors means the complexity of keeping track of whether a rate is annual, semi-annual or continuous is no longer necessary. The only aspect that needs to be tracked is the period over which the discount factor applies. Discount factors can be related to both spot and forward rates.

For compounding rates the discount factor is


$$df = \frac{1}{(1+r)^t}$$

where r is the applicable rate over time t (r must be expressed in the same basis as t —for example, if r is semi annual, then t must be the number of semi annual periods over which r applies). Where the rate r is simple (annual), the discount factor is

$$df = \frac{1}{\left(1+r \frac{d}{365}\right)^t}$$

where d is the number of days over which discounting takes place. For continuous rates, the discount factor is $df = e^{-rt}$.

A curve can be drawn which gives the discount factor corresponding to any rate. These discount curves show the NPV of \$1 at any time when discounted by the applicable rate.



Discount factors are also applicable to forward rates. If r_{ab} is a forward rate from time t_a to t_b , then a forward discount factor df_{ab} can be generated. It provides the discounted value at t_a of one dollar at t_b . The relationship between spot and forward discount factors is straightforward.

$$df_{0b} = df_{0a} \times df_{ab}$$

where the subscripts are as described above.

1.8 Putting it all together—the term structure of interest rates

The interest rates used in the above sections are all considered in isolation. In the marketplace, we are constantly bombarded with different rates or types of rates. These are represented in a multitude of different bases. It is only meaningful to compare these rates if they are converted into a uniform basis, and have a consistent underlying type. When this is done, we obtain an *interest rate term structure* or *yield curve*. The interest rate term structure is just a time consistent view of interest rates. Usually it is a curve showing the rates which are effective over different time periods. For example, if we have a set of annual effective rates specified for periods of one, two, three years and so forth, then a curve through these is a representation of an interest rate term structure. Please refer to Chapter 3 for examples.

There is no single universal term structure. Rates differ depending on the currency, types of instrument and the way the rates are represented. Within a single currency, there are different term structures for classes of interest rates. The interest rate term structure may also be derived from market rates rather than using them directly. These aspects will be discussed further in the chapter on the yield curve.

1.9 Summary—interest rates


In this chapter so far we have covered various forms of interest rates. We have shown the different bases that rates can be expressed in, and how to convert from one base to another. The process of discounting and grossing up have been illustrated. These allow cashflows at any time to be valued at any other time. The concept of net present value shows the worth of future cashflows at the spot date. Continuously compounding rates have been introduced as a tool to simplify calculation. Discount factors are also used to represent the value of cashflows at various times. The difference between spot and forward rates is shown. Using the principle of no arbitrage, we can derive forward rates from spot rates and vice versa.

2. BOND PRICING, DURATION AND CONVEXITY

2.1 Bonds

A bond is a medium to long-term financial instrument. It is usually issued by a party in order to raise funds. Over the life of the bond, the issuer makes periodic interest payments. These are referred to as coupons. At maturity, the issuer is under an obligation to repay the original principal of the bond. The

life of a bond is usually not less than one year (bills or short-term notes are used for shorter periods). It is not uncommon for bonds to have maturities in excess of 20-30 years. Coupon payments are generally made quarterly, semi-annually or annually.



The value of a bond is determined by the size and timing of the cashflows. It is also highly dependent on the quality of the bond. For example, government or sovereign bonds are usually regarded as high quality instruments. They will be more expensive than lower quality bonds such as those issued by a corporation or an individual.

Some of the factors that affect the value of a bond include:

1. *Coupon size and timing*: the larger and more frequent the coupons, the greater the value of the bond.
2. *Maturity*: this is the period over which coupon payments are made.
3. *Issuer quality*: the risk associated with default is relatively lower if the issuer is of high quality. Such bonds will attract a premium.
4. *Current interest rates and outlook*: these affect the value of the bond on the secondary market.
5. *Liquidity*: there can be a premium for bonds that are easily tradeable.

Bonds are referred to as *fixed interest or fixed income instruments*. This is because, once the bond has been issued, all future payments are known. The only disruption to these will be if there is some sort of crisis event where the issuer cannot make an interest or principal payment, or a payment is delayed. The probability of such a default event is priced into most bonds. As this probability of default changes, the premium or discount associated with it will vary.

Many bonds, once issued, trade in the secondary market. Their market value will depend on prevailing economic conditions as well as the current state of the issuer. To enable bonds to trade, the market has developed conventions for their valuation. These all derive from the principle of discounting cashflows (which is discussed earlier in this chapter). This makes sense, as a bond is just a collection of cashflows. The only additional parameter is the risk to the bond holder associated with actually receiving these cashflows.

2.2 Pricing bonds

In the market place, bonds are usually quoted on the basis of either a yield, or a capital price. Either method can be derived from the other. The choice of method of quotation is just a convention of each particular market.

2.3 Bond yields

The yield of a bond—commonly referred to as its yield to maturity—is basically a measure of the return the bond holder can expect for the outlay involved in purchasing the bond. Yield is quoted as a rate in the same basis as the coupon payments that make up the bond. Most government and investment bonds pay semi-annual coupons—thus the quoted yield is semi-annual. There are a number of representations of the formula for pricing bonds. These generally differ in terminology only. They represent the net present value of the cashflows that make up the bond when they are all discounted by the bond yield.

Bonds are commonly priced as a value per \$100 of principal. If the current price is above \$100, the bond is said to be valued at a premium. If it is less, then the bond is at a discount. Whether a bond is at a premium or a discount depends on the relativity of the current yield to the coupon size of the bond. Where the yield is greater than the coupon, the bond will be at a discount. Where it is less, the bond is said to trade at a premium.

The price per \$100 principal is

$$P = v^{\frac{f}{d}} (c(x + a_n) + 100v^n)$$

where

- c = the periodic coupon payment in dollars per \$100 principal
 $c = \text{coupon}/(\text{coupons per annum})$
- n = the number of complete periods from the next coupon to maturity
- f = the number of days to the next coupon date
- d = the number of days from the last coupon date to the next
- v = $1/(1+i)$ where i is the periodic effective interest rate ($i = \text{yield}/\text{frequency}$ ie: $i = \text{yield}/2$ for semi annual coupon bonds)

$x = 0$ if ex-interest, $x = 1$ if cum-interest³

$$a_n = (1-v^n)/i$$

To value a bond, the above formula is used. Given a yield and the structural details, the value of the bond can be determined.

Example: pricing a bond

What is the price per \$100 of a bond with the following characteristics:

Maturity: 15 January, 2001

Settlement: 14 March, 1997

Coupon: 8.5% paid semi-annually

Yield to maturity: 7.14%?

The following quantities can be derived from the above data:

Next Coupon	15/7/97
Previous Coupon	15/1/97
i	3.570%
v	0.965531
n	7
a_n	6.098603
f	123
d	181
c	4.25

Substituting these into the bond formula, and using $x = 1$ (the bond is *cum-interest*), the price of the bond is \$105.844 per \$100 principal. This means that an investor would pay \$105.844 to receive semi annual coupons of \$4.25 every six months from 15 July, 1997. A final payment of \$104.25 (the \$100 principal plus a final coupon) is received by the investor on 15 January, 2001. At this stage, the bond ceases to exist.

The effect on the price of the bond at different yield levels can be seen in the table below.

Yield	Price
7.00%	\$106.321
7.14%	\$105.844
7.50%	\$104.628
8.00%	\$102.969
8.50%	\$101.343
9.00%	\$ 99.748
9.50%	\$ 98.185

As the yield increases, the cashflows are discounted at a higher rate. This means that the bond will have a lower net present value.

- Ex-interest means that the next interest payment is not included in the price of the bond. By convention, there is an ex-interest period before any coupon payment. The duration of this is set by market convention. If the bond trades during this period, the original holder, rather than the purchaser, still gets the coupon.

2.4 Characteristics of bonds

To understand the characteristics of a bond, it is useful to understand how the pricing formula works, and the sensitivity of the bond to various parameters. This allows the risks of the bond to be quantified and managed. It also allows the effect of market conditions to be observed and reacted to.

2.5 Principal and interest

As time progresses, a bond will accrue interest. At coupon dates, this interest is paid out, and the process repeats. The value of a bond can be split into principal and interest components. Interest accrues by the same amount each day. The daily accrual is equal to c/d using the terminology from the example above. Thus, on any day (before the bond goes ex-interest), the accrued interest is:


$$\text{Accrued Interest} = c \left[\frac{d-f}{d} \right]$$

The capital price is just the gross price less the accrued interest

$$\text{Gross Price} = \text{Capital Price} + \text{Accrued Interest}$$

For the bond in the example above, with $f = 123$ and $d = 181$, the accrued interest is \$1.362. The capital price is \$104.482.

Because of the process of accruing interest over a period, then paying it at coupon dates, the gross price of a bond will sawtooth with time if the yield is constant. The capital price will move much more smoothly.



In many markets, bonds are quoted using capital price rather than yield. From the example above, this bond would be quoted as having a (capital) price of \$104.482. The capital price is usually rounded to an even multiple (often eighths, sixteenths or thirty-seconds) of a dollar. The convenience of

doing this is that the settlement proceeds can easily and unambiguously be agreed upon, without the use of a pricing formula.⁴ In cases where bonds are quoted this way, participants in the market still need to obtain yields. This allows them to compare different bonds, and to gauge their returns. In order to calculate the yield given the price of a bond, the formula for the bond price needs to be solved iteratively. There is no closed form solution for the yield of a bond given its price.

2.6 The bond pricing formula


It is useful to examine the bond pricing formula in more detail. This gives an understanding of the fundamental principle of how bonds are valued. It provides a basis for valuing bond-like instruments, where there may be special conditions on the cashflows (for example, uneven coupon periods or partial principal repayment).

The bond formula discounts all cashflows at the yield to maturity. This is done in two stages. First, cashflows are discounted to the next coupon payment date. These are then discounted from the payment date to the settlement date. The term


$$v^{\frac{f}{d}} = \left(\frac{1}{1+i}\right)^{\frac{f}{d}}$$

is a discount factor from the next coupon date to the settlement date. The rest of the equation discounts all cashflows to this date.

4. The reason for using capital price to quote bond prices is mostly historical. Before computers or financial calculators were available, market participants rarely agreed on settlement proceeds when only a yield was agreed. To remove this problem, price was quoted directly. Accrued interest is easily and unambiguously determined. It is still the case that for complex instruments such as prepayable securities—where the notion of yield is ambiguous or dependent on other factors or assumptions—that capital price is a more convenient way to quote instruments.



Discounting each bond cashflow by the yield to maturity of the bond gives its market value in dollars. It assumes that each cashflow is reinvested at the yield to maturity. In reality, the cashflows will all be reinvested at rates available in the market. If a bond were issued with a ten year maturity, then the first cashflow would be paid nine and a half years to maturity, the second at nine years and so on. The first cashflow could thus be reinvested at a 9.5 year rate, the second at a 9 year rate. The penultimate coupon would be reinvested for just six months. The rates for all these periods are definitely not the same. They may not even equal the bond yield. To reflect a more accurate reinvestment assumption, the zero coupon curve is used. A zero coupon curve constructed out of bonds of similar characteristics would be appropriate for this purpose. The short end of the curve would reflect cash and bill rates, further out, the bonds would be used.



If the zero coupon curve is used, then each cashflow is discounted at the appropriate zero coupon rate. If the curve is constructed for bonds of the type being valued, then this price should be exactly the same as the market price of the bond calculated using the yield to maturity and the bond formula. If there is a difference between the price using the formula, and that using the zero coupon curve, it either represents a difference in the type of instrument being valued compared to those on which the curve is based,⁵ or a mispricing in the market.⁶

The zero coupon framework for valuing bonds is not used in the marketplace directly by traders. This is because it requires significant information in its derivation. It is well suited to risk management, relative value, arbitrage analysis and other purposes. It is also very useful for the valuation of non-standard bonds. At the expense of losing the convenience of using a single yield to maturity (or capital price), a more realistic reinvestment assumption can be applied. This is much closer to what would be attained in reality if all cashflows were to be reinvested.

When the zero curve is upward sloping, the yield to maturity will be lower than the zero coupon yield at maturity. This is because, when discounting on the zero curve, earlier cashflows are discounted at low yields. The latter ones will require a higher yield if the market price of the bond is to be retrieved. The converse is true for an inverse yield curve.

5. For example, a corporate bond valued on a government bond curve will have its price overstated. The government curve does not reflect the higher credit risk—and hence discount—associated with corporate bonds compared to government issues.
6. This may be where a bond is relatively cheap compared to similar bonds. If this occurs, an *arbitrage* opportunity may be exploited.

2.7 Risk parameters for bonds

The holder of a bond or bond portfolio is exposed to changes in the value of the instruments that make up the portfolio. These changes may be due to structural issues such as liquidity or the creditworthiness of the issuer. Other risks are market-related, such as the effect of changing yields, or variations in the interest rate term structure. These risks are quantified and managed by calculating various sensitivity parameters for bonds.

2.8 Risk management

An investor holding a bond or a portfolio of bonds is exposed to changes in value of the components of the portfolio. Unless the investor plans to hold all instruments to maturity, and is not reinvesting the interest cashflows from coupon payments, then exposure to prevailing rates and other market conditions must be considered. There are also structural exposures such as the creditworthiness of the issuer which will not be considered further here. The focus is on interest rate risk-related issues.

The fundamental measures which are used by most portfolio managers to quantify the interest rate risk or exposure in their portfolio is the sensitivity to rate or yield changes, and the timing or length of the cashflows in their portfolio. The former is usually quantified by calculating the sensitivity to a yield shift of one basis point. This is referred to in the market as the PVBP or *present value per basis point*. The PVBP is calculated for bonds by shifting the yield to maturity and measuring the dollar change in the bond price. It is usually quoted as an amount per million dollars face value of the bond.⁷ A measure commonly used to express the length of a bond or portfolio is the *modified duration*.


2.9 Duration and modified duration

Measures of duration are useful in that they give a simple means for judging the length of time exposure of a bond or portfolio. These measures aggregate all the bond or portfolio cashflows. Because of this aggregation, they give no information about the structure and timing of cashflows. They are a first order measure, whose use is widespread in the financial industry.

Duration is a concept introduced by Frederick R Macaulay and is one that bears his name.⁸ *Macaulay duration* is essentially the time weighted average of the cashflows of a bond. Graphically, it is illustrative to consider the duration as the fulcrum point on a timeline where the cashflows balance. At the duration point, exactly half the dollar value of a bond will have been paid when referenced to the present.

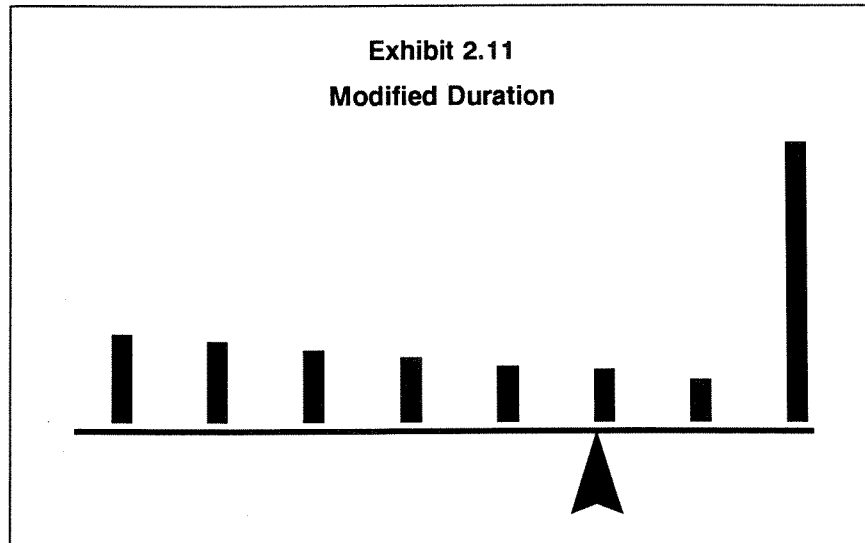
7. In some markets, the PVBP is referred to as the DV01 or *dollar value of one basis point*.

8. For examples, see Das, *Swaps and Financial Derivatives* (2nd ed, 1994), pp 1055-1060, and references contained therein.



Duration is useful in that it illustrates where the cashflows occur. A zero coupon bond would have a duration exactly equal to its time to maturity. Using the fulcrum analogy, the pivot point would be where the cashflow occurs. If coupons are included, then, as the coupon size increases, the fulcrum would need to be brought closer to the settlement date. An annuity (a stream of equal cashflows) would have the fulcrum halfway to maturity if the first cashflow were on the settlement date. There is no effect on duration calculations due to the yield to maturity or reinvestment of coupons.

A more useful duration-type measure is the *modified duration*. Extending the balance analogy above, this is the pivot point where the net present value of the cashflows balance. The modified duration is the point where an investor would receive half the present market value of a bond.



Under the discounting process, cashflows at later dates will be worth less proportionally than near flows. This serves to bring the balance point closer to the settlement date. The modified duration is less than the Macaulay duration for any coupon bond (they are equal for a zero coupon bond).

Numerically, both modified and Macaulay duration are simple calculations.

$$DMac = \frac{\sum_i c_i t_i}{\sum_i c_i}$$

where

$DMac$ = Macaulay duration

c_i = cashflow that occurs at time t_i .

$$DMod = \frac{\sum_i t_i \frac{c_i}{(1+r)^{t_i}}}{\sum_i \frac{c_i}{(1+r)^{t_i}}} = \frac{\sum_i t_i \frac{c_i}{(1+r)^{t_i}}}{price}$$

$DMod$ is the modified duration. All symbols are as above, and r is an appropriate rate by which the cashflows are discounted. The *price* can be either the capital price or the gross price of the bond. The modified duration will vary depending on which is chosen.⁹ If the modified duration is calculated using zero coupon rates then r is replaced by the zero coupon rate

9. *Capital price modified duration* is often chosen as its behaviour with time is more regular than when gross price is used. This is because of the discontinuity in the gross price of a bond when a coupon is paid. When this occurs, the gross price modified duration will increase discontinuously.

appropriate to each cashflow. Modified duration based on yield to maturity will differ slightly from that based on zero coupon rates.

Example: duration and modified duration

What is the duration and modified duration of the following bond

Maturity: 25 January, 2001

Settlement: 25 June, 1996

Coupon: 8.25% paid annually

Yield to maturity: 7.90%?

Date	Time (y)	Cashflow (c _i)	c _i t _i	Discounted Value (DV _i) _i	DV _i t _i
13/06/96	0	-	-	-	-
25/01/97	0.619	8.250	5.108	7.871	4.873
25/01/98	1.619	8.250	13.358	7.294	11.811
25/01/99	2.619	8.250	21.608	6.760	17.706
25/01/00	3.619	108.250	391.776	82.209	297.528
Total		133.000	431.851	104.134	331.918

This gives a value for the Macaulay duration of $431.851/133.0 = 3.247$ years, and a modified duration of $331.918/104.134 = 3.187$ years. Note that the bond price from above is \$104.134.¹⁰ The maturity of the bond is 3.619 years. As the coupon size decreases, the duration of the bond will approach maturity.

Exhibit 2.12

Effect of Coupon Size on Duration and Modified Duration

Coupon	DMac	Dmod	Maturity
0%	3.619	3.619	3.619
2%	3.508	3.488	3.619
4%	3.412	3.376	3.619
6%	3.329	3.280	3.619
8%	3.256	3.197	3.619
8.25%	3.247	3.187	3.619
10%	3.191	3.124	3.619

10. Using the bond formula, the price is \$104,147. This difference is exactly one day's worth of accrued interest. This is because of the leap year 1996, which means an extra day's accrual. This is picked up in the formula, but not by direct discounting of cashflows.

2.10 Duration of portfolios

As for a single bond, duration and modified duration can be calculated for portfolios. The method of calculation is exactly as for the single bond case. A portfolio duration is just the weighted average of the component durations.

$$Duration (total) = \frac{\sum Duration_n price_n}{\sum price_n}$$

This applies for both Macaulay and modified duration (the two must not be combined).

It is useful to know the (modified) duration for a portfolio, as this gives an approximation for its aggregate life. Duration is important for portfolio managers who are benchmarked against bond indices.

2.11 Interest rate sensitivity—PVBP

When yields (or zero coupon rates) change, bonds change in value. The PVBP quantifies sensitivity to this change. Usually this is done by shifting the yield and valuing the bond. The difference between the shifted and original price suitably normalised gives the PVBP. When hedging bonds with futures or switching from one bond to another, traders often make sure that they stay PVBP matched. This means that their sensitivity to yield changes (assuming that the yields all change by the same amount) is not altered, provided the changes are small.

The theoretical value for PVBP is

$$PVBP = \frac{\partial P}{\partial i}$$

This can be related to the modified duration:

$$DMod = \frac{PVBP}{Capital Price} = \frac{1}{Capital Price} \frac{\partial P}{\partial i}$$

Modified duration is a useful risk management measure. It gives an aggregate length for a bond or portfolio. The interest rate sensitivity will be similar to that for a zero coupon bond maturing at the modified duration. The sensitivity will not be exactly the same as the cashflow timing can be very different. This is especially prevalent if a portfolio contains bonds across the complete maturity spectrum. The similarity is only first order. Portfolio managers are often set modified duration limits within which their portfolios must be maintained. This sets the basic risk profile of the portfolio, whilst still leaving the manager the flexibility to choose specific instruments or maturity profile.

2.12 PVBP for zero coupon bonds

The value of a zero coupon bond is

$$\frac{\$100}{(1+r)^t}$$

So, per \$100:

$$PVBP \approx 100 \left[\frac{1}{(1+r)^t} - \frac{1}{(1+r+\Delta r)^t} \right]$$

where

Δr is a perturbation to the zero coupon rate.

Using a Taylor series expansion of this, it can be shown that $PVBP \propto t$. As the length of the zero coupon bond increases, its PVBP will increase proportionally. This is a useful rule of thumb. Most coupon bonds pay coupons substantially less than the final payment at maturity, so the PVBP for such bonds also increases with time to maturity.

2.13 Convexity

Modified duration and PVBP are useful measures of the sensitivity of bonds to changes in the underlying yield. They are calculated for small yield changes. What happens when the change in yield is large? It is normally the case that most of the time markets only move a few basis points from day to day. Occasionally, there are shocks where the market can move tens or even hundreds of basis points. What happens to the value of bonds under such moves?

The relationship between bond price and yield is not linear. A yield move up will not produce the same magnitude of price change as the same move downwards. The price change for a ten basis point move will not be exactly ten times more than for a one basis point move. The PVBP is not constant over all yields. Convexity is a measure of how the PVBP changes with yield. It is a measure of the non-linearity of the price behaviour of bonds:

$$Convexity = \frac{\partial^2 P}{\partial i^2} = \frac{\partial PVBP}{\partial i}$$

Convexity becomes useful when bonds with similar PVBP are being compared. A larger convexity implies that the PVBP will change more with yield than does a lower convexity. This is a useful property, and can provide benefit to a portfolio manager. This will be illustrated in the section on portfolio management.




Exhibit 2.14
Convexity for a 20% Coupon Bond

Yield	Yield Change (bp)	Gross Price (\$)	Price Change (\$)	PVBP per \$m
10.00%	—	181.004	—	—
10.01%	1	180.888	0.116108	116.108
10.10%	10	179.848	1.155831	115.583
11.00%	100	169.950	11.053433	110.534

The more convex a bond or portfolio is, the more the PVBP will change as yield changes become large. This has major implications when hedging bond portfolios.

2.14 Evaluating Duration, PVBP and Convexity

Although there exist derivative formulae for modified duration, PVBP and convexity, in practice it is more convenient to evaluate them numerically. This also has the advantage of quantifying them over a yield change that is realistic in the market.¹¹ The following are suggested for these calculations

$$PVBP = \frac{1}{2\Delta i} [P(i - \Delta i) - P(i + \Delta i)]$$

11. The market is more likely to move, say, 1, 2 or 5 basis points than the infinitesimal move that the differential formulae apply to.

$$\text{Convexity} = \frac{1}{\Delta i^2} [P(i - \Delta i) - 2P(i) + P(i + \Delta i)]$$

$$D_{\text{Mod}} = \frac{PVB P}{\text{Capital price}}$$

The bond price needs to be calculated at only three yield levels to obtain all these quantities. This is simple to implement, and gives values that are relevant to market moves.

2.15 Time effects

As a bond moves towards maturity, it accrues interest. On coupon payment dates, this interest is shed. The capital price will also change. When the bond matures, its capital price will be exactly \$100. If a bond is trading at a premium, its capital price will decrease; while the converse is true for discount bonds. This drift towards a capital price of \$100 is called the *pull to par*.

Exhibit 2.15
Pull to Par for Various Coupon Bonds (Yield 8% Semi-Annual)

Maturity (y)	Coupon		
	6%	8%	10%
10	86.410	100.000	113.590
8	88.348	100.000	111.652
6	90.615	100.000	109.385
4	93.267	100.000	106.738
2	96.370	100.000	103.630
0	100.000	100.000	100.000

Over time, components that change the price of a bond are the pull to par, interest accrual and yield changes. It is useful to split any change into these components. Pull to par depends on the difference between the coupon and the yield to maturity. Where this difference is large, the pull will be greater. The pull to par also increases as the bond gets closer to maturity.

2.16 Non-standard bonds

2.16.1 Constant accrual bonds

It is often the case that a bond is not exactly the same as the standard bonds described above. The standard bond pays a fixed coupon every period. Interest accrues each day, and is paid at the coupon date. The coupon size is always the same. If there are differences in the number of days per period (there invariably are as a standard year has an odd number of days), the interest accrued per day will differ slightly.

With constant accrual bonds, the daily accrual is specified. This means that the actual interest payment will differ depending on the number of days in the period.

Example: constant accrual bond

A bond pays a coupon—nominally 10% per annum. The coupon is paid semi-annually, with the same accrual per day.

With this bond, the accrual is based on a standard year of 365 days. The daily accrual (per \$100 face value) is $\$10/365 = \0.027397 per day. If the coupon periods are 182 days and 183 days long in a normal year, then the interest payments are \$4.986 and \$5.014 respectively. In a leap year, both periods have 183 days. The payments are 5.014% for both periods. In leap years, the bond pays \$10.028 in interest.¹²

Pricing of these bonds is commonly done with the bond formula modified slightly for the first coupon.

$$P = v^d \left((c_n dx + ca_n) + 100v^n \right)$$

Here c_n is the daily interest accrual. The next coupon will be $c_n d$. Pricing assumes that subsequent coupons are all c . Other terms are as per the standard formula. Technically, the value of each coupon should be accrued separately. This is not done, as it will make only a small difference to the bond price, with a large amount of extra complexity.

2.16.2 Bonds with long or short first coupon

The coupon of a standard bond is paid periodically. It is often the case that, when a bond is first issued, the time to the first coupon payment is not an exact period. In this case, the first coupon will accrue interest over either more or less days than if the bond were continuing from a normal payment. In such cases, the initial coupon is usually based on a daily accrual (as described in the section above). The term d (the number of days in the period) is adjusted for the non-standard period. Once this coupon is paid, the bond is priced as a standard bond.

$$P = v^{d1} \left((c_n d1.x + ca_n) + 100v^n \right)$$

for the first period of $d1$ days, after which the standard formula is used.

Example: Short first period bond

A bond paying an annual coupon of 8.25% each year on 15 July is initially issued on 1 March. How much will the first coupon payment be?

The standard accrual for this bond is $\$8.25/365 = \0.022603 per day. The first period is 136 days long, so the first coupon will be \$3.074 instead of \$8.250. After this is paid, all subsequent coupons will be \$8.25.

12. For a conventional bond, the interest payment would be \$5.00 irrespective of the period length. In 183 day periods, the daily accrual is $\$5.00/183 = \0.02732 , and for 182 day periods it is \$0.02747.

2.16.3 Short/long last period

As with a non-standard first period, the last coupon may be delayed or paid early. The final principal repayment may also be shifted. In this case, the formula is adjusted for this. If the last coupon is shifted by l days, and the final payment is shifted by m days, the standard bond formula becomes:

$$P = v^{\frac{f}{d}} (c(x + ca_{n-1}) + cv^{(n+l/365)} + 100v^{(n+m/365)})$$

2.16.4 Bonds with amortising principal

The bonds considered so far all pay periodic interest. The principal is paid at maturity. It is often the case that some or all of the principal will be repaid during the life of the bond. If the repayment schedule is known, the bond can be priced as a series of bonds of different periods. This is best illustrated by example.

Example: An amortising bond

Consider a bond with the following characteristics:

Maturity	5 years
Coupon	10%
Frequency	Semi-Annual
Amortisation	20% every year

The cashflows per \$100 principal for this bond are

Period	Interest	Principal Paid	Remaining Principal	Total Cashflow
1	\$ 5.00		\$100.00	\$ 5.00
2	\$ 5.00	\$ 20.00	\$100.00	\$ 25.00
3	\$ 4.00		\$80.00	\$ 4.00
4	\$ 4.00	\$ 20.00	\$80.00	\$ 24.00
5	\$ 3.00		\$60.00	\$ 3.00
6	\$ 3.00	\$ 20.00	\$60.00	\$ 23.00
7	\$ 2.00		\$40.00	\$ 2.00
8	\$ 2.00	\$ 20.00	\$40.00	\$ 22.00
9	\$ 1.00		\$20.00	\$ 1.00
10	\$ 1.00	\$ 20.00	\$20.00	\$ 21.00

To price this bond, we can split it into a series of five bonds each of face value \$20. These bonds mature at the end of years 1, 2, 3, 4 and 5 respectively. The cashflows for these bonds are as follows:

Period	Bond 1	Bond 2	Bond 3	Bond 4	Bond 5	Total Cashflow
1	\$ 1.00	\$ 1.00	\$ 1.00	\$ 1.00	\$ 1.00	\$ 5.00
2	\$ 21.00	\$ 1.00	\$ 1.00	\$ 1.00	\$ 1.00	\$ 25.00
3		\$ 1.00	\$ 1.00	\$ 1.00	\$ 1.00	\$ 4.00
4		\$ 21.00	\$ 1.00	\$ 1.00	\$ 1.00	\$ 24.00
5			\$ 1.00	\$ 1.00	\$ 1.00	\$ 3.00
6			\$ 21.00	\$ 1.00	\$ 1.00	\$ 23.00
7				\$ 1.00	\$ 1.00	\$ 2.00
8				\$ 21.00	\$ 1.00	\$ 22.00
9					\$ 1.00	\$ 1.00
10					\$ 21.00	\$ 21.00

Each of these is a standard bond of \$20.00 face value. The price of the complete amortising bond is the sum of the prices of each of the component bonds.

Any fixed amortisation of principal can be priced by decomposing the bond into a series of standard bonds of different maturities.

2.16.5 Unknown amortisation of principal

There is a major class of amortising securities where the exact schedule for repayment of the principal is not specified when the bond is issued. There will probably be minimum and maximum boundaries for repayment, but within these, repayment is not known. Such instruments are usually asset-backed securities. These are based on interest payments provided by asset holders. The most common asset-backed securities are mortgage-backed bonds. A mortgage-backed bond is funded by mortgage repayments from property owners. Because a property holder can repay principal at a rate faster than the minimum required under the loan agreement, the holder of a bond backed by these repayments has an uncertain schedule for principal repayment. There are many different structures for mortgage-backed securities. These will not be covered here.

In order to price securities where the principal repayment schedule is unknown, the repayment rate needs to be modelled. This can be a very complex process. In the case of mortgage-backed securities, some models will examine the underlying pool of assets, the demographics of the mortgagors and the prevailing economic environment.

Asset-backed securities that are prepayable contain an additional risk to the investor over the normal risks associated with bonds. This is prepayment risk. This is because the actual repayment rate will differ from that used in the model by which prepayable bonds are priced. If these securities are trading at a premium, then where prepayments are faster than anticipated, the bond holder will be penalised. Where the bond is at a discount, faster repayment will be of benefit to the holder.

Chapter 3

Interest Rate and Yield Curve Modelling

*by Satyajit Das (with a contribution from Roger Cohen)**

1. INTRODUCTION

Interest rates and the process of discounting future cash flows to price and value financial transactions is fundamental to capital markets. Accurate, consistent and reliable interest rates are therefore essential to all financial transactions. This is true irrespective of the type and complexity of the instrument.

In essence, interest rates are the pure price of time designed, through the discounting process, to equate cash flows occurring at different future dates to facilitate valuation and comparison of different transactions. In practice, interest rates are used for the following range of transactions:

1. Pricing and valuation of financial instruments or transactions—entailing the valuation of instruments, such as bonds or derivatives on fixed income instruments, by allowing analysis of the returns from different sets of cash flows through comparison of their discounted present value. In addition, the use of interest rates to value and price derivatives in other asset classes, such as currency, equity, and commodities.
2. Relative value and arbitrage—covering the use of interest rates and discounting to assess the relative value of traded or untraded instruments and to identify arbitrage opportunities.
3. Assessing or forecasting economic expectations—involving the analysis of various types of information available from the yield curve, such as forward rates, to assess market expectations of the path of future interest rates and the term structure of future interest rates which allow the formation and testing of expectations about future economic activity and inflation rates.

However, the process of deriving interest rates or discount factors is far from simple and unambiguous. A whole body of work has developed to assist in the determination of the interest rates to be utilised for the various identified purposes. In this chapter the basic issues relating to the derivation of interest rates and yield curves for use in the pricing and valuation of transactions is examined.

The structure of this chapter is as follows: the interest rates to be utilised are first considered, including identification of the concept of discount factors and the various types of interest rates (par, forward and zero) and their interrelationship, as well as the calculation of forward and zero coupon rates from the available yield curve. The second part of the chapter deals with the

*This chapter is written by Satyajit Das. Roger Cohen contributed Exhibits 3.14 and 3.21.

problems of deriving a suitable yield to enable the calculation of the various interest rates, including approaches to interpolation (linear and splines) and the issues in practical curve construction. The chapter concludes with a review of current best practice in interest rate and yield curve modelling.

2. INTEREST RATES

2.1 The concept of interest rates and discount factors

The concept of interest rates and discount factors or present value are interrelated. Central to the concept of interest rates and discount factors is the fact that value in financial transaction is given by cash flow, which is defined in terms of two vectors: amount and the time at which the cash flow occurs.¹

The concept of discounting these cash flows which may occur at different points of time is designed to enable the value of these individual items to be calculated at a determined point of time (for example, today) to allow comparability. In essence, this requires the cash flow to be moved in time to determine an *equivalent* cash flow as at the relevant date. Using the fundamental homogeneity and uniformity of cash, the current value of cash can be given by its present value, which is intuitively an amount which, if invested at the relevant interest rate, will give a value equivalent to the stated cash flow as at the date on which the cash flow occurs in the futures.

This can be stated more precisely as:

$$C_{t_0} = C_{t_1} * DF_{t_1}$$

Where

C_{t_0} = Cash flow at time t_0

C_{t_1} = Cash flow at time t_1

DF_{t_1} = Discount factor (or present value) for cash flows as at time t_1 .

The discount factor (DF) generated or utilised to calculate a discount factor is essentially the present value of \$1 at a specific future time. In theoretical terms, this is merely the price of the relevant zero coupon bond, discounted at the zero coupon rate (or, if appropriate, the par or coupon yield to maturity).

Exhibit 3.1 sets out mathematically the relationship of discount factors to the relevant interest rate for each future period. Discount factors, under conditions of positive interest rates, will be less than 1 and greater than 0 and are inversely related to yield with reference to maturity. *Exhibit 3.2* sets out the shape of the interest rate curve and the discount rate curve.

1. To be strictly accurate an additional vector which is required to define value is any contingency or conditionality relating to the cash flow, eg, in the case of an option.

Exhibit 3.1**Discount Factors and Interest Rates****1. Simple Interest**

$$DF = 1 / (1 + R_{t1} * t/N)$$

Where

DF = Discount factor for time t1 at rate R_{t1}

R_{t1} = Interest rate as at time t1

t = Number of days (or t1 - t0)


N = Number of days in a year (360 or 365)

2. Compound Interest

$$DF = 1 / (1 + R_{t1})^{(t/N)}$$

3. Continuously Compounded Interest

$$DF = e^{-R_{t1} * t/N}$$



The advantage of discount factors is that each discount factor is unique. It is not affected by market quotation conventions prevalent in relation to interest rates such as the periodicity of compounding (annual, semi-annual, quarterly or continuous) or the day count basis utilised (actual/365, actual/360 or bond basis).

In effect, the problem of yield curves can be expressed as the problem of either deriving term structure of interest rates or discount factors.

2.2 Types of interest rates

2.2.1 Overview

There are in practice three separate types of interest rates:

1. *Par rate*—which is defined as the interest rate on a coupon paying instrument out of today which is the standard interval rate of return formulae which discounts all payments on a coupon bond or instrument at the same interest rates.
2. *Forward rate*—which is defined as the interest rate on a coupon paying instrument out of a nominated date in the future.
3. *Zero rate* (also known as zero coupon rates, spot rates or pure interest rates)—which is defined as the interest rate on an instrument which pays no coupon and entails the exchange of a cash flow today for another (larger) cash flow at a nominated future date.

The par rate is the only observable interest rate in markets. Forward rates and zero rates are more difficult to observe directly. However, in jurisdictions where there are traded markets in interest rate futures and/or zero coupon securities,² it may be possible to directly observe certain forward and zero rates.

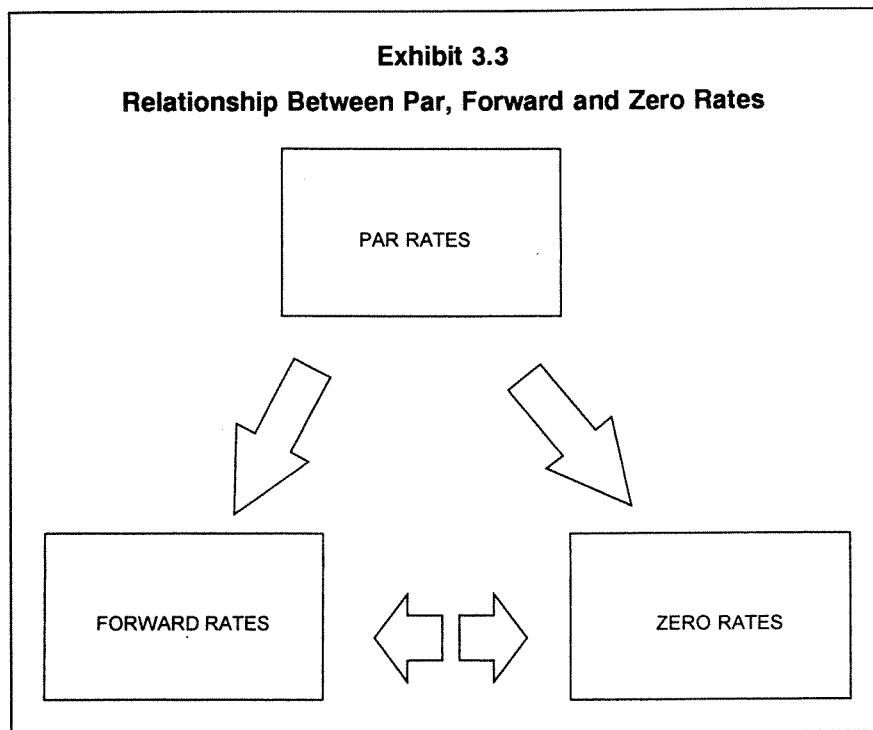
The three types of rates are clearly and unambiguously interrelated as they are, in effect, different perspectives on the same set of interest rates. In practice, the par interest rates which are observable are utilised to calculate the implied forward and zero rates.

The formal interrelationship between par, forward and zero rates can be stated as follows:

- forward interest rates are the interest rates at which par yields are reinvested;
- zero rates are the par interest rates with the reinvestment risk removed; and
- forward rates as between two points in time are implied by the zero rates at those two points.

Exhibit 3.3 sets out diagrammatically this relationship.

2. For example, the US Treasury STRIPS market which allows the unbundling of a bond into individual coupon and principal components and allows separate trading in these components.



The difference between the rates (in particular, par and zero rates) can be seen from a consideration of their role in valuation. Valuation of all financial transactions assumes the use of identified and specific interest rates or discount/present value factors to discount or present value cash flows identified with individual transactions. The two alternative types of interest rates available are:

1. par rates to maturity; and
2. zero rate to maturity.

As noted above, the par rate to maturity is usually directly observable, being the market quoted rate for the relevant securities of the required maturity. In the case of derivative or swap transactions, the relevant par rate to maturity is, typically, the quoted swap rate for the relevant maturity or, if unavailable, the interpolated yield based on available swap yield curve information. In contrast, the zero rate is not directly observable and is usually estimated from the existing par rate curve for the relevant instrument.

Traditionally, financial instruments have been valued utilising par yield to maturity. However, use of par rates creates a number of problems:

- coupon effect;
- assumptions on reinvestment rates; and
- absence of an unambiguous and unique interest rate for each maturity.

The coupon effect refers to the phenomenon observed in markets that the par interest rates of bonds or other financial instruments with the same maturity but different coupons may vary significantly. These differences may

be caused by factors such as the interest rate risk or differential interest rate volatility of the securities, tax or clientele effects.

Utilisation of the par rate implies that the actual realised return only equals the normal redemption par yield to maturity if reinvestment rates on all intermediate cash flows, typically the coupons, are actually equal to the redemption yield. The realised yield, therefore, would only be equal to the par yield where the security is a zero coupon security, that is, a security which has no intermediate cash flows, as there is no potential reinvestment risk in the transaction. In practice, reinvestment rates on coupon cash flows will not equal the redemption yield. Theoretical forward rates are the only true measure of available reinvestment rates, and even then, the forward rates implicit in the yield curves at any point in time do not guarantee that these reinvestment rates are actually achieved.

In addition, the use of coupon of par rates creates an ambiguous relationship between yields and maturities. The use of par rate technology does not facilitate the identification of an *unique* interest rate and, by implication, discount factor for a particular maturity.

For example, assume the following yield curve exists:

Maturity (years)	Par yield % pa
0.25	7.25
0.50	7.55
1.00	7.92
1.50	8.23
2.00	9.05

Under these circumstances, a two year security, which pays intermediate coupons, say, every six months, will be valued by discounting all payments at 9.05% pa. However, for an identical security, with a maturity of 1.5 years, all cash flows, including intermediate coupons, would be discounted at a different rate, namely, 8.23%. Consequently, the rates applicable for years 0.5, 1.00 and 1.5 can be, either, 9.05 or 8.23% pa depending, solely, on the final maturity of the security. Because of these problems, par rate to maturity valuation does not imply an unambiguous relationship between the interest rate and the relevant maturity.

The identified problems of par rate to maturity technology are substantially overcome by utilising zero rates. As noted above, the zero rate can be defined as the interest or discount rate which applies between a cash flow now and a cash flow at a single date in the future, which is equivalent to the yield on a pure discount bond or zero coupon security (hence, the reference to zero rate). Utilising zero rates allows, for example, a two year yield to be directly related to a pure two year security, being a pure zero coupon security with a single cash flow in two years time.

The zero rate eliminates the coupon effect. The use of zero rates to discount or present value cash flows does not involve any assumptions as to the reinvestment rate applicable to any intermediate cash flows. In addition, the zero coupon rate has the advantage that each maturity is identified with a single unambiguous interest rate, being the rate of a pure single payment instrument. These factors allow zero rates to be utilised to value and ultimately manage entire portfolios of financial instruments (bonds and

derivatives) as a series of cash flows, each of which is valued at a unique rate.

2.3 Derivation of forward rates

Forward interest rates can be calculated from the current yield curve. If suitably spaced yields and either synthetic or actual securities are available, then forward rates can be estimated.

The forward rates can be calculated based on the theoretical construct that securities of different maturities can be expected to be substitutes for one another. Investors at any time have three choices. They may invest in an obligation having a maturity corresponding exactly to their anticipated holding period. They may invest in short-term securities, reinvesting in further short-term securities at each maturity over the holding period. They may invest in a security having a maturity longer than the anticipated holding period. In the last case, they would sell the security at the end of the given period, realising either a capital gain or a loss.

According to a version of the pure expectations theory of interest rate term structure (see discussion below), investors' expected return for any holding period would be the same, regardless of the alternative or combination of alternatives they chose. This return would be a weighted average of the current short-term interest rate plus future short rates expected to prevail over the holding period; this average is the same for each alternative.

Forward rates may be calculated from the currently prevailing cash market yield curve, as any deviation from the implied forward rates would create arbitrage opportunities which market participants would exploit. This arbitrage is undertaken by buying and selling securities at different maturities to synthetically create the intended forward transaction. By simultaneously borrowing and lending the same amount in the cash market but for different maturities it is possible to lock in an interest rate for a period in the future. If the maturity of the cash lending exceeds the maturity of the cash borrowing the implied rate over the future period, the forward-forward rate, is a bid rate for a forward investment. Similarly, if the maturity of the cash borrowing exceeds the maturity of the cash lending, then the resulting forward-forward rate is an offer rate for a forward borrowing. This process of generating forward rates is set out in *Exhibit 3.4*.

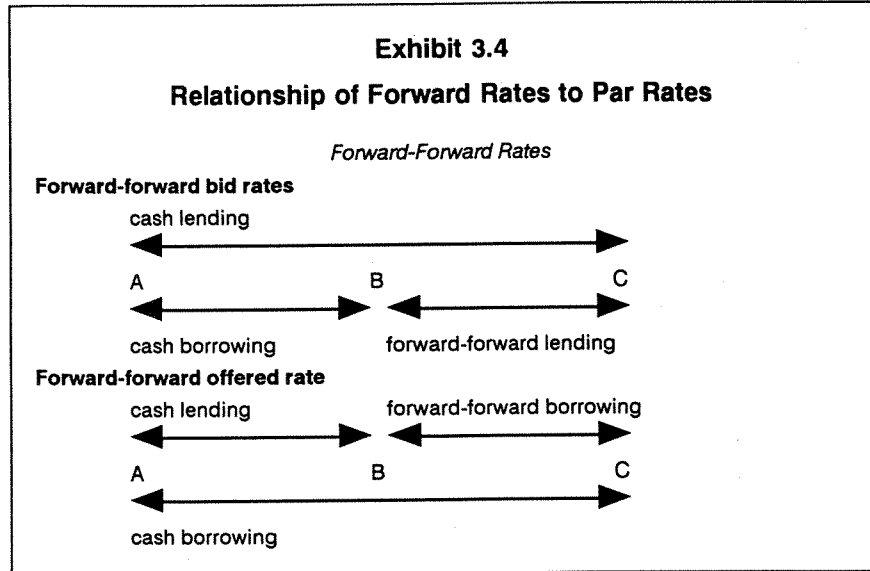


Exhibit 3.5 sets out the mathematical relationship between par interest rates and forward rates. *Exhibit 3.6* sets out examples of calculating forward rates.

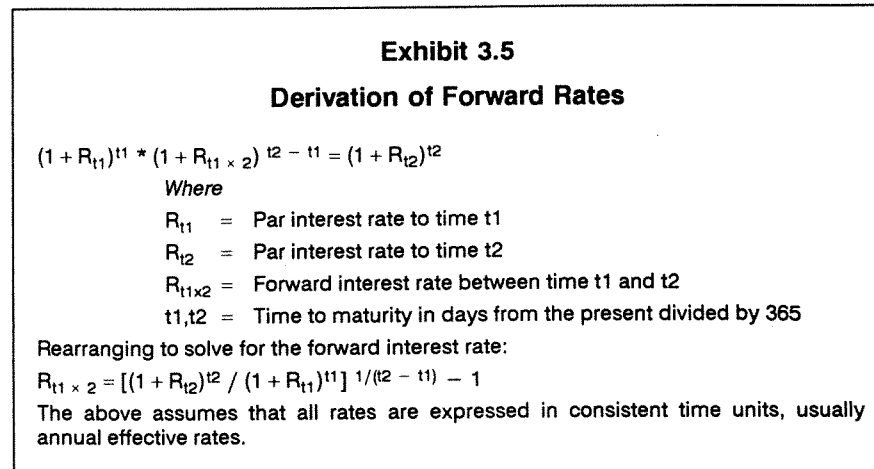


Exhibit 3.6
Derivation of Forward Rates—An Example

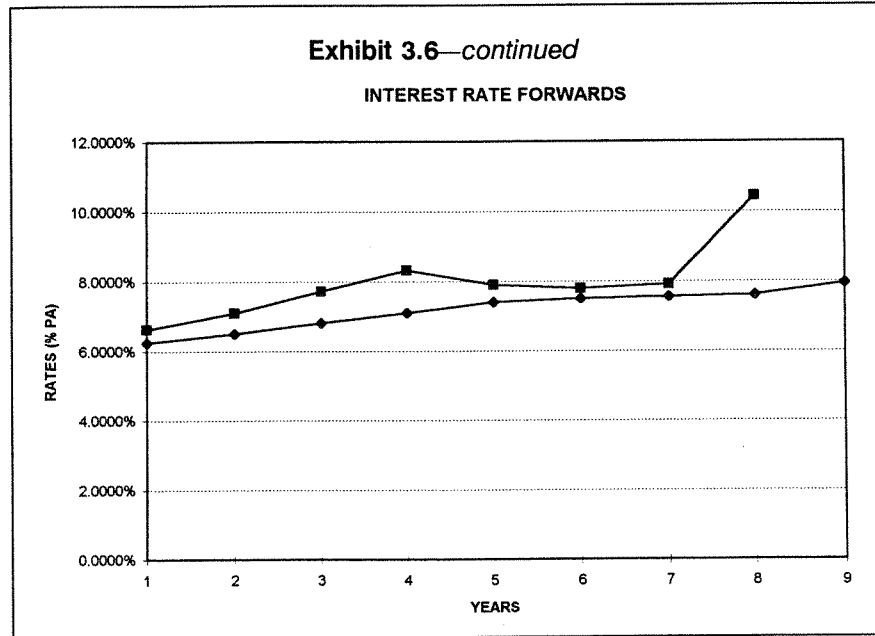
23-Jul-96

CALCULATION OF FORWARD/FORWARD INTEREST RATES

DAYS	(3) RATES (ANNUAL)	(4) (1+R) ^{1/2}	(5) (1+R) ^{1/11}	(6) (4)/(5)	(7) 1/(2-1)	FORWARD RATES FOR PERIOD	
						(6) ^{1/11}	(6) ^{1/11} - 1
0							
31	6.2500%	1.01600	1.00516	1.01078	5.98361	6.627%	6.627%
92	6.5000%	1.03372	1.01600	1.01744	3.96739	7.101%	7.101%
184	6.8000%	1.05284	1.03372	1.01850	4.05556	7.716%	7.716%
274	7.1000%	1.07400	1.05284	1.02010	4.01099	8.308%	8.308%
365	7.4000%	1.09478	1.07400	1.01934	3.96739	7.898%	7.898%
457	7.5000%	1.11570	1.09478	1.01911	3.96739	7.799%	7.799%
549	7.5000%	1.13682	1.11570	1.01884	4.05556	7.906%	7.906%
639	7.6000%	1.16532	1.13682	1.02507	4.01099	10.440%	10.440%
730	7.9500%						

**FORWARD
RATE
MATRIX**

DAYS	RATES (ANNUAL)	31	92	184	274	365	457	549	639	730
0										
31	6.2500%	6.627%								
92	6.5000%	6.912%	7.101%							
184	6.8000%	7.209%	7.405%	7.716%						
274	7.1000%	7.507%	7.705%	8.013%	8.308%					
365	7.4000%	7.592%	7.754%	7.974%	8.102%	8.308%				
457	7.5000%	7.628%	7.763%	7.930%	8.000%	8.102%	7.898%			
549	7.5000%	7.669%	7.785%	7.925%	7.977%	7.977%	7.867%	7.799%		
639	7.6000%	8.026%	8.161%	8.340%	8.464%	8.464%	8.503%	8.708%	7.906%	
730	7.9500%								8.112%	10.440%



The following characteristics of forward rates should be noted:

1. Forward rates lie above (below) the par rates where the yield curve is positive (negative). This means that the forward rates cross from one side to the other of the par curve where the yield curve changes shape.
2. Forward rates have greater momentum than par rates; that is, the rate of change of the forward rates is more attenuated than that of the par curve.
3. Forward rates can be more volatile than par rates; that is, a small change in par rates can lead to a proportionately larger change in the forward rate.

It is important to note that forward rates, when regarded as forecasts of future short-term interest rates, require a number of theoretical and practical assumptions. From a theoretical perspective, this approach assumes the absence of transaction costs and assumes the validity of the pure expectations theory of the term structure of interest rates. In particular, the forward rate as calculated from the current cash market yield curve contains no compensation for risk and, in particular, includes no liquidity premium. In practice, the last condition is violated as forward rates are generated from the observed interest rate term structure, which, typically, incorporates a liquidity premium.

2.4 Derivation of zero rates

2.4.1 Basic methodology

The actual computation of the zero rate yield curve is complex. In theory, for each future payment of a coupon security, there exists a zero rate that discounts that payment to its present value. These rates constitute the zero rate curve, points along which represent the yield to maturity of a zero coupon bond for the appropriate maturity rate. This zero coupon yield curve is estimated from the existing par or coupon yield curve. This is completed by calculating equilibrium zero coupon rates which value each component of the cash flow of conventional coupon securities in an internally consistent fashion, such that all par bonds would have the same value as the sum of their cash flow components.

The zero coupon rates are calculated using an iterative methodology whereby the zero coupon rate is determined from a known yield curve for the successive points in time (often referred to as bootstrapping). An alternative technique for deriving the zero rates is using the implied forward rates.

2.4.2 Calculating zero coupon rates through bootstrapping

The bootstrapping approach involves a series of distinct steps:

- separate a coupon bond into a series of zero coupon bonds;
- utilise available zero rates to price components; and
- solve for the unknown zero rate within the constraint that the market value of the bond must be equal to the value of the components using zero rates.

Exhibit 3.7 shows the simple calculation of a zero coupon rate. Given that a one year bond has a coupon and yield to maturity of 8.00% pa semi-annual, a total price of \$1m is derived. However, if the six month discount security has a yield of 7.00% (not 8.00%) and the first coupon is discounted accordingly at 7.00% (which is a known zero coupon rate), then the 12 month payments must be discounted at a rate higher than the rate (8.02% pa semi-annual in this case) to maintain the equilibrium price of \$1m.

In a similar way, break even zero rates for each subsequent maturity can be derived through iteration. Known zero rates are used to derive the succeeding zero using the same logic to generate a complete yield curve of zero rates from the par rate curve.

The zero rate could, in theory, also have been derived from the discount factor. The relevant discount factor for the 1 year rate is in fact the present value of the final cash flow in 1 year divided by the final cash flow which can be used to solve for the zero rate. This is shown in *Exhibit 3.7*.

2.4.3 *Calculating zero coupon rates through implied forward rates*

As an alternative method, it is possible to determine the zero coupon rate curve by using the forward rates implicit in the current yield curve and assuming compounding of intermediate cash flows at the implicit forward rates. The basic concept is to use forward rates to reinvest intermediate cash flows to synthesise a zero coupon bond and derive the zero rate which equates the two cash flows. *Exhibit 3.8* sets out an example using the same data as in *Exhibit 3.7*. In theory, both approaches should yield identical results provided a consistent yield curve is utilised.

The zero rate using the forward rate through the pyramid technique can also be derived using discount factors. This is done by taking the six month discount factor and multiplying it by the discount factor calculated from the forward rate for 6×12 forward rate (in the above example). The product is effectively the zero rate discount factor for 1 year. This is shown in *Exhibit 3.8*.

**Exhibit 3.8
Derivation of Zero Rates—Implied Forward Technique**

CALCULATING BREAK-EVEN ZERO RATES—FORWARD RATES

YEARS	PAR RATES (%PASA)	FORWARD RATES (%PAA)	DISCOUNT FACTOR	BOND	CASH FLOWS RE-INVESTMENT	FINAL	ZERO RATE (%PA ANNUAL)	ZERO RATE (%PAS/A)	DISCOUNT FACTORS FROM ZERO RATES	FORWARD RATES
0.00	7.00%	7.12%	0.956916	(1,000,000)		(\$1,000,000)	7.12%	7.00%	0.966184	
0.50	8.00%	8.16%		40,000		\$1,081,801	8.16%	8.02%	0.924384	
1.00				1,040,000	\$1,801					0.924556

2.4.4 Characteristics of zero coupon rates

Exhibit 3.9 sets out examples of zero rates for hypothetical yield curves.





Exhibit 3.9—continued

YEAR	DAYS TO PAYMENT	PAR RATE	ZERO RATE
0.50	182	9.0000%	9.0000%
1.00	365	8.7500%	8.7445%
1.50	547	8.6500%	8.6420%
2.00	730	8.5500%	8.5375%
2.50	912	8.4750%	8.4585%
3.00	1095	8.4000%	8.3781%
3.50	1278	8.3250%	8.2964%
4.00	1461	8.2500%	8.2135%
4.50	1643	8.1250%	8.0710%
5.00	1826	8.0000%	7.9275%
5.50	2008	7.9500%	7.8735%
6.00	2191	7.9000%	7.8181%
6.50	2373	7.8500%	7.7616%
7.00	2556	7.8000%	7.7039%
7.50	2739	7.7500%	7.6452%
8.01	2922	7.7000%	7.5856%
8.50	3104	7.6500%	7.5250%
9.01	3287	7.6000%	7.4635%
9.50	3469	7.5500%	7.4012%
10.01	3652	7.5000%	7.3381%

ZERO COUPON YIELD CURVES



The following characteristics of zero coupon rate and the corresponding zero coupon yield curve should be noted:

1. Theoretical zero coupon rates are always above (below) the relevant par or coupon yield curve for a normal or positively (inverse or negatively) sloped yield curve. This reflects the fact that a coupon bond is a collection of zero coupon bonds and the yield to maturity on a coupon bond is simply the average of the zero coupon rates on the constituent zero coupon securities. Consequently, if yield is increasing in a normally sloped yield curve, then each constituent zero element of the coupon bond will have a yield which is less than or equal to that on a zero with a

maturity that is the same as the coupon bond dictating that the yield on the coupon bond must be less than a zero of the maturity. A reverse logic is applicable in the case of negatively sloped or inverse yield curves.

2. The steeper the curve the more steep is the zero coupon rate curve.
3. Zero coupon rates can be more volatile than par rates as each zero rate is dependent on each forward rate leading up to the maturity of the zero coupon rate. A movement in any of the rates results in a movement in the zero coupon rate.


3. YIELD CURVE MODELLING

3.1 Overview

All par rates and zero rates assume and require the existence of a known yield curve. However, much of the problems in derivation of interest rates and discount factors revolves around the issues in generating a complete yield curve of rates for generation of the suitable zero rates used for valuation.

Exhibit 3.10 sets out a typical yield curve which illustrates the difficulties with defining the complete yield curve under most conditions. The curve highlights the following problems:

- data noise—the yield curve may be obscured by the presence of noisy data points whereby there may be a number of yields for similar or identical maturities, which has the effect of increasing the difficulty of determining the *true* interest rate for any maturity; and
- data sparseness—the yield curve may be sparse; that is, it may have significant gaps between observable interest rates. This makes it difficult to specify interest rates for maturities *between* the observed data points.



The yield curve described is fairly typical with the problem of noise at the shorter maturity and sparseness in longer maturities. The problems described may be caused by a number of factors including the institutional structure of the market, such as the regulatory framework and liquidity factors, as well as tax factors which may affect trading and valuation of financial instruments.

It should be noted that the above assumes homogeneity in terms of default risk or credit quality of the complete set of interest rates; that is, the rates are all risk free rates or of an identical credit quality. In practice, the interest rates may not be homogenous. This problem is considered below in the context of yield curve construction in practice.

The yield curve modelling problem is separable into two separate and distinct problems:

1. interpolation—that is, the generation of a complete yield curve from the data available; and
2. term structure of interest rates—that is, an understanding of the yield curve shapes and interest rate evolution over time.

The second naturally influences the first process.

3.2 Interpolation techniques

As noted above, the process of interpolation requires the use of available interest rates at various points in the yield curve to generate a complete term structure of yields from which the relevant zero rates can be stripped.

There are a number of interpolation techniques:

1. *Stepped model*—all points on the yield curve are given by the nearest actually observed interest rates.
2. *Linear interpolation*—all points on the yield curve are created by a straight line drawn between each actually observed interest rate.
3. *Non-linear interpolation*—all points on the yield curve are fitted to actually observed interest rates using either regression or spline techniques.

In practice, linear and non-linear techniques are the most important interpolation practices used. Each of these is discussed below.

The need and importance of creating an accurate and consistent yield curve using the choice of interpolation techniques available is best achieved when the yield curve generated satisfies the following criteria:

1. *Fit*—that is, the yield curve generated is consistent with and closely tracks *observed market interest rates*.
2. *Low in noise*—that is, the curve has the appropriate degree of fit in that it is not volatile in response to noisy data (usually where the curve is over fitted).
3. *Consistent*—that is, the par, forward and zero rate derived from the curves are consistent with the observed and theoretical behaviour of these rates. In essence, they are arbitrage free.
4. *Smoothness*—that is, the par, forward and zero rates derived are smooth in that they do not show sudden and unexpected changes and volatility.

In practice, all the criteria identified are unlikely to be satisfied *simultaneously*. In addition, the appropriate trade-off between the criteria is not readily definable. This necessarily introduces a substantial degree of subjectivity in the choice of method and the generation of the yield curve.

3.3 Linear interpolation

Linear interpolation requires the use of straight lines as between any two points of the observed yield curve to estimate the interest rate between these points. *Exhibit 3.11* sets out an example of linear interpolation.

Exhibit 3.11			
Linear Interpolation			
Given the following seven and ten year interest rates, the benchmark interpolated bond rate for an eight year rate is calculated as follows:			
Maturity	Yield		
7 Years	7.30% pa		
10 Years	7.47% pa		
Interpolated yield is calculated as follows:			
Interest Rates	Maturity	Days (between)	Blending Factor
7 year rate	15/4/19X4		$556 / (389 + 556) = 0.588$
		389	
8 year rate	9/5/19X5		
		556	
10 year rate	15/11/19X6		$389 / (389 + 556) = 0.412$
Maturity	Interest Rate (%pa)	Blending Factor	Blended Rate (%pa)
7 year rate	7.30	0.588	4.292
10 year rate	7.47	0.412	3.078
		8 year Interpolated Yield	7.37

It is important to note that linear interpolation *on interest rates* is equivalent to *exponential interpolation* on discount factors. Consequently, it is usually done with interest rates rather than discount factors.

The major advantage of linear interpolation is its simplicity and ease of calculation. The disadvantages include:

- the tendency to produce inaccurate rates where the yield curve is changing slope reflecting an inherent tendency for discontinuity (kinks) at each maturity point where the yield curve is not linear in slope;
- the difficulty of generating rates where there is sparse or noisy data; and
- the prospect of generating yield curves which are inconsistent with term structure models of interest rates and also inconsistent with the concept of yield curves which change shape continuously.

3.4 Non-linear interpolation models

3.4.1 Introduction

The concept of non-linear interpolation is predicated on the use of mathematical techniques to generate a fitted yield curve through observed interest rate points. This is undertaken with the objective of fitting a yield curve which reflects the optimality criteria identified and is consistent with the term structure of interest rate assumptions usually made. As noted above, two types of models are generally utilised: regression-based models; and cubic spline-based models.

3.4.2 Regression-based models

A number of models have emerged which seek to use regression techniques, usually non-linear least square regression techniques, to create a fitted yield curve. The models are generally similar in approach, differing in:

- the form of the equation; and
- the number of terms.

Two popular models are the Bradley-Crane model (described in *Exhibit 3.12*) and the Elliot-Echols model (described in *Exhibit 3.13*).

Exhibit 3.12

Regression-based Yield Curve Models—Bradley-Crane Model

The Bradley-Crane model has the following form:

$$\ln(1 + R_M) = a + b_1(M) + b_2 \ln(M) + e$$

Where

R_M = Observed interest rate for maturity M

M = Maturity of the interest rate

The model implies that the natural logarithm (\ln) of one plus the observed yields for term to maturity of length M are regressed on two variables, the term to maturity and the natural log of the term of maturity. The last term (e) represents the unexplained yield variation. Once the estimated values of a, b_1 and b_2 are obtained, specific maturities of interest can be substituted to obtain estimated yields at these maturity points.

Source: Stephen P Bradley and Dwight B Crane, "Management of Commercial Bank Government Security Portfolios: An Optimisation Approach under Uncertainty" (1973) (Spring) *Journal of Bank Research* 18.

Exhibit 3.13**Regression-based Yield Curve Models—Elliot-Echols Model**

The Elliot-Echols model has the following form:

$$\ln(1 + R_i) = a + b_1 (1/M_i) + b_2 (M_i) + b_3 (C_i) + e_i$$

Where

R_i = Yield to maturity

M_i = Term to maturity

C_i = Coupon rate of the i th bond

The model implies that the natural logarithm (\ln) of one plus the observed yields are regressed on three variables, the inverse of maturity, the term to maturity and the coupon. The last term (e) represents the unexplained yield variation. Once the estimated values of a , b_1 , b_2 and b_3 are obtained, specific maturities and coupons can be substituted to obtain estimated yields at these maturity points.


The Elliot-Echols Model is useful where it is sought to fit yield curves directly to yield data for individual bonds rather than to an homogenised yield series. This might be desirable as a means of avoiding possible distortions created in the process of arriving at the synthetic yield series.

Source: Michael E Echols and Jan Walter Elliot, "A Quantitative Yield Curve Model for Estimating the Term Structure of Interest Rates" (1976) *Journal of Financial and Quantitative Analysis* 87.

An example using a regression-based model is set out in *Exhibit 3.14*.

Exhibit 3.14**Example of Regression-based Model**

Regression-based models require the fitting of a functional form to the yield curve. The function is chosen so that it reflects the general shape of the term structure. Parameters that specify the exact form of the function are evaluated to minimise the difference between observed market data and the values given by the function.

**Fitting a Curve**

The yield curve is specified as a function of rates $r_1, r_2, r_3 \dots r_m$ and time t .

$$df(t) = F(t, r_1, r_2, r_3, \dots, r_m)$$

Where $df(t)$ is the discount factor at time t . The form of the function is found by a least squares minimisation of the differences between the observed market rates and the values given by the function. This requires minimisation of:

$$\sum_{i=1}^n [F(t_i, r_1, r_2, \dots, r_m) - \text{Market Price}(t_i)]^2$$

This now gives a function for which discount factors can be obtained for any term.

Example: An Exponential Curve for Yield

Here we specify the curve of the form:

$$df(t) = a_1 e^{-r_1 t} + a_2 e^{-r_2 t} + \dots + a_n e^{-r_n t}$$

To obtain the term structure function, values for the coefficients a_1, a_2, \dots, a_n must be found. This is done using the least squares minimisation shown above. Once these are obtained, we have a function for the yield at any time.

To illustrate this consider a term structure function containing seven terms

$$df(t) = a_1 e^{-r_1 t} + a_2 e^{-r_2 t} + \dots + a_7 e^{-r_7 t}$$

The values of $r_1, r_2, r_3, r_4, r_5, r_6$ and r_7 are specified. Market rates for 1 through 10 years are used as observations. To define the yield curve, only the coefficients $a_1, a_2, a_3, a_4, a_5, a_6$ and a_7 need to be found. This is done using the minimisation above.

Exhibit 3.14—continued

Time (y)	Yield
1	7.52
2	7.57
3	7.7
4	7.8
5	7.88
6	7.95
7	7.995
8	8.03
9	8.05
10	8.06

Market Observed Rates




Exhibit 3.14—continued

Coefficient (i)	ai	ri
1	5.627834	0.00274
2	4.530894	0.019178
3	-1.28394	0.082192
4	-0.25905	0.246575
5	-0.69042	1
6	2.108328	2
7	4.998892	10

This gives a smooth curve through the known data points. At each known rate, the difference between the market and that given by the curve is small (less than 0.006).

The defined discount function is:

$$df(t) = 5.627834e^{-0.00274t} + 4.530894e^{-0.019178t} - 1.28394e^{-0.082192t} - 0.25905e^{-0.246575t} - 0.69042e^{-t} + 2.108328e^{-2t} + 4.998892e^{-10t}$$

Using this function, discount factors—and hence yields—can be obtained for any time t .

These regression-based models are generally useful in avoiding some of the problems of linear interpolation techniques. In practice, the models represent a compromise between too few terms (which tends to create a smooth curve) and too many terms (which tends to overfit the curve creating a noisy curve).

3.4.3 Cubic splines

A number of models have emerged which use polynomial functions to model and create a fitted yield curve. The basic technique used is that of a cubic spline.

The concept of spline techniques is based on creating a yield curve which does not oscillate to a significant degree (that is, it is not noisy) and is relatively smooth. In practice, this is created using splines which are pieces of elastic material which are constrained so as to pass through a given series of points but are allowed to assume other shapes in between the points specified. In theory, the spline will take the shape that minimises its strain energy which is consistent with the mathematical definition of smoothness specified in determining the fit of the yield curve.

Using splines there are two choices in fitting a yield curve:

1. Use a single high order polynomial—this is generally not favoured because there is an inherent tendency for the curve to take untractable shapes between data points.
2. Use a number of lower order polynomials which are then linked to create a complete yield curve—this is generally the favoured methodology because of its inherent flexibility and its inherent satisfaction of the condition that it go through all observed interest rate data points.


The latter technique is referred to as a piecewise cubic spline technique. Using this technique, the complete yield curve is generated as follows:

1. The observed yield curve in terms of observed data points is divided into a series $(n-1)$ where n is the number of observed data points) of pairs; in effect, a series of paired adjacent yield curve points and rates.
2. The yield curve between any of these observed pairs is then specified as a polynomial which is unique.
3. Each polynomial which specifies the yield curve shape between two unique points is related to the adjacent or neighbour polynomial so that the slope and/or the rate of change of slope is equal at the common data point between the two polynomials. In effect, the first and (optionally) the second derivative of the two polynomials are equated.

Exhibit 3.15 sets out the mechanics of implementing a piecewise cubic spline methodology of interpolation.

Exhibit 3.15**Piecewise Cubic Spline Technique of Interpolation**

The piecewise cubic spline is calculated for the following three interest rates (R) as at time t_1 , t_2 and t_3 . The technique is also applicable to discount factors. The yield curve segment is set out in the following diagram:



The yield curve is divided into two separate pieces, which are described by unique splines.

The section between t_1 and t_2 is described by:

$$R t_1 = a + b t_1 + c t_1^2 + d t_1^3$$

The section between t_2 and t_3 is described by:

$$R t_3 = m + n t_3 + p t_3^2 + q t_3^3$$

There are eight constants within the two polynomials which must be calculated. It is now necessary to ensure that the two polynomials go through the common data point and that their slope and the rate of change of slope is equal. This is done as follows:

The section at t_2 is described by both the above polynomials:

$$R t_2 = a + b t_2 + c t_2^2 + d t_2^3$$

$$R t_2 = m + n t_2 + p t_2^2 + q t_2^3$$

The first and second derivatives at t_2 must be equal, therefore:

$$b + 2 c t_2 + 3 d t_2^2 = n + 2 p t_2 + 3 q t_2^2$$

$$2 c + 6 d t_2 = 2 p + 6 q t_2$$

To ensure that the change in slope is equal to zero at the ends, the following equation must also be satisfied:

$$2 c + 6 d t_2 = 2 p + 6 q t_2 = 0$$

There are now six equations that can be solved using multiple regression techniques to generate the optimal piecewise cubic spline.

The above process generates a spline which passes through the data points which by definition is the smoothest interpolated function which fits the observed data.

The major advantages of piecewise cubic splines include:

- the fitted yield curve passes through observed data points and avoids the discontinuity or kinks where the yield curve changes slope;
- the curve is generally smooth; and
- the curve provides a robust estimation technique in a variety of market conditions.

The disadvantages of piecewise cubic splines include:

- the need for sufficient data points to allow a good fitted curve to be generated;
- the computation is somewhat complex; and
- the solution of cubic polynomials with multiple regression techniques exhibit the problems of multi-collinearity (that is, the function introduces uncertainty due to the linkages between each segment of the yield curve) and when solved using multiple regression techniques the definition of accuracy of the result is not unambiguous.

3.4.4 Other forms of yield curve interpolation

Two other forms of interpolation models are sometimes used: basis splines and Laguerre functions.

Basis splines are similar to the piecewise cubic spline technique described. The major benefit in using basis splines is that they avoid some of the problems with piecewise cubic splines described above. This is the result of the fact that basis splines go to zero at defined points reducing the linkage issues identified above. Typically, the third order basis splines are used, as they satisfy the required criteria of smoothness.

Basis splines are generated in a systematic manner with second order splines being generated from first order splines and third order splines being generated from second order splines. *Exhibit 3.16* sets out the mechanics of using basis splines to create a yield curve. The process follows the following logic:

- Each spline function is specified with a defined range. Outside this range it has a zero value. The points at which the spline is zero are referred to as knot points.
- One spline will end where another spline commences across the yield curve.
- Where knot points have been specified, each spline is weighted using multiple regression techniques. This is predicated on the fact that market bond prices can be expressed in terms of the sum of the discounted bond cash flows and the discount factor at each point in the yield curve coinciding with a cash flow is capable of definition in spline functions and weights. This will allow the bond price today to be expressed as a function of unknown function weights and the product of the spline function values and the bond cash flows. This allows the regression to be performed

nominating the bond price as the dependent variable and the function cash flows products as the independent variable. The regression process is then used to estimate the function weights.

Exhibit 3.16

Basis Spline Technique of Interpolation

The basis spline is calculated by specifying the following:

First, discount factors are specified as the sum of weighted spline functions:

$$DF(t) = \sum_{l=1}^e W_l S_l(t)$$

Where

DF(t) = Discount factor for time t

W_l = Function weight

S_l = Spline function l

Secondly, the bond price is specified as the sum of the discounted bond cash flows:

$$P_i = \sum_{j=1}^1 C_j \sum_{l=1}^e W_l S_l(t)$$

Where

P_i = Price of bond i

C_j = Bond j cash flow at time j

Finally, prices are expressed as the function weights and cashflow function product which allow determination of the unknown weights:

$$P_i = \sum_{e=1}^e W_l \sum_{j=1}^n C_j S_l(t)$$

Source: David Cox, "Yield Curves And How To Build Them" (1995) 4 *Capital Market Strategies* 29-33.

The basis spline, some commentators have argued, has better properties for fitting yield curves than the simpler piecewise cubic spline techniques. However, the basis spline techniques have a number of disadvantages:

- they are complex and computationally difficult;
- the shape of the fitted yield curve is sensitive to the location of the knot points. It appears necessary to ensure an even number of bonds are available between the knot points, which, in practice, is difficult to satisfy; and
- basis splines demonstrate instability and volatility and provides inaccurate estimates where there are gaps in the yield curve and sparse data.

An alternative interpolation technique is to utilise Laguerre functions. *Exhibit 3.17* sets out a description of interpolation techniques using Laguerre functions.

Exhibit 3.17**Polynomial-based Yield Curve Models—Laguerre Model**

Laguerre functions consist of a polynomial multiplied by a polynomial decay function in the following form:

$$I_t = (a_0 + a_1 * t_1 + a_2 * t_2 + \dots + a_n * t_n) * e^{-b * t}$$

Where

I_t = Interest rate for maturity t

t_n = Time to maturity

a_n, b = Constants

Where Laguerre functions are utilised for term structure modelling the decay function eventually dominates the polynomial component. This means that the long term rates stabilises, as predicted by a Laguerre function. This property provides Laguerre models with an advantage over other models where the estimates of long term rates continues to increase or decrease with time.

The advantages of Laguerre functions include:

- they provide a range of flexible shapes which are consistent with observable interest rate data; and
- they are consistent with theoretical work on yield curve shape and there is some evidence for their applicability to interest rate data.

Source: B F Hunt, "Modelling The Term Structure", paper presented at Conference on Options on Interest Rates (organised by IIR Pty Ltd), Sydney, March 1992.

3.4.5 Interest rate models

A newer approach to modelling the yield curve entails the use of term structure models. The key feature of these models is the use of assumed stochastic processes to drive the term structure of interest rates.

These models have the following characteristics:

- the models entail explicit recognition of the uncertain element in interest rate structure; that is, interest rates are probabilistic rather than deterministic;
- the models entail linking the term structure of interest rates to specified stochastic processes and nominated stochastic factors;
- the evolution of these factors over time in accordance with the assumed process determines interest rates; and
- the model generated interest rates satisfy certain no arbitrage conditions.

There are a large number of competing models.³ *Exhibit 3.18* sets out an example of these types of interest rate models.

3. See John Hull, *Futures Options and Other Derivatives* (3rd ed, Prentice Hall, Englewood Cliffs, New Jersey, 1997), Ch 17.

Exhibit 3.18**Interest Rate Models**

The Vasicek model specifies the following stochastic model for interest rates:

$$dr = \alpha (\gamma - r) dt + \sigma dz$$

Where

dr = Change in the short-term interest rate

α = Parameter (greater than 0) which describes the speed at which r revert to a long-run average value

γ = Long-run value of r

r = Short-term interest rate

dt = Short-term interval

σ = Volatility of r

dz = Random variable chosen from a normal distribution with mean 0 and variance dt

The process specified identifies that the change in the short term rate r over the interval dt will have two components:

1. A deterministic component ($\alpha (\gamma - r) dt$) whereby r will revert to a long run value at a speed parameter (α).
2. A stochastic component (σdz) which will change randomly.

The structure of the first term implies that if r is close to (away from) its long run value, the deterministic term will be small (large). This term reflects the premise of mean reversion whereby interest rates tend towards some normal rate. The stochastic term will be larger as the time over which change occurs increases. The structure is designed to be consistent with the general pattern of evolution of interest rates in capital markets.

The specified process for interest rate changes allows the derivation of valuation formula for a discount bond, which in turn facilitates the solution for the value of interest rate derivative products.

Source: O A Vasicek, "An Equilibrium Characterisation of the Term Structure" (1977) 5 *Journal of Financial Economics* 177-188.

The model described is a relatively simple single factor model incorporating mean reversion. The major variations include two factor models (such as a short term and a long term interest rate), the inclusion or exclusion of mean reversion, and the imposition of arbitrage free conditions. Models commonly utilised include the Heath-Jarrow-Morton⁴ model and the Hull-White model.⁵

The major application of these models is in pricing interest rate derivatives, in particular, options. The research into interest rate models is largely predicated on these demands. They are also related integrally to the non-linear interpolation techniques identified above. This relationship is predicated on the fact that an appropriate curve is fitted to observed market data. The fitted curve then allows the construction of an interest rate yield curve model which obtains estimates which are consistent with the market

4. See D Heath, R Jarrow and A Morton, "Contingent Claim Valuation With A Random Evolution Of Interest Rates" (1991) *Review of Futures Markets* 54-76; "Bond Pricing and the Term Structure of Interest Rates: A New Methodology" (1992) *Econometrica* 60 at 1, 77-105.
5. See Hull, op cit n 3.

data. In this sense, the interest rate models reflect an extension of the interpolation techniques to allow the provision of solutions, both analytical or numerical, for the value of interest rate derivative products.

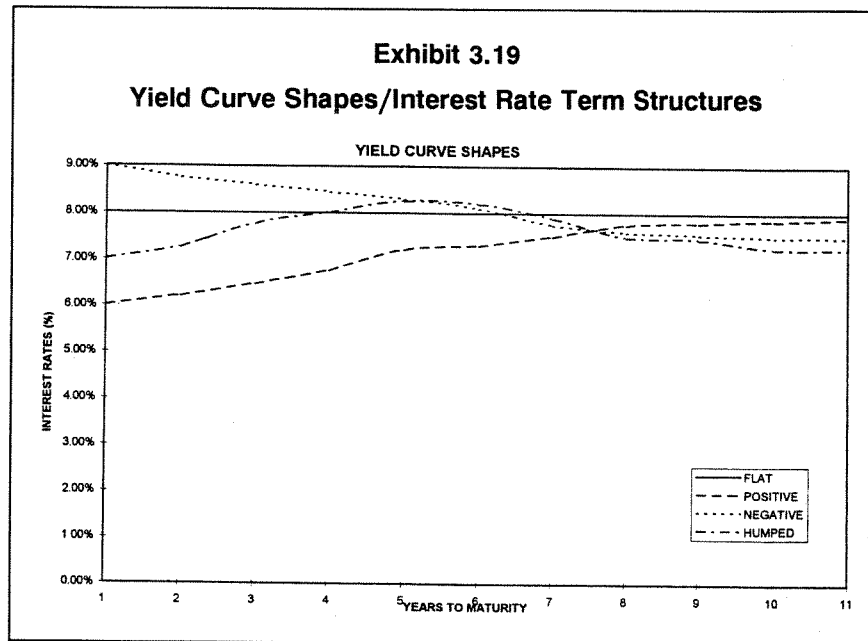
4. TERM STRUCTURE OF INTEREST RATES

Interest rates deal with the process of valuation of cash flows at different future times. Term structure deals with the pure price of time in the application of different interest rates at different future times. The process of interpolation assumes implicitly or explicitly a term structure model of interest rates. Understanding of the term structure of interest rates is essential in yield curve modelling.

The term structure of interest rates can be defined as the structure of interest rate applicable for cash flows of a homogenous credit quality for different maturities. The types of term structure (or yield curve shapes) observed (see *Exhibit 3.19*) include:⁶

- positive—interest rates increase with maturity;
- negative—interest rates decrease with maturity;
- flat—interest rates are the same across all maturities; and
- humped—interest rates increase with maturity but peak and decrease from their maximum level with further increases in maturities.

6. For a summary of interest rate term structure theories see Richard W McEnally and James V Jordan, "The Term Structure of Interest Rates" in Frank J Fabozzi and T Dessa Fabozzi (eds), *The Handbook of Fixed Income Securities* (4th ed, Irwin Professional Publishing, Chicago, 1995), pp 779-830.



The following theories of the determinants of the term structure of interest rates are usually considered:

1. *Expectations Hypothesis*—in its pure form, the expectations hypothesis states that expected interest return from securities of all maturities are equal. This implies that the return for a particular maturity represents the expected holding period return from investing in shorter term securities which are reinvested at maturity at the implied forward rates. This is based on the implicit assumption that the implied forward rate is the market consensus expected short term interest rate as of the future date. In effect, this theory is consistent with the assumption that longer term interest rates are an average of expected short term interest rates.
2. *Liquidity Preference Hypothesis*—this states that investors prefer shorter maturities in preference to longer maturities and require a premium in the form of a higher interest rate or lower price (the liquidity premium). This is consistent with the assumption of increased price risk with increases in maturity reflecting the potential instability in their capital values upon liquidation prior to maturity if required. The liquidity premium is usually assumed to increase with maturity but at a decreasing rate.
3. *Preferred Habitat/Market Segmentation Hypothesis*—this assumes that the market for shorter dated securities is segmented from that for longer dated securities reflecting the investment preferences of the underlying investors arising from their asset liability matching requirements and risk preferences. This implies that differences in the market structure as embodied in differential demand/supply equilibrium for securities of different maturities dictate the interest rate term structure.

The theories of interest rate term structure are not mutually exclusive. Newer attempts to develop models to encompass the key determinants of

term structure usually combine elements of each of the theories. For example, the biased expectations hypothesis states that interest rates reflect the combined impact of future interest rate expectations and also liquidity premia for increased maturities.

The interest rates models described above are not true term structure theories but represent new ways of modelling the yield curve. In reality, the interest rate models are consistent with the key theoretical approaches identified because each represents a special case of stochastic processes used to generate the yield curve.

Exhibit 3.20 sets out a summary of the relationship between the observable interest rate term structures and the theoretical models identified.

Exhibit 3.20				
Summary of Interest Rate Term Structure				
Type Of Term Structure	Flat	Positive	Negative	Humped
Expectation Theory	Short-term interest rates are expected to remain the same	Short-term interest rates are expected to increase	Short-term interest rates are expected to decrease	Short-term interest rates are expected to increase and then decrease
Liquidity Premium	No liquidity premia	Positive liquidity premia	Negative liquidity premia	Positive liquidity premia followed by negative liquidity premia
Preferred Habitat/Market Segmentation	Equilibrium in demand supply across all maturities	Excess of supply over demand in longer maturities	Excess of supply over demand in shorter maturities	Excess of supply over demand in intermediate maturities
Biased Expectations	Short-term rates are expected to decrease but are offset by increasing liquidity premia	Short-term rates are expected to remain the same or increase moderately and are accentuated by increasing liquidity premia	Short-term rates are expected to decrease but the rate of decrease is offset by an increasing liquidity premia	Short-term rates are expected to stay the same or increase and then decrease (the decrease being sharper than the increase), with the increase being accentuated by and the fall retarded by an increasing liquidity premia

In practice, the yield curve interpolation model utilised must be consistent with the observed interest rate term structure prevalent in the relevant market to avoid inaccuracies and poor predictive performance.

5. YIELD CURVE CONSTRUCTION IN PRACTICE

The process of derivation of appropriate interest rates and discount factors can, as shown above, be reduced into two separate and distinct processes. The first process entails the use of yield curve modelling processes to derive a complete set of interest rates. The second process entails the generation of zero rates from the yield curve which can be utilised for the valuation of financial instruments.

In practice, a number of additional considerations are relevant. A major factor underlying these uncertainties is the incomplete and imperfect nature of financial markets generally and the difficulties in nominating objective criteria to select optimal models for constructing accurate yield curves. These problems include the following:

- The difficulty in identifying yield curves which are homogenous in terms of credit or default risk. In practice, a series of interest rates derived from similar but not perfectly credit homogenous yield curves are combined.
- The problems of defining the characteristics of fit for an estimated yield curve because the criteria of fit, consistency and smoothness can be applied at different levels. For example, there are in reality multiple curves which can be fitted to satisfy the smoothness criteria: the par curve; the forward curve; the zero curve; or the discount factors generated off any of these curves. Linear interpolation often creates irregular forward curves while splines can create regular smooth curve, for one or more of these sets of interest rates or discount factors, *but not for all curves*.

The problems identified are not necessarily capable of perfect solution. In practice, practitioners use compromises reflecting trade-offs between market structure, data integrity, estimation accuracy, computational efficiency and cost effectiveness. A major criteria is the issue of hedge-ability; that is, the capacity to hedge the components of the yield curve and the interest rate risks assumed in the course of pricing, valuing and trading financial instruments off the selected yield curve. In essence, practitioners will favour the construction of a yield curve which not only is nearest the theoretical paradigm but one which also facilitates trading and hedging activities.

Against this set of constraints, most practitioners utilise two separate yield curves for valuation purposes:

1. A risk free curve—usually constructed from the available series of interest rates on government securities of the relevant tenor.
2. A risk adjusted (the swap) curve—usually constructed from a mixture of instruments (including short term inter-bank rates, near term short term interest rate futures contracts or forward rate agreement (FRA) prices/rates, and interest rate swap rates. The swap curve is used to value credit risk affected (that is, non government risk financial instruments) where appropriate, incorporating adjustment spreads where the underlying

instrument is considered to have fundamentally different risk or other characteristics.

The derivation of the risk free curve follows the established procedures identified. Key considerations include, depending on the market, noisy data (reflecting a large number of government securities of identical or similar maturities with different coupons trading at different yields) and data sparseness (whereby there may be significant gaps in the yield curve). In practice, these are overcome by constructing a fitted curve (using one or other of the techniques identified) and generating the required zero rates from that curve.

The generation of the swap curve is more complex. As noted above, this curve is constructed by combining a series of interest rates from different instruments. The following example, which uses the US market, is indicative of both the approach and the key considerations which are relevant. It should be noted that the problem of yield curve construction in other markets is necessarily more complex and less readily soluble than those in the US example, reflecting the relative maturity, liquidity and efficiency of that market.

In practice, the swap curve in US\$ is constructed as follows:

- The cash rate (usually based on an interbank rate such as LIBOR) to the first IMM eurodollar delivery date (the near month contract) is taken. Some interpolation is usually required, as the period may not coincide with the traded maturities for cash, which are usually overnight, one week, one, three, six et cetera months.
- The next series of rates taken are the eurodollar futures rates. The number of successive eurodollar futures contracts varies but in practice will be between 12 and 20 quarterly contracts (three to five years).⁷ For each contract, the traded futures price is deducted from 100 to determine the forward rate which is then incorporated into the yield curve. This process is repeated for each contract.
- Beyond the futures contracts, available interest rate swap rates are utilised to complete the yield curve.
- Certain dominance rules are specified; for example, the eurodollar rates may or may not be overridden by the relevant swap rates.

The curve once derived is fitted and zero rates generated in the established manner. *Exhibit 3.21* sets out a simple example of the process described.

7. It is likely in currencies other than US\$ the number of futures contracts used would be lower, say four to eight contracts (one to two years), reflecting the fact that the futures markets in the relevant currency do not allow trading beyond this maturity and/or the liquidity of the market. Technically, in the US\$ market, it is theoretically possible to trade eurodollar futures out to 10 years (40 successive quarters), which would mean it would be possible to derive a 10 year yield curve from the futures rates.

Exhibit 3.21**Constructing a Zero Coupon Curve**

This is an example of how to construct a zero coupon curve using bills, bill futures and swaps. We use the bootstrap method. Interpolation is based on keeping the forward rates constant. Internal calculations will be in terms of continuously compounding rates and discount factors.

The instruments

For the curve, we use instruments whose prices are available in the marketplace. These are bills, bill futures and swaps. We will assume that the bill futures take precedence where there is any overlap.¹

To keep the procedure general, all times will be in days or years. This means that dates can be omitted. In a real application, dates would be converted to days. The spot date is assumed to be day zero. All days/years are relative to the spot date unless otherwise specified.

Bills

These are pure discount instruments whose maturity is a fixed number of days out from the spot date.

# Days	Rate
1	6
30	6
60	5.98
150	6

The price of a bill is just its face value discounted over the appropriate number of days. Thus we can directly obtain discount factors for bills. If the bill has d days to maturity then the discount factor is

$$df = \frac{1}{\left(1 + r \frac{d}{365}\right)}$$

This gives a continuously compounded zero rate of

$$r(\text{zero}) = \frac{-365}{d} \ln(df)$$

- i. This is a realistic assumption as they are the most liquid of the instruments used. Also, short dated swaps are usually hedged using bill futures.

Exhibit 3.21—continued

# Days	Rate	df	Cts Zero Rate
1	6	0.999836	5.999507
30	6	0.995093	5.985254
60	5.98	0.990266	5.950799
150	6	0.975936	5.927221

Bill futures

Bill futures are just like bills, except that they start out of a forward date. In this example, we have a strip of bill futures. Where one bill future matures, the next begins. From the table below, the first bill future starts at day 40 and matures on day 132. The next starts on day 132 and expires on day 223 and so forth. The yield of a bill future is just 100 minus its price.

Days to Start	Days to Expiry	Price	Yield
40	132	94.000	6
132	223	93.920	6.08
223	314	93.660	6.34
314	405	93.350	6.65
405	496	93.110	6.89
496	587	92.940	7.06
587	678	92.790	7.21
678	769	92.640	7.36
769	860	92.520	7.48
860	951	92.420	7.58
951	1042	92.330	7.67
1042	1133	92.220	7.78

It is a simple matter to get the forward discount factor and the forward rate continuously compounded for a bill future. These are given by the same formulae as for bills. If the bill future starts at d_s and expires at d_e , then

$$df_s^e = \frac{1}{(1 + r \frac{(d_e - d_s)}{365})}$$

$$r_s^e(\text{zero}) = \frac{-365}{(d_e - d_s)} \ln(df_s^e)$$

Exhibit 3.21—continued

These are forward discount factors and rates. They need to be linked to spot rates. This will be illustrated below when we combine all the components of the curve.

Days to Start	Days to Expiry	Price	Yield	Cts Zero Rate	df
40	132	94.000	6	5.955082	0.985102
132	223	93.920	6.08	6.034379	0.985068
223	314	93.660	6.34	6.290415	0.984439
314	405	93.350	6.65	6.595475	0.983691
405	496	93.110	6.89	6.831492	0.983112
496	587	92.940	7.06	6.998586	0.982703
587	678	92.790	7.21	7.145964	0.982342
678	769	92.640	7.36	7.293288	0.981981
769	860	92.520	7.48	7.411109	0.981693
860	951	92.420	7.58	7.509266	0.981452
951	1042	92.330	7.67	7.597587	0.981236
1042	1133	92.220	7.78	7.705509	0.980972

Swaps

Swaps are priced as fixed coupon instruments whose coupon is equal to the swap rate. In other words they are par bonds. Conventionally, swaps of three years and less pay quarterly coupons while swaps of longer maturity pay coupons semi annually. It is convention to quote rates for swaps of maturity one to five years, then for seven and ten year swaps. For maturities of six, eight and nine years, we shall interpolate the rates linearly.ⁱⁱ

# Years	Rate	Freq
1	6.5	4
2	7	4
3	7.2	4
4	7.25	2
5	7.35	2
6	7.415	2
7	7.48	2
8	7.52	2
9	7.56	2
10	7.6	2

The swaps will be priced when the curve is assembled.

ii. It may be more appropriate to use higher order interpolation. Alternatively, the curve can be built using only the quoted swap rates. The other rates obtained from the curve.

Exhibit 3.21—continued

Putting it together

We have chosen to construct the curve giving precedence to bill futures where there is overlap. Thus the bill of 150 days will be omitted. The 60 day bill will not be used explicitly, but will be used to obtain a bill rate to the beginning of the bill futures strip. Swaps of three years and less are not needed either.

Up to 30 days, the bills provide zero coupon instruments directly. Once a rate at 40 days is known, the bill futures can be used to compute the curve out to 1133 days. The swaps are then required for further dates.

As we know zero rates from the bills to 30 and 60 days, a 40 day bill rate can be implied by interpolation. In this example we assume that forward rates are constant. Using the 30 and 60 day bill rates, the discount factor and forward rate from 30 to 60 days is given by:

$$df_{30}^{60} = \frac{df_0^{60}}{df_0^{30}} = \frac{0.990266}{0.995093} = 0.995149$$

$$r_{30}^{60} = \frac{365}{(60 - 30)} \ln(df_{30}^{60}) = 5.916344\%$$

This forward rate is used from 30 to 40 days (as the interpolation keeps the forward rates flat). Thus the 40 day discount factor and zero rate can be obtained.

$$df_{30}^{40} = df_0^{30} e^{-0.05916344 \cdot 10/365} = 0.993481$$

The process so far has used bill rates directly to obtain the zero curve to 30 days. The 60 day bill is used to obtain the curve to the beginning of the bill futures strip. The futures can be combined directly to generate the curve out to 1133 days.

	Days (Start)	Days (Expiry)	Zero Rate	Spot DF	Fwd DF	Fwd Rate
Bill	30	60	5.985254	0.995093	0.995149	
Bill	60		5.950799	0.990266		5.916344
Futures Splice	40	60	5.968027	0.993481	0.99838	5.916344
Futures	40	132	5.968027	0.993481	0.985102	5.955082
	132	223	5.959005	0.97868	0.985068	6.034379
	223	314	5.989763	0.964067	0.984439	6.290415
	314	405	6.076895	0.949065	0.983691	6.595475

Exhibit 3.21—continued

The forward discount factors for the bill futures have already been obtained above. These can be combined with the discount factor to the beginning of the futures strip to build the curve. For the first future

$$df_0^{132} = df_0^{60} df_{60}^{132} = 0.993481 \times 0.985102 = 0.97868$$

and the continuous zero rate is

$$r_0^{132} = \frac{-365}{132} \ln(0.97868) = 5.959005\%$$

This can be continued for all the bill futures.

# Days	df	Cts Zero Rate
1	0.999836	5.999507
30	0.995093	5.985254
60	0.990266	5.950799
150	0.975936	5.927221
132	0.985102	5.955082
223	0.985068	6.034379
314	0.984439	6.290415
405	0.983691	6.595475
496	0.983112	6.831492
587	0.982703	6.998586
678	0.982342	7.145964
769	0.981981	7.293288
860	0.981693	7.411109
951	0.981452	7.509266
1042	0.981236	7.597587
1133	0.980972	7.705509




Exhibit 3.21—continued

The Zero Curve to the End of the Bill Futures Strip

The swaps are now required. The first swap to be used is the four year swap. The swap rate of 7.25% implies the following cashflows

Years	Cashflow
0	-100.000
0.5	3.625
1	3.625
1.5	3.625
2	3.625
2.5	3.625
3	3.625
3.5	3.625
4	103.625

As the curve has already been generated out to 1133 days (or 3/104 years), this can be used to value the cashflows to three years. The remaining two cashflows are required to generate the curve.

Years	Days	Cashflow	DF	NPV
0	0	-100.000	1	-100
0.5	182	3.625	0.970623	3.51851
1	366	3.625	0.940189	3.408186
1.5	547	3.625	0.908889	3.294723
2	731	3.625	0.876684	3.177981
2.5	912	3.625	0.845036	3.063255
3	1096	3.625	0.81323	2.947957
3.5	1278	3.625		
4	1462	103.625		

The NPV of the swap must equal exactly zero if the cashflow on the spot date is included. Of the cashflows where the curve is available, the NPV is -80.5894. The complete equation is

$$df_{1278} \times 3.625 + df_{1462} \times 103.625 = 80.5894$$

This cannot be solved directly as there are two unknown discount factors. If we use our interpolation rule keeping the forward rate constant, then an iterative solution can be generated.ⁱⁱⁱ

This gives a solution for the forward rate of $r_{1096}^{1462} = 8.022778$ and discount factors $df_{1278} = 0.781339328$ and $df_{1462} = 0.750369643$. This solves the equation for the swap price above. This procedure is then repeated for all the other swaps.


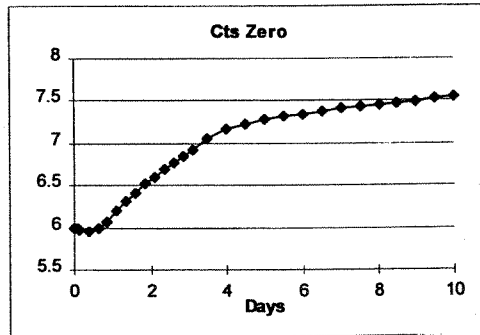
iii. Any common numerical method such as Newton-Raphson can be used.

Exhibit 3.21—continued

Instrument	Years ^{iv}	Discount Factor	Continuous Forward Rate	Continuous Zero Rate
Spot		1		
1 Day Bill	0.003	0.9998	5.9995	5.9995
30 Day Bill	0.082	0.9951	5.9848	5.9853
Futures	0.110	0.9935	5.9163	5.9680
	0.362	0.9787	5.9551	5.9590
	0.611	0.9641	6.0344	5.9898
	0.860	0.9491	6.2904	6.0769
	1.110	0.9336	6.5955	6.1934
	1.359	0.9178	6.8315	6.3105
	1.608	0.9019	6.9986	6.4172
	1.858	0.8860	7.1460	6.5150
	2.107	0.8701	7.2933	6.6071
	2.356	0.8541	7.4111	6.6922
	2.605	0.8383	7.5093	6.7703
	2.855	0.8226	7.5976	6.8426
	3.104	0.8069	7.7055	6.9119
4y Swap	3.501	0.7813	8.1037	7.0471
	4.005	0.7504	8.0228	7.1699
	4.501	0.7223	7.6876	7.2269
5y Swap	5.005	0.6948	7.6876	7.2733
	5.501	0.6689	7.6782	7.3098
6y Swap	6.005	0.6435	7.6782	7.3407
	6.501	0.6190	7.8404	7.3789
7y Swap	7.005	0.5950	7.8404	7.4121
	7.504	0.5725	7.7330	7.4334
8y Swap	8.008	0.5506	7.7330	7.4523
	8.504	0.5295	7.8649	7.4763
9y Swap	9.008	0.5089	7.8649	7.4981
	9.504	0.4892	7.9816	7.5233
10y Swap	10.008	0.4699	7.9816	7.5464

iv. Note that the number of years is not exactly integral for the swaps. This is because when this example was generated, it was assumed that there were 365 days per year. Exact calendar dates were used.

Exhibit 3.21—continued




The graph showing the forward rates illustrates the interpolation method. The forward rates are stepped. They change discontinuously. Where there are frequent and liquid instruments (the bills and bill futures), the change is regular. Once the instruments become less liquid, there is some irregularity in these rates. The zero coupon rates are much smoother. It is rare that a long dated forward instrument will require pricing. If it does, the spread will have to be large to account for the irregularity in the forward rates. Alternatively an interpolation method where the forwards are kept smooth will need to be employed.

Obtaining discount factors

Once the curve has been generated, we have discount factors (and rates) at the points where cashflows from the underlying instruments occur. The curve will price these instruments. This in itself is not useful as the prices are already used in the curve construction. Indeed this is a circular situation. The value of the zero curve is in having a model which gives discount factors at any time in the future. To obtain these, we can either graphically retrieve them, or use the interpolation method implicit in the curve.

Exhibit 3.21—continued

**The Discount Factors**

Using the interpolation method to obtain discount factors requires minor calculation. As an example, we will obtain the discount factor at 7.4 years.

Instrument	Years ^v	Discount Factor	Continuous Forward Rate	Continuous Zero Rate
7y Swap	7.005	0.5950	7.8404	7.4121
	7.504	0.5725	7.7330	7.4334
8y Swap	8.008	0.5506	7.330	7.4523

We know the discount factor to 7.005 years. We also know that the forward rate for any period between 7.005 and 8.008 years is 7.7330% (as our interpolation method keeps the forward constant). We can obtain the discount factor from 7.005 to 7.4 years.

$$df_{7.005}^{7.4} = e^{-0.07733t} = 0.969952$$

This then allows the discount factor and zero rate to be calculated.^{vi}

Instrument	Years ^{vii}	Discount Factor	Continuous Forward Rate	Continuous Zero Rate
7y Swap	7.005	0.5950	7.8404	7.4121
7.4y	7.4	0.577089	7.7330	7.429182
	7.504	0.5725	7.7330	7.4334
8y Swap	8.008	0.5506	7.7330	7.4523

Discount factors can be obtained to any dates by this process.

- v. Note that the number of years is not exactly integral for the swaps. This is because when this example was generated, it was assumed that there were 365 days per year. Exact calendar dates were used.
- vi. In this equation, $t = 0.394520$.
- vii. Note that the number of years is not exactly integral for the swaps. This is because when this example was generated, it was assumed that there were 365 days per year. Exact calendar dates were used.

Exhibit 3.21—continued**Pricing cashflows**

Arbitrary cashflows can now be valued by obtaining the discount factors to the dates where the cashflows occur. To price any instrument, decompose it into cashflows, then use this method. If there are liquidity or other conditions, the zero curve can be modified to account for these. Also the difference between market prices and the price on the curve can be calculated. This is useful to assess the premium or discount the market is building in to non standard instruments. As an example, corporate bonds can be valued on a zero coupon curve built using government bonds. The discount for credit liquidity and other factors can be quantified.

The approach described requires careful consideration of the following factors:

- The nature of the instruments is different, including differences in credit risk and instrument features. For example, the use of futures contracts introduces the following factors: the payment of deposits and margins; the differential credit risk of the clearing house; and, in the case of eurodollar futures, the problems of the fixed tick point value (0.01% is equal to US\$25) or the negative convexity.
- The problem of overlapping dates. This can be illustrated by comparing and contrasting the use of FRAs versus futures. The FRAs out of spot will trade at regular runs (3 × 6; 6 × 9 et cetera) which allows the calculation of the relevant forward interest rates and discount factors which are incorporated in the yield curve. In contrast, the eurodollar futures contracts are standardised. They are traded to pre-specified dates and are cash settled against three month LIBOR. As these dates may not be precisely three months apart, the prospect exists for gaps or overlaps between the end of the reference period and the maturity date of the next eurodollar contract. The practical import of this is that the discount factors cannot be multiplied or forward rates compounded as the interest periods are not exactly linked. This requires adjustments to the futures rates.
- The forward rates embodied in eurodollar futures prices reflect some of the inherent biases in futures prices, including the impact of margin payments and the negative convexity. Where interest rates are expected to increase, the holder of a short (long) futures position will receive (be required to pay) margin payments which can be invested (must be funded) at higher interest rates. This uncertainty forces the futures rate to trade above the theoretical arbitrage free forward rate. In practice, the bias in futures prices and the negative convexity necessitate additional adjustments based on the expected volatility of interest rates.
- The selection of the transition point between the futures curves and the swap curve is relatively arbitrary reflecting institutional and market considerations as well the margining and convexity issues identified.

The presence of these factors means that the generation of the relevant yield curve is unlikely to be an objectively verifiable activity.

The major practical problem to arise relates to the fact that more than one equally tractable yield curve can be created from the same set of market data reflecting differences in interest rate selection and mode of adjustment for the problems identified. In reality, the problems are generally likely to be confined to the shorter end, particularly the transition points from one set of interest rates to the next source. A major area is the transition from the futures or FRA rates to the swap rates.

Given that the zero rates are likely to be different, either a selection must be made between the yield curves or adjustment made to the set of rates to be used. One possible basis for selection between the curves is by application. Where an adjustment is to be made it is necessary to either segment the curve to ensure the absence of overlap or create blended curves. There are problems with each approach:

- the use of different curves differentiated by application will create different values and prices for different transactions opening up the possibility of value loss; and
- the adjustments will either create severe discontinuity in the curve and irregularities in the rates or require complex and highly subjective adjustments.

In practice, there is little uniformity in approach to these issues. Each institution generally employs its own set of techniques to deal with the problem, reflecting the nature of the market and the purpose for which the generated zero rates are to be utilised.

5. SUMMARY

The pricing, valuation and trading of financial instruments, irrespective of whether it is a simple fixed income security or a derivative transaction, requires and assumes the availability of interest rate or discount factors extending across the maturity spectrum. In practice, the practitioner is required to choose between different interest rates to value these transactions. The best practice method now applied universally in capital markets is to discount the transaction cash flows using zero rates derived from the relevant yield curve. The zero rate itself requires the construction of the yield curve using sophisticated mathematical techniques within a framework of economic theory which is consistent with observed interest rate behaviour. While the requirement for accurate zero rates is now recognised, deficiencies in market structure and data availability present significant challenges to the derivation of accurate, consistent and computationally efficient yield curves.

Part 3

Derivative Pricing

Chapter 4

Pricing Forwards and Futures Contracts¹

by John Martin

1. INTRODUCTION



A forward is any contract which obliges you to buy or sell a financial instrument or physical commodity at some date in the future at an agreed price. For our purposes, forwards includes OTC forward contracts and exchange-traded futures contracts, and includes instruments such as:

- foreign exchange forward contracts;
- forward rate agreements;
- forward bonds;
- short-term interest rate futures;
- bond futures;
- stock index futures; and
- commodity futures contracts.

Forward contracts represent an extremely useful starting point for all derivative valuation. As we will see in other parts of this book, more complex derivatives such as swaps and options can be decomposed into portfolios of forward contracts. As a consequence, the valuation of these more complex instruments will be based partly on the forward valuation techniques developed in the following chapters.

While forward contracts are the simplest form of derivative, they represent the largest derivative type by outstanding face value and volume. *Exhibit 4.1* highlights the relative importance of forwards versus swaps and options.

1. Parts of this chapter are based on the discussion on forward contracts in J S Martin, *Derivative Maths* (IFR, 1996).



As *Exhibit 4.2* demonstrates, foreign exchange and interest rate products dominate forward transactions. It is interesting to note that most of this volume is comprised of OTC foreign exchange forwards and exchange traded interest rate futures. The volume of equity forwards is relatively low and is mainly stock index futures. Like all derivatives markets, forward markets have seen very rapid growth since the mid 1980s.

2. CHARACTERISTICS OF FORWARDS

2.1 Forwards versus cash transactions

The definition of when a transaction is a cash transaction and when it is a forward is important. We might expect any transaction which settles today to be a cash transaction and a forward is anything settling from tomorrow onward. Unfortunately this is not always the case and, depending on the underlying financial asset, a “cash” transaction can range from today for a money market transaction to several weeks, or longer, in some securities markets. Some examples of cash instrument settlement days are listed below:

Cash Market	Settlement Date
Money market transactions	Today
Euromarkets	Today + 2 business days
Foreign exchange	Today + 2 business days
Stock exchanges	Today + 5 business days
Crude oil	Today + 1 month
Some property markets	Today + 6 weeks

All of these represent “cash” transactions, however, in each market the convention applying to the settlement of these transactions changes according to market convention. The market convention usually reflects the ease with which settlement can be arranged and/or the relative complexity of changing ownership of a financial asset. For example, money market instruments in most currencies are lodged on electronic networks, where both payment and transfer of ownership can occur in a matter of minutes—same day settlement is easily possible. However, where change of ownership involves different time zones and bank accounts such as in foreign exchange, or the shipment of commodities in the oil market, or the completion of legal documentation as in property, the time to settlement becomes longer.

When considering a forward, the market convention on the time to settlement underlying a cash or spot price has to be known. A forward transaction does not commence till the settlement day passes the cash settlement date. So, in the foreign exchange market, a forward is a transaction which settles after two business days.

2.2 Forward price and value

A forward contract gives us the right to buy or sell a financial asset at some date after the normal cash settlement date at an agreed price. The attraction of forward contracts is that they provide a method of obtaining a fixed price on an asset regardless of movements in the cash price between the trade date and the settlement date.

As in the cash market, a forward *price* is agreed between the buyer and seller which reflects the relative cost or benefits to both parties of delaying

settlement of that transaction. Once this price is agreed then the market replacement cost, or the *value* of reversing, will also change as market conditions change. An important feature of forwards, and derivative valuation in general, is that there are always two components to valuation:

1. What is the forward price?
2. What is the value of an open transaction based on this price?

The forward price and cash price are usually different. When valuing a forward transaction we first need to determine the prevailing market price for a forward, and then determine the present value based on that forward price. This is simplified in futures markets, as a transparent forward price is the subject of trading and, in liquid markets, is a fair reflection of the price at which open contracts can be reversed.

Determining the forward price is not only a requirement of valuation. It is also an essential tool in comparing different forward instruments such as forwards and futures, whether you are a market maker, arbitrageur or hedger. For example, suppose you know that the true "fair" forward price of a security is \$100 and the futures market is trading at \$99. This suggests that an arbitrage opportunity of \$1 exists if you buy the futures contract and enter into another contract to sell it forward. Similarly, if you are a market maker in forward instruments you need to be able to calculate a constantly updated forward price to quote to your clients.

In this chapter we review the following:

- general forward pricing and valuation (Section 3);
- interest rate forwards and futures (Section 4);
- foreign exchange forwards (Section 5); and
- equity forwards (Section 6).

2.3 Terminology

In these chapters on forwards and futures we will use the following standardised terminology:

Term	Description
Valuation Date	The date on which a valuation is being performed (usually today)
Forward Expiry Date	The date on which the forward contract expires and the obligation to buy or sell forward falls due
Forward Settlement Date	The date on which the forward obligations arising from the forward contract must be settled in cash (usually the forward expiry date or soon after)
Forward Period	The number of days between the valuation date and the forward expiry date
Cash or Spot Price	The price paid on the valuation date for a "cash" purchase of the underlying asset
Forward Price	The price agreed to be paid for the underlying asset on the forward expiry date
Asset Income	The income paid to the owner of a financial asset—usually during the forward period. Examples of asset income include coupons and dividends.

3. GENERAL FORWARD PRICING AND VALUATION


3.1 Deriving the forward price

The easiest way to understand forward pricing is to break it down into its underlying components. Like most derivatives, forward transactions can be reproduced by a series of physical positions. This can probably be best explained by some simple examples.

3.1.1 Example: buying a painting

Calculating the forward pricing is the same question as how much should I pay to buy something in the futures. To remove the complications sometimes presented by financial instruments it is often easier to understand derivative pricing using a tangible good such as a piece of artwork. In both of these examples we ignore any compounding effects and assume that there are no other benefits from the asset apart from any described.

Suppose an art dealer offers to sell you a painting today for US\$1 million today or US\$1.2 million in one year's time. Would you buy it today or in one year? (As a hint the current one year interest rate is 10% pa.)



In *Exhibit 4.3* the problem is represented diagrammatically. The art dealer is offering to buy the painting at a cash price of \$1 million or a forward price of \$1.2 million. We wish to take the deal which is financially beneficial.

A useful way of looking at the transaction is in terms of the cash and asset balances. If we buy the painting today we give up \$1 million. This is either financed by borrowing the funds or reducing our existing cash balances, which at the end of the year incurs an interest cost or reduces interest income by \$0.1 million. While the forward purchase avoids spending cash today, it requires paying \$1.2 million in one year. At the end of one year both deals give us the same asset, however the cash cost of the forward purchase is \$0.1 million higher than the spot purchase—so we would buy the painting today.

While a simple example, it displays an important idea: a forward transaction can be replicated by purchasing the asset today and borrowing the money to finance it. The “fair” forward price is then the cash price plus the

interest cost over the life of the forward transaction. In this example, the fair one year forward price is \$1.1 million.



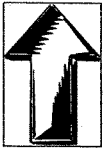



This example shows the forward price on an asset that provides no cash return in the form of coupons or dividends. Most financial assets provide an income so we need to incorporate that into our pricing framework. In the following example we use another simple tangible asset.

3.1.2 Example: buying a warehouse

You are given the opportunity to buy a warehouse as an investment for DEM 1 million today or DEM 1.1 million in one year's time. Would you buy it today or in one year? (As a hint, the warehouse is currently earning rental income of 2% pa and you can borrow DEM for 1 year at 14% pa.)

This time we need to take account of the cash income. In this case, if we buy the property today we receive DEM 0.02 million in income, which we do not receive if we purchase the warehouse forward. The rent has the effect of reducing the net cost of "carrying" that property.

Exhibit 4.4
Buying a Warehouse Forward

Today	or	1 year
		
		
		
\$1 million		\$1.1 million

Spot Purchase	Balance	Net Cashflows Balance	
Assets	Warehouse		Warehouse
Cash	-\$1m	-\$0.12m	-\$1.1m
Forward Purchase			
Assets	0		Warehouse
Cash	0		-\$1.1m
Net Cash difference of today versus forward			-\$0.02m

Using the same cash balance approach as in the above example we can see from *Exhibit 4.4* that, while the rental income reduces the net cost of buying the property the net cash cost, and the “fair” forward price, at the end of the year is DEM 1.12 million. So, in this example we would buy the property forward.

3.2 Some conclusions regarding the forward price

While simple, these two examples display the fundamentals of forward pricing. Essentially the “fair value” forward price makes buyers and sellers indifferent between buying and selling the underlying asset today or in the future based on the current market cash price, cost of financing the asset and the expected return on the asset. In other words the forward price is essentially a summary of the net financial obligations of owning an asset.

The examples also highlight four of the key features of the forward prices:

1. *Replication*: A forward purchase of an asset can be replicated by buying the asset in the cash market, financing the purchase by borrowing the cash required and then receiving any asset income.
2. *Fair price*: The “fair” forward price is given by the cash price plus the net cost of financing the asset over the term of the forward contract.
3. *Interest effect*: The interest cost tends to *increase* the forward price versus the cash price.
4. *“Dividend” or “coupon” effect*: Any cash return on the asset over the term of the forward contract tends to *decrease* the forward price versus the cash price.

These four general rules should apply to all forward prices on financial assets regardless of whether it is an interest rate, foreign exchange or equity product provided they operate in freely operating markets. It is worth noting that these relationships start to break down when you move away from financial assets, particularly to consumable commodities such as agricultural and energy products. This is because the decision to have the physical commodity today or in the future is not just a financial or investment decision, the decision to buy a commodity today or in the future also has to take into consideration when the commodity is required for consumption.

While these examples were from the point of view of a forward purchase, the same logic applies to a forward sale except it works in reverse. Looking at the painting example, if we agree to sell the painting forward we forgo receiving the cash today and any interest earned over the year. Correspondingly, if we sell the painting forward we want to ensure that we will receive a cash amount, which is, at least, equivalent to the cash price plus interest—giving us the same price as suggested by the “buy” example.

The forward price can be “synthetically replicated” using physical transactions. A forward purchase can be replicated by buying today and financing the purchase over the forward term by borrowing. While economically, the price achieved by the replication is equivalent to a forward which usually makes the derivative more attractive. First, the transaction costs of the physical replication are often greater due to the number of extra “legs” and also physical transactions often have higher execution costs than derivatives. Secondly, forward transactions are a future commitment and are “off-balance” sheet. The transactions in the physical replication are included in the balance sheet which has the effect of increasing assets and liabilities and possibly increasing capital costs.

3.3 Cashflow timing and the cost of carry

A cash and forward purchase provide ownership of the asset at different times. However, providing the only benefit offered by these assets is their income stream, and the repayment of principal in debt instruments, then this benefit of ownership now or in the future is only notional.² As we have seen the forward price incorporates both the net interest cost of holding the asset

2. This is an important condition of forward pricing. If there are other tangible or intangible benefits of owning an asset such as the appreciation of artwork, or the ability to use it for consumption, then, unless these factors can be quantified, the forward price is unknown.

as well as the asset return. As a result the difference between cash and forward transactions is only a difference of cashflow profiles:

1. *Cash purchase cashflow profile*: Requires cash today, however, it will provide income between today and the forward date.
2. *Cash sale cashflow profile*: Receives cash today but will miss out on any income between today and the forward date.
3. *Forward purchase cashflow profile*: Does not require cash till the forward settlement date and as such misses out on any asset income.
4. *Forward sale cashflow profile*: Forgoes receiving cash till the forward settlement date, however it will provide income between today and the forward date.

These cashflow profiles are interesting and explain an important part of the use of forward transactions. Any of the four alternatives represent the exchange of cash and an asset. However, if opposite spot and forward transactions are combined (for example, a spot sale and a forward purchase) then two offsetting cashflows at different points in time are created. In effect, we are creating transactions, using an underlying financial asset, which have the cashflow profile of a borrowing or lending. This is a fundamental driving force in forward markets globally and is a key reason for transactions such as "repos" (see section 4) and FX swaps (see section 5).

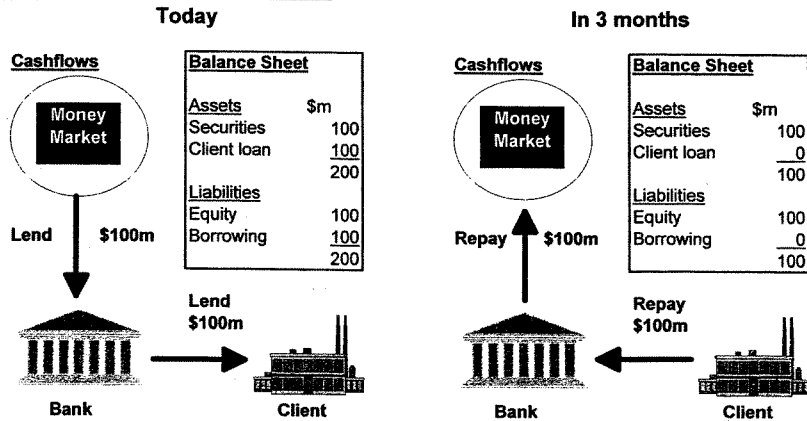
Suppose you are a small merchant bank which owns \$100 million of liquid government securities. Your organisation has a requirement for \$100 million in short-term funding needs over the next three months. The traditional way to finance this would be to borrow in the short-term money market. This creates three problems: your ability to borrow is dependant on your credit standing, as is the cost of borrowing, and the transaction will increase gearing. If, however, you entered into an agreement to sell the securities today and then buy them back in three months, you have created the underlying borrowing required. Further, the ability to raise the funds, and the cost of those funds, is primarily determined by the government securities, not your own credit rating. This type of transaction is often described as "security lending", "sell and buy", "repurchase or reciprocal purchase agreement (repo)", or as a "liquidity swap"—we will refer to it as a "repurchase agreement". Both the money market and "sell and buy" transactions are summarised in *Exhibit 4.5*. The cashflows are identical, however, the balance sheet and gearing consequences are very different—with the money market doubling total assets and liabilities while the other security transaction has no effect on the totals, just the composition.

Exhibit 4.5

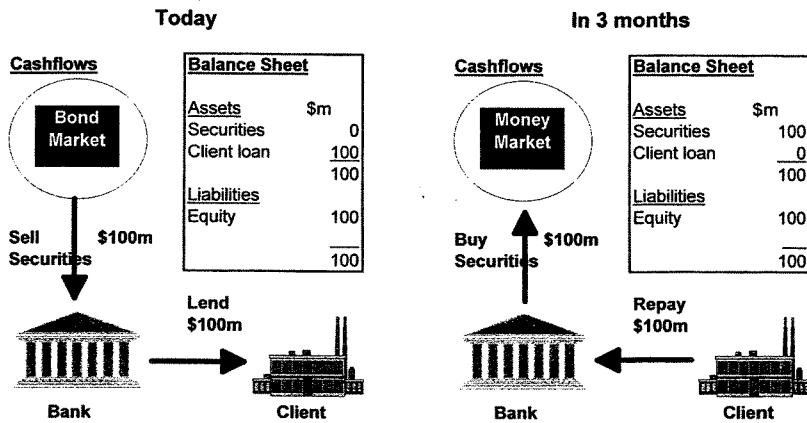
Using Forwards to Raise Finance

Your bank's only asset is \$100m worth of securities and all its liabilities are shareholders funds. It requires \$100m in funding to make a loan to a client. You can borrow the funds or enter into an agreement to simultaneously sell and buy the securities.

Option 1: Borrow in the money market



Option 2: Sell securities today and buy forward



Whereas the cost of borrowing in the money market is given by its interest rate, in the security transaction it will be given by the difference between the spot and forward price. As we know the forward price of the securities is determined by the current cash price plus the net cost of financing those positions. Given the forward purchase can offer the government securities as collateral the implied interest cost should be lower than the direct money market borrowing.

From the point of view of large holders of financial assets repurchase agreements are a balance sheet efficient and, potentially, low-cost form of financing. For example, an investment bank which is a market maker in securities and holds large bond portfolios will use repurchase agreements to finance its bond holdings in a similar way to the example in *Exhibit 4.5*. We will discuss the intricacies of these agreements in the interest rate and foreign exchange forward chapters.

3.3.1 *Cost of carry*

The key pricing concept in forward transactions is the net financing cost of creating a synthetic replica using cash instruments. The usual terminology for this net financing cost is the “cost of carry”.

$$\text{Forward Price} = \text{Cash Price} + \text{Cost of Carry}$$


$$\text{Cost of Carry} = \text{Interest Cost} - \text{Asset Income}$$

3.4 Forward pricing formulae

We now know the general relationship between cash and forward prices. In this section we will develop three simple formulae for pricing forward transactions depending on the nature of the income stream generated by the underlying financial asset during the period of time to the forward expiry date. The three forms of asset considered are:

- an asset which pays *no* income;
- an asset which pays *constant* income; and
- an asset which pays “*lumpy*” income.

The difference in these three income streams is summarised in *Exhibit 4.6*. Distinguishing by asset income allows us to develop three models which can price most financial assets. As a result, we find that we can apply the same “*lumpy*” pricing models to bonds and shares—even though, apart from the income streams, the underlying instruments have very different characteristics.



3.4.1 Forward pricing on an asset which pays no income

If the asset pays no income between the day we calculate the forward price and the expiry of that forward contract, then the forward price is the cash price adjusted for the interest cost only. Examples of these types of financial assets includes precious metals, some commodities, and artwork.

It is worth noting that this formula also applies to financial assets which pay an uneven or “lumpy” income (see section 3.4.3 below), but will not pay any income between the pricing date and the forward expiry date.

Forward Price: Asset Pays No Income

Using simple interest the calculation is as follows

$$F = S \times (1 + r \times f / D)$$

Where

- F = Forward price
- S = Cash or spot price of the underlying instrument
- r = Interest rate to forward date (preferably zero-coupon rate)
- D = Day count basis (365 or 360)
- f = Number of days to the forward expiry date

Examples of underlying assets: precious metals; artwork; and some commodities.

An important assumption in all of the forward formulae is that the interest rate, r , is a simple interest rate over the term to expiry of the forward. Formally this means the interest rate in the model is an effective zero-coupon rate. While zero-coupon rates will be discussed in detail in later chapters the important point about zero-coupon rates is that they pay no interest between today and the maturity date, and there is no risk associated with the re-investment of coupons.

In practice, most money market instruments are zero-coupon. So pricing forward transactions with up to six months to expiry is accurate. For longer terms, pricing using a coupon paying interest rate should be viewed as an estimate—accurate pricing requires zero-coupon yields.

3.4.2 *Forward pricing on an asset which pays constant income*

The assumption in this formula is that the underlying financial asset pays income at an even, constant rate over the life of the forward contract. In practice that means the asset will pay income for every day that it is held. Examples of this type of instrument include discount money market instruments, foreign exchange and broad-based equity indexes.

Forward Price : Asset Pays Constant Income

Using simple interest, the calculation is as follows

$$F = S \times (1 + (r - q) \times f / D)$$

Where

- F = Forward price
- S = Cash or spot price of the underlying instrument
- r = Interest rate to the forward expiry date
- D = Day count basis (365 or 360)
- f = Number of days to the forward expiry date
- q = Asset income expressed as a % per annum

Examples of underlying assets: money market instruments; foreign exchange; and broad-based equity indexes.

3.4.3 *Forward pricing on an asset which pays “lumpy” income*

In this formula it is assumed that the underlying financial asset pays income only at certain points over its life. Typically, the asset income is accrued over a period and then paid at the end of the period—this gives a “lumpy” appearance to the income cashflows. From the point of view of forward pricing the important consideration is how many income payments will occur during the life forward term.

Some examples of this type of instrument include bonds and shares.

Forward Price: Asset Pays "Lumpy" Income

Using simple interest and one income payment, the calculation is as follows

$$F = S \times (1 + r_1 \times f_1 / D) - c \times (1 + r_2 \times f_2 / D)$$

Where

- F = Forward price
- S = Cash or spot price of the underlying instrument
- r_1 = Interest rate to the forward expiry date
- r_2 = Interest rate between the income payment and forward expiry dates
- D = Day count basis (365 or 360)
- f_1 = Number of days to the forward expiry date
- f_2 = Number of days between the income payment and forward expiry dates
- c = Asset income expressed in the same units as the cash price

Examples of underlying assets: bonds and shares.

Sample calculations are provided in *Exhibit 4.7*, where the forward price is calculated on a security under each of the three asset income assumptions.

Exhibit 4.7**Forward Price Example**

You intend to buy a security 180 days forward. The current spot price is \$90 and the six month interest rate is 6.7% pa (A/360). Calculate the forward price under the following three asset income scenarios:

- no income;
- income paid at rate of 8% pa on a constant basis; and
- a lump sum of \$4.50 will be paid in 91 days—assume the three month interest rate in three months is also 6.7% pa.

i) No income

$$S = \$90 \quad r = .067 \quad f = 180 \quad D = 360$$

$$\begin{aligned} F &= S \times (1 + r \times f / D) \\ &= 90 \times (1 + .067 \times 180 / 360) \\ &= 93.015 \end{aligned}$$

ii) Income = 8% pa constant

$$S = \$90 \quad r = .067 \quad f = 180 \quad D = 360 \quad q = .08$$

$$\begin{aligned} F &= S \times (1 + (r - q) \times f / D) \\ &= 90 \times (1 + (.067 - .08) \times 180 / 360) \\ &= 89.415 \end{aligned}$$

iii) Income = lump payment of \$4.50

$$S = \$90 \quad r_1 = .067 \quad r_2 = .067 \quad f_1 = 180 \quad f_2 = 89 \quad D = 360 \quad c = 4.50$$

$$\begin{aligned} F &= S \times (1 + r_1 \times f_1 / D) - c \times (1 + r_2 \times f_2 / D) \\ &= 90 \times (1 + .067 \times 180 / 360) - 4.5 \times (1 + .067 \times 89 / 360) \\ &= 88.44046 \end{aligned}$$

As the results demonstrate, the nature of the income payment has a considerable impact on the forward price. And as the cost of carry concept tells us, where the asset income exceeds the cost of financing the security (scenarios i and ii) the forward price is lower than the cash price.

3.5 Valuation of forward contracts

In the previous sections we have examined how to determine a forward price based on cash market information. Now that we can generate a forward price we can determine the present value of open forward contracts. We divide this calculation into two steps: determining the value on the forward expiry date; and then determining the present value.

3.5.1 Valuation on the forward expiry date—the forward value

In the financial mathematics of financial assets, value is given by the present value of all future coupon and principal cashflows. As a result, the present value generally represents a premium or discount to the face value of the asset. This is not the case with forward contracts—when a forward contract is initially executed its value is zero, as the forward price this value can change to be positive or negative.

A forward contract represents a commitment to purchase or sell a financial asset, they are not financial assets in their own right. Unlike a financial asset which has value arising from future cashflows, the value of a forward contract only arises from the benefit or loss arising from the obligation to buy or sell the underlying asset.

If we think of the cashflows of a forward contract on an interest bearing security, then as well as the future coupon and principal repayments there is the initial cashflow associated with the purchase or sale of the security on the forward settlement date. So, if a bond is purchased forward, the cashflows consist of a cash payment on the forward settlement date and then cash receipts in the form of coupons and principal repayments over the remaining life of the bond.

At the time a forward contract is executed the forward price and the value of all of the future cashflows created by the bond after the forward expiry date are equal. However, as interest rates change the relative values of the forward price and all of the future cashflows are different. This relationship for a forward purchase of a bond contract can be summarised as follows:

$$\text{Forward Value} = \text{Forward Bond Value} - \text{Forward Contract Price}$$

where forward bond value is the value of all of the cashflows created by the bond after the forward expiry date and the forward contract price is the price agreed under the forward contract. As the forward price is fixed, the contract value will change as the forward bond value changes: if the yield to maturity on the bond falls (increasing the forward bond value) then the forward contract value will rise above zero and if bond yields rise the forward contract value will fall below zero.

Over the life of the contract the forward contract price is fixed. The forward bond value is simply calculated by determining the current forward price using the appropriate formula from Section 3.4 above. So, if the forward contract price was \$110 and the current forward price is \$120, the value of a forward purchase on the forward expiry date would be positive \$10.

This relationship holds for all forward contracts. The value as at the forward expiry date is the difference between the forward contract price and the current forward price. This relationship also explains the risk profile, or potential for profit or loss, of forward contract. This relationship is described as the “pay-off” of the forward contract and the graphical representation as a “pay-off diagram”. We will utilise this concept regularly in our investigation of derivative value.

Exhibit 4.8 provides an example of a full forward pricing and valuation exercise. It also illustrates the sensitivity of the contract value to changes in the current forward price using a pay-off graph.

3.5.2 *Forward contract valuation today—present value*

A common mistake in using forward contracts is forgetting that the forward valuation occurs on the forward expiry date. As we know from the time value of money, cash today is worth more than in the future. The implication is that, depending on the time period involved, the forward valuation overstates the true present value. We need to be very careful when dealing with forward contracts to ensure we know whether we are calculating forward or present values. This is an important consideration when comparing futures and forwards and is discussed in Section 3.6 below.

Calculating the present value of a forward contract can be performed using one of the present value formulae from Chapter 3. If we can apply a simple interest rate, then the present value of the forward contract value is:

$$\text{Present Value} = \text{Forward value} / (1 + r \times f / D)$$

where r is a simple interest zero-coupon rate between today and the forward expiry date. As already noted, if the interest rate that is being used is a money market interest rate it already is zero-coupon interest rate and can be entered directly into this calculation.

The distinction between forward and present values is demonstrated in *Exhibit 4.8*.

Exhibit 4.8**Forward Price and Valuation**

You have entered into the forward contract discussed in Exhibit 4.7 where the asset pays no income at a price of \$93.015. You decide to calculate the value of this contract after 30 days have passed. In that time interest rates have risen to 8% pa and the cash price of the security has declined to \$84.2. Calculate the current forward price, the forward valuation and then the present value of this contract.

Current forward price

$$\begin{aligned}
 S &= \$84.2 & r &= .08 & f &= 150 & D &= 360 \\
 F & & & & & & & \\
 &= & & & & & & S \times (1 + r \times f / D) \\
 &= & & & & & & 84.2 \times (1 + .08 \times 150 / 360) \\
 &= & & & & & & \underline{87.0067}
 \end{aligned}$$

Value at forward expiry date


$$\begin{aligned}
 \text{Forward value} &= \text{Current forward price} - \text{forward contract price} \\
 &= 87.0067 - 93.0150 \\
 &= \underline{-6.0083}
 \end{aligned}$$

Present value of forward contract

$$\begin{aligned}
 \text{Present value} &= \text{forward value} / (1 + r \times f / D) \\
 &= -6.0083 / (1 + .08 \times 150 / 360) \\
 &= \underline{-5.81448}
 \end{aligned}$$

Profit and loss risk profile of the forward contract — the pay-off diagram

Forward Purchase - Payoff Diagram



3.6 Valuation differences between forwards and futures

Futures contracts are a standardised form of forward contracts: a futures contract represents a commitment to buy or sell a fixed amount of the underlying financial asset, at a single date in the future. An OTC and futures contract with the same forward expiry date should have the same forward price.

While the pricing and valuation methodologies for the forwards and futures are similar, there are some key differences that need to be considered. In the following two chapters we will deal with specific differences between comparable OTC and exchange traded contracts.

The differences arise from the fact that futures contracts are subject to daily mark-to-markets and upfront initial margins (or performance bonds). These margining requirements have been created to minimise the risk that the clearing house of futures exchanges take to position takers. The effect of these margins, however, is to alter the valuation of futures versus forwards as they cause cashflows prior to the forward expiry date. In terms of the forward valuation/present valuation distinction in Section 3.5, a futures contract generates present values rather than forward values.

3.6.1 *The impact of daily mark-to-markets on valuation*

While equivalent forward and futures contracts will have the same value at the forward expiry date, it is often not realised that prior to the expiry date they have different present values. It is still common practice in many organisations to directly compare the valuations of forward and futures contracts without recognising that there is a difference in timing.

The valuation effect of daily mark-to-markets is to create a cashflow timing difference: in essence the forward values are being paid early. On any given day, the present value of a futures contract (represented by the mark-to-market gain or loss) is the same as the forward value on an equivalent forward contract.

For example, suppose we buy identical futures and forward contracts for expiry in one year. Today the difference between the current forward price and contract price at which we bought both contracts gives a value equivalent to \$100 profit. The present value of these two instruments will reflect the relative timing of the values. If the one year rate is 10%, then the present values are:

$$\begin{aligned} \text{Futures Profit} &= \$100 \\ \text{Forward Profit} &= \$100 / (1 + 10\%) \\ &= \$90.91 \end{aligned}$$

The difference in these values is important, as the futures contract in fact demonstrates greater sensitivity to movements in the forward price. In this example we can say that for every \$1 movement in the forward value, the futures contract will generate a present value change of \$1, while the present value of the forward contract will only change by 0.9091.

3.6.2 The impact of initial margins on valuation

Initial margins are a security deposit that must be paid by both buyers and sellers of futures contracts to the clearing house. These initial margins are held to cover any losses incurred by a defaulting position holder. In general the initial margins are set to cover the losses generated by a very large movement in the futures price over a 24 hour period. As a result the more volatile the futures price, the higher the level of initial margins.³

The impact of initial margins on valuation is not as straightforward to quantify as the daily mark-to-market because it depends on the level of initial margins and the rules of the futures exchange clearing house with respect to the types of collateral (for example, cash, government securities, precious metals, shares) allowed and the interest payment policy. These costs can be divided into two categories: interest cost and capital cost.

3.6.2.1 Interest cost

If the clearing house only accepts cash then the cost of initial margins is equivalent to the interest spread between the cost of funding the initial margin deposits and the interest paid on that deposit by the clearing house:

$$\text{Interest Cost} = \text{Initial Margin} \times (\text{Funding Cost} - \text{Clearing House Rate})$$

Suppose you enter a futures contract with a face value of \$100 million and an initial margin requirement of \$3 million. Your company borrows at the overnight money market rate while the clearing house pays the overnight rate minus 0.50% pa on your initial margin deposit. The additional funding cost of these initial margins is equivalent to 0.75% pa or \$15,000 annually. In terms of the total contract value this has increased the cost of financing the position by 0.015% pa or 1.5 basis points.

To incorporate the interest cost into the forward pricing formula we need to increase the cost of carry to reflect the additional financing cost of the initial margin. In this example we would add 1.5 basis points to the interest rate, which will have very little effect on the forward price, as is shown in the calculation below where the forward price from *Exhibit 4.8* is recalculated using an interest rate of 8.015% pa:

$$\begin{aligned} F &= S \times (1 + r \times f / D) \\ &= 84.2 \times (1 + .08015 \times 150 / 360) \\ &= 87.0119 \end{aligned}$$

This represents a change in the forward price of 0.0052—a very small impact. Often the interest cost is ignored by market participants because it is viewed as relatively unimportant.

If the clearing house accepts the lodgment of other forms of collateral such as bonds, shares and money market instruments without imposing any charges—a common practice in most large exchanges—then the interest cost will be the difference between the return on the asset and your organisation's cost of funds.

It is quite common for financial institutions to consider that providing securities as initial margin collateral incurs no cost. This is because they

3. For more detail on how initial margins are derived, see the Appendix to Chapter 7 in Martin, *op cit* n 1.

already hold the assets used for collateral for regulatory or investment reasons—they are simply re-using assets already held.

3.6.2.2 Capital cost

Whether initial margin collateral is in the form of cash or some form of security, there is a capital cost. By providing this collateral to the clearing house, the position taker is potentially transferring ownership of the assets and there is the possibility that those assets will not be repaid—initial margins represent a credit risk to the clearing house and some capital has to be set aside for that possibility.

The capital cost will depend on the capital allocation policies of the organisation involved. However, for a bank the Bank for International Settlements (BIS) Capital Adequacy Standards would view the initial margins deposited with the clearing house as a deposit with a corporation—requiring an allocation of capital equivalent to 8% of the deposit. Given the assumed cost of capital of an organisation then we can calculate the capital cost as follows:

$$\text{Capital Cost} = \text{Initial Margin} \times 8\% \times \text{Cost of Capital}$$

So, on the example above, if the cost of capital is assumed to be 20% pa then the annual capital cost is as follows:

$$\begin{aligned} \text{Capital Cost} &= \$3,000,000 \times 8\% \times 20\% \\ &= \$48,000 \text{ pa} \end{aligned}$$


In the same manner as the interest cost, the capital cost is considered an additional cost in financing the futures position and increases the cost of carry. Once again the capital cost is often viewed as too small to worry about and is ignored by some market participants.

4. INTEREST RATE FORWARDS AND FUTURES

4.1 Introduction

As we saw in the previous section, interest rate forwards and futures represent the largest single category of volume in financial derivatives. This is a recognition of the importance of these instruments as a day-to-day hedging and trading tool for all participants of the financial market place. While the volume of OTC forward contracts in some currencies is substantial (for example, US\$ Treasury Bonds), *Exhibit 4.9* shows that the global volume of interest rate futures is considerably higher than forwards. This is a reflection of the fact that short-term and long-term interest rate futures are the primary forward interest rate instruments in most currencies, whereas OTC forwards tend to be for more specialised use.

Exhibit 4.9
Composition of Interest Rate Forwards
Outstanding Volume - 1993



Source: GAO

In this chapter we will develop pricing and valuation models for forward and futures contracts on interest bearing financial assets.⁴ These models will be divided into categories reflecting the different characteristics of forwards on short-term and long-term debt securities. The categories used are as follows:

- forward rate agreements (section 4.2);
- short-term interest rate futures (section 4.3);
- bond forwards (section 4.4); and
- long-term interest rate futures (section 4.5).

For each model our starting point is the generalised model developed in Section 3. These models will then be adapted to the specific cashflow and convention characteristics of each instrument.

4.2 Forward rate agreements

4.2.1 General description

Forward rate agreements (FRAs) are the predominant form of OTC forward on short-term interest rate securities. They represent an agreement between two parties who wish to "fix" the interest rate on an underlying short-term security at a future date. FRAs do not have physical delivery, instead, any profits and losses are realised by way of a cash settlement at the end of the forward period. While the underlying instrument in an FRA is usually a short-term instrument with a term of three or six months, the forward period can range from one month to several years. An FRA is agreed

4. Another description of the underlying assets is "debt securities".

in terms of a forward interest rate as opposed to a forward price and the pricing formulae need to be adjusted to reflect this.

The liquidity of FRAs in most countries is very high, with most financial institutions providing market making services. Reflecting the level of activity, standard documentation and dealing terms and conditions have been developed in most countries.⁵ A summary of the general terms and conditions is provided in *Exhibit 4.10*.

5. Examples of this documentation and terms and conditions can be obtained from most bankers' associations or the local branch of ISDA in the relevant country. Otherwise, see S Das, *Swaps and Financial Derivatives* (2nd ed, IFR, 1994), pp 89-96.

**Exhibit 4.10
General FRA Terms and Conditions**

Item	Description
Broken period	A settlement period which differs in length from that used in fixing the interest settlement rate.
Buyer/borrower	The party wishing to protect against a rise in interest rates.
Cash settlement	There is no delivery under an FRA, instead any profits or losses are realised as a cash settlement on the settlement date.
Contract amount	The notional sum on which the FRA is based (that is, the principal).
Contract/trade date	The date the FRA is entered into.
Contract rate	The rate of interest agreed between the parties on the contract date (that is, the forward rate).
Maturity date	The date that the settlement period ends (that is, the maturity date of the security which notionally underlies the FRA).
Run	The period or term of the notional underlying security, usually three or six months.
Seller/lender	The party wishing to protect against a fall in interest rates.
Settlement date	The expiry of the forward period, the start of the notional underlying security and the day the settlement sum is paid.
Settlement period	The term of the notional underlying security represented by the number of days between the settlement date and the maturity date.
Settlement rate	The mean rate quoted by the specified reference banks for the settlement period of the notional underlying security. In US\$ based FRAs this is commonly given by the Reuters page "LIBO".
Settlement sum	The amount representing the difference between interest calculated at the contract rate and the settlement rate.

Source: BBA and AFMA

Under an FRA the two parties agree the interest rate (the “forward rate”) applying to a notional principal amount of an underlying money market security at a forward settlement date. Depending on how the relevant interest rate moves between the trade date and the forward settlement date, one of the parties will owe the other party a net settlement amount equivalent to the difference between the forward rate and the actual rate for the forward settlement date. The party which benefits from a fall in interest rates is defined as the “lender” or “seller”. The other party which benefits from a rise in interest rates is the “borrower” or “buyer”.

FRAs are normally quoted in terms of monthly combinations of the time to the forward settlement date and the time to maturity of the notional underlying security. For example, an FRA with one month to forward settlement on a three month security is referred to as a “1 × 4”. On quote vendor services such as Reuters, FRA dealers generally provide indications of FRA rates in terms of the standard combinations set out below:

Tenor	Rate	Description of Forward
1 × 4	7.35	A three month security starting in one month
3 × 6	7.25	A three month security starting in three months
6 × 9	7.23	A three month security starting in six months
3 × 9	7.24	A six month security starting in three months
6 × 12	7.20	A six month security starting in six months

An example of an FRA transaction starting in three months on a three month security is summarised in *Exhibit 4.11*.

Exhibit 4.11
A 3 × 6 Forward Rate Agreement

<p>Trade Date</p> <p>Agree forward rate of 6% pa on US\$10m.</p>	<p>Forward Settlement Date</p> <p>Actual settlement rate is 8% pa.* Buyer pays seller \$49,019.61.</p>	<p>Notional Maturity Date</p> <p>Maturity is only notional.</p>
---	---	--

Settlement amount	=	$\frac{(.08 - .06) \times 90 / 360 \times 10m}{1 + .08 \times 90 / 360} = 49,019.61$
-------------------	---	--

Note: *The settlement rate is commonly set two business days prior to the settlement date.


4.2.2. *Synthetic replication of an FRA*

To understand the pricing of an FRA we will look at how an FRA can be synthetically replicated using cash instruments. The example in *Exhibit 4.11* is from the point of view of the buyer/borrower, and it is agreeing on the interest rate of a three month borrowing commencing in three months time. This can be replicated by borrowing on the trade date for six months and investing the funds raised for three months till the forward settlement date as is shown in *Exhibit 4.12*.

Exhibit 4.12


Synthetic Replication of an FRA Buyer/Borrower Position

Trade date



Borrow \$10m for six months from a bank where interest is paid quarterly at 5.75% pa. Invest for three months at 5.50% pa. This creates a net interest cost, or cost of carry, over the first three months equivalent to approximately 0.25% pa.

In three months



The three month investment matures and is repaid using the funds received. The borrowing is repaid in three months. As interest rates have risen a profit is crystallised on the transaction.

Exhibit 4.12—continued

<i>Forward rate calculations</i>	
Investment	=
at three months	= $10\text{m} \times (0.055 \times 90 / 360)$ 137,500.00
Borrowing interest	=
at three months	= $(1 + 0.0575 \times 90 / 360)$ 143,750.00
Net cost at	=
three months	= 6,250.00
Principal left at	=
three months	= 9,993,750.00
Borrowing at	=
six months	= 10,143,750.00
Forward rate	=
	= $(10,143,750 / 9,993,750 - 1) \times 360 / 90$ 6.00%
<i>Settlement sum calculation</i>	
In the synthetic replication the settlement sum will be equivalent to the net cashflows after the borrowing has been repaid early at three months.	
Settlement amount	=
at three months	= Principal left at three months minus the present value of the remaining interest on the borrowing (at 8% pa) $9,993,750 - 10,143,750 / (1 + .08 \times 90 / 360)$ _____48,897.06 *

Note: * There is a small rounding difference of \$122.55 between this calculation and Exhibit 4.11. This arises because the forward interest rate in this example is actually 6.00375% pa.

As with any forward transaction, a cost of carry will be created depending on the difference between the three and six month interest rates. The forward interest rate will be given by the six month interest rate adjusted for the cost of carry. So, if the six month interest rate is 5.75% pa⁶ and the three month rate is 5.5% pa then the cost of carry is 0.25% pa. The forward interest rate is consequently 6.00% pa. At the end of three months there has been a net interest cost of \$6,250 which effectively means the principal amount of the original borrowing remaining is \$9,993,750.

To replicate the cash settlement of the FRA on the forward settlement date, the borrowing is repaid three months early using the principal amount available at three months. The cost of repaying this borrowing is the present value of the principal or interest accrued at the end of the borrowing. The net cashflow arising from repaying the borrowing is \$48,897.06, or approximately the same as the same strategy using an FRA in *Exhibit 4.11*.

Using the terminology developed in Sections 2 and 3:

1. The forward rate in this example is the forward price and is derived using the now familiar cost of carry concept.
2. The settlement sum or amount is the same as the forward value.
3. The present value of an FRA is given by taking the present value of the settlement sum.

4.2.3 A model for FRA forward interest rates

FRAs and interest rate futures introduce a number of new characteristics for the general model developed in Section 3:

- an underlying instrument with a limited life, where
- the calculation is expressed in terms of interest rate not price.

In this section we will convert the generalised price formulae from Section 3 into a formula that generates a forward interest rate on a security which accrues interest for a limited period.

An FRA is an instrument in which the underlying asset is cash provides a constant income in the form of interest payments. We know that the underlying asset is a security which pays interest on a principal amount, S , from today until its maturity date. This future value of the cashflows on this security can be expressed as:

$$FV = S \times (1 + q \times d / D)$$

where the asset income, q , is the yield to maturity on the asset expressed as a per cent per annum and d is the number of days from today until maturity of the asset.

Using the forward pricing model with constant income in Section 3.4.2 we know this can be expressed as:

$$F = S \times (1 + (r - q) \times f / D)$$

where r , the financing cost, is the interest rate over the forward period. The cash price, S , is the principal value of the security and the forward price is

6. To avoid compounding differences this rate, while fixed for a term of six months, is compounded quarterly.

this amount adjusted for the cost of carry. Our aim is to express this same concept in terms of a forward interest rate calculation. Essentially we need to incorporate the cost of carry over the forward period into the interest calculation from the forward settlement date to the maturity date of the underlying security.

In simplistic terms, the interest on the forward security will be equivalent to the difference between the interest earned between today and the forward settlement date and between today and the maturity date of the underlying security. Given that we know the values of q and r , then the amount of interest earned by the forward security can be expressed as:

$$\text{Forward Interest} = S \times (q \times d / D - r \times f / D)$$

The forward interest rate can then be expressed as:

$$\text{Forward Rate} = \frac{\text{Forward Interest}}{\text{Forward Price}} \times \frac{D}{(d-f)}$$

If we insert the formulae above then we have:

$$\text{Forward Rate} = \frac{S \times (q \times d / D - r \times f / D)}{S \times (1 + (r-q) \times f / D)} \times \frac{D}{(d-f)}$$

which simplifies to:

$$\text{Forward Rate } (r_f) = \frac{(q \times d / D - r \times f / D)}{(1 + (r-q) \times f / D)} \times \frac{D}{(d-f)}$$

This formula approximates the calculation performed in *Exhibit 4.12*. However, it ignores the timing of cashflows and the compounding of interest. In most forward interest rate calculations interest rates r and q have different compounding frequencies, which means they cannot be directly compared.

A common method of avoiding the compounding problems is to convert the interest rates into continuously compound rates. This avoids any compounding differences and simplifies the forward rate calculation. We know that the future value of using a continuous rate is as follows:

$$FV = S \times \exp(q \times d / D)$$

Further, we know that the future value of an amount invested for the full term d and an amount invested for the combined term of f and $(d-f)$ must be the same. Using the future value formula we can express this as follows:

$$S \times \exp(q \times d / D) = S \times \exp(r \times f / D + r_f \times (d-f) / D)$$

If we cancel S and take the natural logarithm of both sides of this equation, this simplifies to:

Forward Interest Rate: Continuous Compounding

$$r_f = \frac{q \times d / D - r \times f / D}{d / D - f / D}$$

Where

- r_f = Forward interest rate % pa
 r = Interest rate to the forward settlement date % pa
 q = Interest rate to the maturity date % pa
 D = Day count basis (365 or 360)
 f = Number of days to the forward expiry date
 d = Number of days to the maturity date of the underlying instrument

Examples of underlying assets: FRAs and short-term interest rate futures

Market interest rates are rarely quoted in continuously compounded form, to use this model we need to convert to and from continuously compounded rates. *Exhibit 4.13* provides an example of a spreadsheet which calculates forward interest rates using the continuous compounding method.

Exhibit 4.13 Forward Rate Agreement Calculator		
<i>Spreadsheet example</i>		
Field	Cell	Cell: Formula
Inputs		
Trade date	01-Nov-95	\$E\$8 :
Forward settlement date	30-Jan-96	\$E\$9 :
Underlying maturity date	29-Apr-96	\$E\$10 :
Spot rate to forward settlement date %	5.5000	\$E\$11 :
Frequency (1, 2 or 4)	4	\$E\$12 :
Interest rate for maturity %	5.8050	\$E\$13 :
Frequency (1, 2 or 4)	2	\$E\$14 :
Outputs		
Term to forward settlement in days — f	90	\$N\$16 : +E9-E8
Term to maturity in days — d	90	\$N\$17 : +E10-E8
Continuous rate to settlement date % — r	5.4625	\$N\$18 : =LN(E6/(E7*100)+1)*E7*100
Continuous rate to maturity % — q	5.7224	\$N\$19 : =LN(E8/(E9*100)+1)*E9*100
Forward rate % — continuous compounding	5.9822	\$N\$20 : =(E11*(E5-E3)-E10*(E4-E3))/(E5-E4)
— quarterly compounding	6.0271	\$N\$21 : =4*(EXP(E12/(4*100))-1)*100
— semi-annual compounding	6.0725	\$N\$22 : =2*(EXP(E12/(2*100))-1)*100
— annual compounding	6.1647	\$N\$23 : =(EXP(E12/100)-1)*100

The approach in this model is to take the interest rates based on market rates and adjust them to continuous rates to calculate the forward rate. Once the continuously compounded rate has been calculated this can be re-converted to any compounding basis required.

4.2.4 A model for FRA valuation

The value of an FRA on the forward settlement date is the difference between the agreed contract rate in the FRA and the prevailing reference interest rate for the remaining term to maturity of the underlying security (the "settlement rate").

There are two general methods of calculating the forward amount, depending on the conventions in the money market. In most markets where money market securities are traded in terms of face value or discounted using the discount method, such as in the United States and United Kingdom, the settlement formula is as follows:

FRA Settlement Calculation: Full Face Value	
If $r_s > r_c$ then the settlement sum is seller pays buyer	$= \frac{(r_s - r_c) \times d / D \times S}{1 + r_c \times d / D}$
If $r_s < r_c$ then the settlement sum is buyer pays seller	$= \frac{(r_c - r_s) \times d / D \times S}{1 + r_c \times d / D}$
Where	
r_c	= Contract rate % pa
r_s	= Settlement rate % pa
d	= Settlement period (days till maturity of underlying security)
D	= Day count basis (365 or 360)
S	= The contract amount
FRA markets commonly using this method: US\$ and most European currencies	

This calculation is derived so that all obligations of the FRA can be terminated on the forward settlement date rather than the maturity date of the notional underlying security. That explains why the difference between the contract and settlement interest amounts is calculated at the maturity of the notional underlying security and then present values this difference to the forward settlement date.

In markets where money market securities are traded at a discount to face value using the yield method, such as in Australia and New Zealand, the forward value is based on a discounted proceeds calculation as follows:

FRA Settlement Calculation: Discounted Face Value

If $r_s > r_c$ then the settlement sum is

$$\text{seller pays buyer} = \frac{S}{1 + r_c \times d / D} - \frac{S}{1 + r_s \times d / D}$$

If $r_s < r_c$ then the settlement sum is

$$\text{buyer pays seller} = \frac{S}{1 + r_c \times d / D} - \frac{S}{1 + r_s \times d / D}$$

Where

- r_c = Contract rate % pa
- r_s = Settlement rate % pa
- d = Settlement period (days till maturity of underlying security)
- D = Day count basis (365 or 360)
- S = The contract amount

FRA markets commonly using this method: A\$ and NZ\$

4.2.4.1 Forward value

Prior to the forward settlement date, the forward value is given by the relevant settlement calculation above. However, instead of the settlement rate, r_s , the prevailing forward rate, r_f , is used. As we noted in Section 3, at initial execution the forward value of an FRA will be zero as the contract rate and the prevailing forward rate are the same. As time passes and the forward rate changes so will the forward value of the FRA. To illustrate this point, *Exhibit 4.14* extends the previous example and examines the change in value of the FRA contract over the forward period using both settlement calculations.

Exhibit 4.14

Forward and Present Value of an FRA

You have entered into the FRA from Exhibit 4.11 as the buyer/borrower. Two months have passed. Calculate the forward rate, the forward value and the present value using the market data provided. Calculate values using both the full face value and discounted face value settlement methods.

Forward Period Remaining

3 months → Enter into 3 × 6 buyer/borrower FRA at 6% pa

1 month	Market data	1 month rate	=	6.00%
		4 month rate	=	6.50%

Forward rate calculation

f	=	30
d	=	120
r	=	5.99% pa (continuous compounded rate)
q	=	6.43% pa (continuous compounded rate)
rf	=	$.0643 \times 120 / 360 = .0599 \times 30 / 360$
		$120 / 360 = 30 / 360$
		6.58% pa (continuous compounded rate)

Then converted to a quarterly rate
 rf = 6.63%

Exhibit 4.14—continued

	<i>Forward values</i>
rc	= 6.00%
dc	= 90
Full face value	= $(0.0663 - 0.06) \times 90 / 360 \times 10m$ $1 + 0.663 \times 90 / 360$ 15,493.20
Discounted face value	= $\frac{10m}{1 + .06 \times 90 / 360}$ 15,264.24
	= $\frac{10m}{1 + .06 \times 90 / 360}$
	= 15,264.24
	<i>Present values</i>
Discount the forward values to today using the prevailing one month interest rate	
Full face value	= $\frac{15,493.20}{1 + .06 \times 30 / 360}$ 15,416.12
Discounted face value	= $\frac{15,262.24}{1 + .06 \times 30 / 360}$ 15,188.30

4.2.4.2 Present value

The present value of an FRA is easily calculated once the forward value has been generated using the standard present value formulae. An example of this calculation is provided in *Exhibit 4.14*. The present value should be calculated using a zero-coupon interest rate.

4.2.5 FRA risk characteristics

A forward rate agreement is the right to purchase or sell a short-term money market instrument at some date in the future. Correspondingly, the sensitivity to movements in interest rates of an FRA is very similar to the money market instruments.

In the case of an FRA the duration and convexity is equivalent to the underlying instrument. The point value of a basis point (PVBP), however, is less than that of the underlying instrument. An FRA generates gains and losses on the forward settlement date equivalent to the underlying security, however, these amounts are present valued in the PVBP and so are consequently smaller.

An example of the PVBP is provided later in the chapter in *Exhibit 4.18*, where an FRA PVBP is calculated and compared to short-term interest rate futures contracts.

4.3 Short-term interest rate futures

4.3.1 General description

Short-term interest rate futures represent standardised, exchange traded forward contracts on money market instruments. In general, most major currencies have one futures contract on a tradeable short-term money market instrument such as a bank deposit or bank bill. The pricing and valuation of these instruments is very similar to FRAs and the two markets can often be viewed as direct substitutes. The global volume in these instruments is enormous, representing the largest single category of futures contract. A list of the major short-term interest rate futures contracts are listed in *Exhibit 4.15* along with total volume for 1994.

Exhibit 4.15
List of Short-term Interest Rate Contracts and Volumes

Contract	Currency	Exchange(s)	1994 Futures Volumes	
			No of contracts	Face Value (Bn)
Bank accepted bills*	A\$	SFE	9,369,008	6,933
3 month Euro-Swiss franc	CHF	LIFFE	1,698,736	1,493
FIBOR futures	DEM	DTB	428,516	305
3 month Euro-deutschmark	DEM	LIFFE	29,312,222	62,575
3 month ECU interest rate	ECU	LIFFE	622,457	814
PIBOR 3 month	FFR	MATIF	13,176,354	2,717
3 month Sterling interest rate	GBP	LIFFE	16,603,152	25,821
3 month Euro-yen	JPY	TIFFE	37,425,846	367
3 month Euro-yen	JPY	SIMEX	6,820,673	67
3 month Euro-lira	ITL	LIFFE	3,456,437	2
Bank accepted bills	NZ\$	NZFOE	608,460	393,963
MIBOR 90	ESP	MEFF	3,730,008	155
3 month Eurodollar	US\$	CME	104,823,245	524,116
3 month Eurodollar (fungible with CME)	US\$	SIMEX	8,687,969	4,344
3 month Eurodollar	US\$	LIFFE	91,738	9,174
1 month Eurodollar	US\$	CME	1,911,184	191,118
90 day T-bills	US\$	CME	1,020,491	510
		Total	239,786,496	1,224,476

Note: *SFE contract upsized from 500,000 to 1,000,000 in April 1995

The benchmark contract for short-term contracts is the Eurodollar contract traded on the Chicago Mercantile Exchange (CME). As the table in *Exhibit 4.15* shows the Eurodollar contract is the most heavily traded reflecting its status as the primary hedging vehicle for short to medium term exposures. It is also traded on the London International Financial Futures and Options Exchange (LIFFE) and the Singapore International Money Exchange (SIMEX). The SIMEX contract is “fungible” with the CME contract, which means contracts traded on the two exchanges can be offset.

The Eurodollar contract was the first of the short-term futures contracts when it was listed in 1981. Most other short-term interest rate futures contracts have been a copy of the Eurodollar contract with only the currency, settlement interest rate and face value changed. The A\$ and NZ\$ bank bill contracts traded on the Sydney Futures Exchange (SFE) and the New Zealand Futures and Options Exchange (NZFOE) respectively, are the only contracts which have different valuation formulae.

In order to familiarise ourselves with short-term interest rate contracts we will firstly review the features of the Eurodollar contract and then examine the differences with other contracts.

4.3.1.1 Eurodollars

The Eurodollar is a cash settled contract on a three month Eurodollar time deposit. The name “Eurodollar” derives from the fact that it is a forward contract on a US dollar money market instrument traded in Europe (or London to be more specific). The importance of the contract reflects the importance of the US\$ in global financing and the willingness of US-based market participants to use futures.

The CME lists contracts to expire in quarterly rests in March, June, September and December. Currently, there are 40 consecutive quarters listed, that is, expiries out to 10 years. The Eurodollar has obvious appeal to corporations, banks and fund managers with short-term interest rate exposures. However, a substantial driving force behind Eurodollar volumes is from organisations with medium-term exposures such as interest rate swap market makers.

The contract expires on the third Monday of the delivery month and is cash settled against a three-month London Interbank Offered Rate (LIBOR) in a similar fashion to a US\$ FRA. If the current month is not a quarterly delivery month then a single “spot” contract is also listed to ensure traders have a very short-term instrument. For example, after the March contract expires an April contract is listed. This “spot” contract concept is currently only offered on the Eurodollar.

The price of a contract is expressed as:

$$\text{Futures Price} = 100 - \text{Interest Rate} \times 100$$

So, if the current interest rate for a Eurodollar deposit starting on the futures expiry date is 5.00% pa, then the futures price is 95.00. The aim of quoting in terms of price rather than yield is primarily to keep interest rate contracts in line with other price-based contracts on bonds, shares and commodities.

A buyer of a Eurodollar contract gains if the futures price rises (interest rate falls) above the price at which they purchase and the seller gains if the

price falls (interest rate rises). Be careful when comparing futures to FRAs as the terminology is opposite; a Eurodollar futures buyer is equivalent to an FRA seller/lender as they both benefit from a fall in interest rates.

While the detailed contract specifications of all short-term interest rate futures contracts are provided in the Appendix to this chapter the major features of the Eurodollar contract are summarised below:

Summary of Eurodollar Futures Specifications	
Feature	Description
Underlying	90 day Eurodollar time deposit
Face Value	US\$1,000,000
Delivery Months	March, June, September, December and spot month
Delivery Method	Cash settled
Settlement Rate	LIBOR rate for three month Eurodollar deposits on the last trading day
Last Trading Day	Third Wednesday of the delivery month
Quotation Method	100 minus the rate of interest
Valuation Formula	Term deposit
Tick Size	The value of each price point is \$25
Margining	Initial margin (currently \$500/contract) and daily mark-to-market

The details of most of the other short-term interest rate contracts are similar except for differing face values as is summarised below:

Contract	Exchange	Face Value
90 Day T-Bills	CME	1,000,000
Bank Accepted Bills	SFE	1,000,000
3 Month Euro-Swiss Franc	LIFFE	1,000,000
FIBOR Futures	DTB	1,000,000
3 Month Euro-deutschmark	LIFFE	1,000,000
3 Month ECU Interest Rate	LIFFE	1,000,000
MIBOR 90	MEFF	10,000,000
PIBOR 3 Month	MATIF	5,000,000
3 Month Sterling Interest Rate	LIFFE	500,000
3 Month Euro-Yen	TIFFE	100,000,000
Bank Accepted Bills	NZFOE	500,000

For most contracts the underlying instrument is the same as the Eurodollar that is, a three month deposit on a discount security where interest is

calculated using the discount method. However, in the case of the SFE and NZFOE contracts the underlying instrument is a bank bill, which is a discount security valued using the yield formula.⁷ This has an impact on valuation, as discussed in section 4.3.3.

4.3.2 *A model for futures prices*

The price of a futures contract is equal to 100 minus the forward interest rate. So, the futures pricing model will be based on the forward pricing models from section 3 and the FRA model from section 4.2.3.

The method of synthetically replicating a futures contract is exactly the same as an FRA (see section 4.2.2). However, as we have already noted, futures contracts have a different cashflow profile to similar forward contracts because of initial margins and the daily mark-to-market of gains and losses.⁸ As a result of initial margins, futures may need to incorporate a small funding cost, while the impact of the mark-to-market is unknown, as it will depend on the level of interest rates over the life of the futures contract.

In summary, the short-term futures contract price is primarily determined by the prevailing forward rate using the formula in section 4.2.3 above. There is, however, an element of the interest rate which will not be known till expiry of the contract due to the unknown funding requirements during the life of the contract. This can be summarised as follows:

$$\text{Futures Price} = 100 - (\text{Forward Rate} + \text{Funding Adjustment})$$

For contracts with a forward period of up to six months the differences between short-term interest rate futures and FRAs are very small, as well as unknown, and can often be ignored.⁹ However, for longer term futures consideration should be given to incorporating the possible funding consequences of a futures contract. It has to be remembered that this is just an estimate; it is common for market users to estimate the “worst case” funding cost requirement and incorporate that into their estimate of the effective forward rate given by short-term interest rates.

- A discount formula calculates interest based on the future face value and then deducts this from the face value of instrument. A yield formula is a present value of the future face value of the contract. For more on this distinction see Martin, *op cit* nl, Ch 3.
- See section 3.6 for a more detailed discussion on this point.
- In this case the spreadsheet model provided in Exhibit 4.13 is appropriate.

Exhibit 4.16

Synthetic Replication of a Sold Eurodollar Futures Position

In the following example the forward interest rate provided by a sold Eurodollars position is compared against the forward interest rate provided by a forward rate agreement. The Eurodollar contract expires in 1 year's time for a term of 90 days and it is sold today at a price of 92.74. Over the year interest rates fall and the futures price converges to a three month rate of 5% pa. The table summarises the resulting cashflows by quarter.

To simplify the analysis all interest rates are continuously compounded and converted to a 365 day basis. Also, the futures price is assumed to remain steady until the end of the quarter, at which time it falls to the price shown in the table.

Dates	
Trade date	24-Oct-96
Forward settle	24-Oct-97
Maturity date	22-Jan-98

Market Rates	
Overnight/3 month/1 year =	6.00%
1.25 year rate =	6.25%

Implied FRA Rate Calculations

f	365		
d	455		
r	6.00%	Forward rate =	7.2639%
q	6.25%		

Exhibit 4.16—continued

	Interest Rates		Synthetic Replication		Futures Replication						
	Qtr	Overnight rate and 3 mth rate	Futures Price	Current fwd rate	Borrow for 1.25 years @ 6.25%	Invest for 1 year	Mark-to-Market	Funding requirement	Interest on initial margins (o/n - 1% pa)	Quarterly interest	Total interest
Today	0	6.00%	92.74	7.26%	1,000,000.00	(1,000,000.00)	1,100.00	(500.00)	8.51	18.72001	
	1	6.75%	92.30	7.70%			1,100.00	600.00	9.46	41.99346	
	2	7.50%	91.86	8.14%			1,100.00	1,718.72	10.42	70.03446	
	3	8.25%	91.42	8.58%			1,050.00	2,860.71	11.38	101.9596	
Forward Maturity	4	9.00%	91.00	9.00%	(1,081,026.40)	1,061,836.55	1,050.00	3,980.75	11.38	104.2797	336.9872
	5	9.00%						4,082.71	11.38		

Effective Forward Rate from Futures Contract
 Investment return after 1 year = 1,061,836.55
 Initial borrowings after 1.25 years = 1,081,026.40
 Extra funding cost of futures = (336.99)
 Total borrowings after 1.25 years = 1,080,689.41
 Effective interest cost = 7.1374%
 Difference between FRA and futures = -0.1264%

Exhibit 4.16 gives an example of the impact that the funding requirements of a futures contract can have on the effective forward rate. In this case we examine the synthetic replication of a single sold Eurodollar contract. As with the FRA synthetic replication we borrow till the maturity date of the underlying (1.25 years) and invest for the forward period (1 year). The additional complication of the future contract is the upfront initial margin of \$500 and the mark-to-markets based on the prevailing forward price. In the example the Eurodollar future is sold at a price of 92.74 with one year till expiry. It is assumed that the forward price rises over the life of the futures contract to settle at 95.00 (a three month interest rate of 5.00% pa). As well as the initial margin this generates substantial funding requirements for a sold position. The interest costs of funding these cashflows is incorporated into the effective futures forward rate calculation. Under this scenario, the effective interest rate by nearly 10 bp over the equivalent FRA forward rate.

The difference between the FRA and the effective forward rate in the futures contract is dependant on the actual path taken by interest rates over the forward period. The effective forward rate on the example above is calculated for a range of outcomes in *Exhibit 4.17*.

Exhibit 4.17

Effective Futures Rate Over a Range of Outcomes

Using the example from Exhibit 4.16, the effective forward rate given by the futures contract is calculated for a range of final three month interest rates.

3 Month Rate After 1 Year	Effective Futures Rate	FRA/Futures Difference
5%	7.3620%	0.0981%
6%	7.3273%	0.0634%
7%	7.2784%	0.0146%
8%	7.2152%	-0.0487%
9%	7.1374%	-0.1264%

It is important to note that these funding issues also face the buyer of a futures contract. However, in the case of the buyer a fall in rates generates positive cashflows and improves the effective forward interest rate.

It is important to realise that these funding problems effect both buyers and sellers of futures contracts. The effective forward rates achieved are the same but they have a different effect. For a seller, if interest rates fall the higher effective forward rate increases its cost of borrowing. However, for a buyer, if interest rates fall, the higher effective forward rate represents an improvement in their investment yield.

This difference in cashflows also gives rise to so-called “convexity adjustments” when hedging OTC products such as FRAs and swaps. This issue will be discussed in the following section on valuation.

4.3.3 Valuation of short-term interest rate futures contracts

As we saw above, the valuation of futures contracts can be a source of considerable confusion. It is important to remember that futures contracts have the peculiar property of constant PVBP—there is no distinction between present and future values. Whereas in cash financial assets and OTC derivatives there is a difference equivalent to the time value of money.

In this section we will consider the formulae for determining the contract value of futures contracts and then examine the impact of constant PVBP when using FRAs.

4.3.3.1 Contract values

All open futures contracts are subject to at least a daily mark-to-market, sometimes more.¹⁰ Whenever a mark-to-market is made the valuation method is unchanged and is based on the same formula used to determine the final settlement value of the futures contract. In the case of a Eurodollars contract, the underlying security is the interest on a 90 day deposit—each point change in the futures price is equivalent to a 0.01% pa change in the interest rate in the underlying security. We can then express the value of this type of contract as:

$$\text{Contract Value}_{\text{ED}} = \text{Face Value} \times (100 - \text{Price}) \times d / D / 100$$

where price is the prevailing futures price. So if the futures price is 96.24 then the contract value of a Eurodollar contract is:

$$\begin{aligned} &= 1,000,000 \times (100 - 96.24) \times 90 / 360 / 100 \\ &= 9,400. \end{aligned}$$

On any given day the mark-to-market gain or loss will be equivalent to the difference in the contract value at the previous mark-to-market price and today's market price.

Most market participants recognise that the contract value formula always implies a constant PVBP, commonly known as the “tick value” in futures markets, equivalent to \$25, regardless of the time to expiry of the futures contract. As we saw in section 3.5.3, this is quite an unusual property, and is commonly referred to as “non-convexity”. In fact it represents a form of slightly negative convexity as convexity relates to the percentage change in price and in these contracts the PVBP is declining as a percentage of the present value as interest rates fall. It is more appropriate to describe this property as “non-dollar convexity”.

This formula can be applied to all short-term interest rate contracts except the two SFE and NZFOE bank bill contracts. In the case of these contracts the contract value is given by a discount security formula using the yield method:

$$\text{Contract Value}_{\text{BB}} = \text{FV} / (1 + (100 - \text{Price}) \times d / D / 100)$$

10. Some exchanges such as the CME and CBOT mark-to-market twice a day, once at the close of business and once at the end of morning trading. Further, most exchange clearing houses require “intra-day” margins when market conditions are volatile—effectively marking-to-market during the day.

So in the case of the SFE contract a price of 96.24 is:

$$= 1,000,000 / (1 + (100 - 96.24) \times 90 / 365 / 100)$$

$$= 990,813.93.$$

In these instruments the PVBP is not constant as it changes slightly according to the level of interest rates. At the price in the example above the tick value is \$24.21, however, at a price of 86.24 the tick value is \$23.06. In practice most market participants assume the tick value is approximately \$24. Because of the underlying yield discount method, bank bill future contracts are displaying a “positive” PVBP relationship as PVBP rises with a fall in yields. It is important to realise that this convexity is arising just from the valuation of the underlying three month bank bill contract—it is not related to the term to forward expiry at all and is just a reflection of the bank bills convexity on the futures expiry date. As such, it shares the same property as the Eurodollar contract: at a given yield the PVBP of the futures contract will be the same regardless of the number of days to the future’s expiry date. In the following section we will refer to this property as “constant PVBP”.

4.3.3.2 Time value of money, hedge ratios and convexity adjustments

Futures contracts do not appear to obey the rules of the time value of money (TVM). At a given yield the value of a contract is the same today as in the future—there seems to be no compensation for the passing of time. Of course, this is not the case. We need to remember that futures contracts represent a highly structured form of financial derivative. The constant PVBP is a result of the risk management practices of exchange clearing houses; it has not been specifically designed to work this way.

One of the “cornerstones” of valuation is that all derivatives must obey the TVM. As a result, when using futures contracts they must be used to conform with the TVM characteristics for the purpose they are being used. So, when using futures contracts to hedge an instrument such as physical financial assets or OTC derivatives, the number of futures contracts needs to be adjusted in accordance with the TVM characteristics of the instrument being hedged. This is an extremely important rule, and if it is not followed then it will lead to over or under-hedging.

In this section we will consider using short-term interest rate futures to hedge FRAs. For a future and FRA with the same forward settlement date and notional maturity date, the forward interest rate represents is very similar—with differences arising due to the funding consequences of the futures contract. So, as market interest rates change the current forward interest rate used to determine both FRA rates and futures prices can be viewed as identical. Any differences in the two contracts will result from the different treatment of present values.

This is highlighted in *Exhibit 4.18* where the PVBP of an A\$ FRA and equivalent bank bill future are compared. While the future value is the same, the FRA PVBP is lower, reflecting the TVM. If we wish to hedge the FRA with futures contracts the aim is to ensure that any profits and losses today on the FRA are offset by the futures contracts. The appropriate amount of futures to hedge the FRA is that amount which equates the PVBPs of the two instruments—in this case 91 contracts. This equating of PVBPs is often referred to as determining the “hedge ratio”. (Note that as time passes the

PVBP of the FRA will rise toward the futures PVBP and the number of futures contracts will need to rise correspondingly.)

Exhibit 4.18

Futures and FRAs: Dealing with Different PVBPs

It is 13 June 1995. Calculate today's PVBP on bank bill futures (BAB) contract and FRA is listed below for face values of A\$1 million. Using this information if you had bought \$100 million face value of FRAs, how many futures contracts would you sell to hedge the price risk?

Current Market Data

Instrument	Tenor	Underlying Days	Expiry Date	Current Yield	PVBP Yield
1. FRA	15/18	90	13-Sep-96	7.58%	7.59%
2. BAB	Sep-96	90	13-Sep-96	7.58%	7.59%

Note: Zero coupon rate to 13 Sep 96 = 7.90%

Calculations

Instrument	Current Value	PVBP Value	Future Value	Present Value	PVBP
1. FRA	981,652.51	981,628.75	(23.76)	(21.61)	21.61
2. BAB	981,652.51	981,628.75	(23.76)	(23.76)	23.76
		Difference		(2.15)	

Exhibit 4.18—continued

Number of futures contract to hedge \$100m FRAs.
 We assume that the futures price and FRA are very closely correlated. And then apply the hedge ratio formula developed in section 3.5.3.

Hedge Ratio	=	PVBP(FRA) / PVBP
	=	21.61 / 23.76
	=	0.9093

So for every \$1 face value of FRA we would sell 0.9099 BAB contracts.

Number of contracts	=	FRA face value × Hedge Ratio / BAB face value
	=	\$100,000,000 × 0.9093 / \$1,000,000
	=	90.93
	=	91 contracts (rounded to nearest whole contract)

This analysis suggests that FRAs and futures can be equated by adjusting for the differences in PVBP. This is generally true, however, there is another small effect that FRA and swap market makers call the “convexity adjustment”.¹¹ As the name implies it is an adjustment to take account of the differences in the convexity characteristics of futures contracts and other closely related OTC derivatives. This adjustment starts with the recognition that the interest earned on futures mark-to-markets is negatively correlated with the futures price. That is, as interest rates fall futures prices rise. From the point of view of a short-seller of futures contracts this means if interest rates fall then they will pay mark-to-market losses. However, the interest rate to fund these losses is *lower* than the rate prevailing when they executed the transaction. On the other hand, if interest rates rise, the seller receives mark-to-market profits and earns a *higher* rate of investment interest. This creates a natural bias in futures contracts which favours short-sellers whether interest rates rise or fall. This bias works against a buyer of short-term futures contracts.

As with the funding adjustment, the exact effect of this convexity effect is not known till the expiry of the futures contracts. Estimating the convexity adjustment depends on determining an expected path for interest rates over the life of the futures contract—not a straightforward task. As with the funding adjustment the approach is to make an estimate of the net convexity approach and convert it into a forward interest rate equivalent. The implied futures yield should then be equivalent to the FRA rate plus the convexity adjustment.

A simple example of showing how the convexity adjustment can be calculated is provided in *Exhibit 4.19*. In this case we look at the outcome of the previous hedging example assuming interest rates either increase or decrease by 1% pa. The convexity effect can be seen at work in this example—regardless of whether interest rates rise or fall the futures generate a net benefit of around 1 bp. This is an extremely simple example as it assumes that the zero-coupon and forward rates rise or fall by 1% pa on the first day and stay there. It is a complex estimation problem for a small increase in accuracy and as such is generally of concern only to FRA and swap market makers.

11. Convexity is a measure of the sensitivity of duration to a movement in the yield on the underlying instrument. For more details see Martin, op cit n 1, Ch 3.

Exhibit 4.19

Futures and FRAs: Estimating Convexity Adjustments

To see the convexity adjustment at work let us look at the cashflows generated by the transactions in Exhibit 4.18. We will examine the cashflows impact if interest rates rose by 1% pa or fell by 1% pa on the trader date and then stayed there till the forward expiry and examine the relative costs of futures and FRAs.

Original Transactions — 23 June 95

Instrument	Amount	Forward Settlement	Maturity Date	Traded Rate/ Price
1. FRA	100,000,000	13-Sep-96	12-Dec-96	7.58%
2. BAB	91 contracts	13-Sep-96	12-Dec-96	92.42

Interest Rates

	Original	After 1% rise	After 1% fall
FRA	7.58%	8.58%	6.58%
BAB	7.58%	8.58%	6.58%
1.25 year ra	7.90%	8.90%	6.90%

Calculations for 1% Rise in Rates

Instrument	Traded value Value (1)	Settlement Value	Settlement Amount	Present Value (2)
1. FRA	98,165,251.11	97,928,214.59	(237,036.52)	(215,545.25)
2. BAB	(89,330,378.51)	(89,114,675.28)	215,703.23	215,703.23
		Difference		157.98

Exhibit 4.19 — continued

Calculations for 1% Fall in Rates

Instrument	Traded value Value (1)	Settlement Value	Settlement Amount	Present Value (2)
1. FRA	98,165,251.11	98,403,437.92	238,186.80	215,545.25
2. BAB	(89,330,378.51)	(89,547,128.51)	(216,749.99)	(216,749.99)
		Difference		(158.75)

Notes: (1) A long position is shown as a positive, a short position is negative
 (2) The futures present value is the mark-to-market, where a gain is positive and a loss is negative.

Transaction Cashflows and Net Benefit of Futures

Date/Cashflows	Interest Rates Up 1% pa Futures FRA	Interest Rates Down 1% pa Futures FRA
Trade date 13-Jun-95		
Traded Contract value	(89,330,379)	(89,330,379)
Value after rate change	(89,114,675)	(89,547,129)
Mark-to-Market	215,704	(216,750)
Expiry date 13-Sep-96		
Settlement	0	0
Interest on mark-to-market	24,258.32	(18,853.23)
Total Future value	239,962.32	(235,603.22)
Futures Benefit — \$	2,925.80	2,583.58
Futures Benefit — % pa	0.0120%	0.0105%
Futures Benefit — bp	1.20	1.05

This example displays the futures/FRA convexity effect. Given the simple scenario used, the convexity adjustment suggests that the futures yield should be a little more than 1 bp greater than the FRA rate.

In practice, the amount of the convexity adjustment is ignored for forward period of up to one year. For longer forward terms the adjustment is in the order of 1 or 2 basis points—gradually rising as the forward period increases.

4.3.3.3 A complete futures pricing model

If we incorporate all of the special features of a futures contract relative to a forward contract we can summarise the “complete” short-term interest rate futures pricing model as follows:

$$\text{Futures Price} = 100 - (\text{Forward Rate} + \text{Funding Adjustment} + \text{Convexity Adjustment})$$

4.4 Forward bonds

4.4.1 General description

Forward bonds are an OTC forward contract on fixed interest bearing bonds and can be likened to an FRA on a long-term interest rate security. Under a forward bond agreement the two parties agree to deliver a specified bond series at a fixed price at a future date. While FRAs relate to a generic money market interest rate such as LIBOR, forward bond agreements relate to a specific bond issue. So, every forward bond agreement must reflect the characteristics (such as issuer, maturity date, coupon and yield) of the underlying bond. This variety of issues combined with the more complex valuation formula, tend to make forward bonds a more complex and specialised transaction than FRAs.


Typically the forward bond market revolves around government and other high quality bond issues as there is already considerable activity in the underlying securities. As we will see, cash, forward and futures transactions in the bond market are closely related.

4.4.2 Synthetic replication


A forward bond purchase can be synthetically replicated in the same way as an investor FRA by purchasing the security today and financing the bond for the forward period. This is illustrated in *Exhibit 4.20*.

Exhibit 4.20
Synthetic Replication of a Forward Bond Purchase

Trade date



Suppose you are a fund manager and you wish to buy bonds but will not receive the funds to buy the bonds for three months. To replicate buying the bonds forward you can buy the bond today for \$100 million and fund the purchase by borrowing for three months at a rate of 5.5% pa.
In three months



After three months the funds are received from fund members, these proceeds are used to repay the money market loan and interest. The cost of buying the bonds forward will reflect the borrowing cost and any cashflows received from the bond. Let us assume a coupon of \$5 million is received at the end of the three months.

Exhibit 4.20—continued

<i>Forward Rate calculations</i>	=		
Original cash cost	=	100,000,000.00	
Bond value at three months	=	95,000,000.00	The bond value is reduced by the coupon payment
Borrowing interest at three months	=	$100m \times (1 + 0.0575 \times 90 / 360)$	
	=	1,437,500.00	
Asset cashflows	=	5,000,000.00	
Net Cash value at three months	=	96,437,500.00	
Forward Price	=		96.4375

As with all forward transactions the forward price represents the current cash price adjusted for the cost of carry. In this case the cost of carry is the difference between the coupons received on the bond and the cost of financing the bond:

$$\text{Cost of Carry} = \text{Financing Cost} - \text{Bond Coupons}$$

A common mistake in forward bond calculations is to use the yield to maturity as the asset return instead of the bond coupon. A forward calculation is concerned with actual cashflows that take place during the forward term—a yield to maturity reflects the asset return over the whole life of the underlying security.

While the example uses a money market interest rate to determine the financing cost this provides only an estimate. If a bond is a government bond then the credit quality of that instrument is likely to be better than most money market instruments. In a synthetic forward purchase the buyer could offer the bond as security and borrow at an interest rate appropriate to the credit quality of the bond. This will be discussed more in the section below on repurchase agreements.

4.4.3 A model for forward bond prices

The synthetic replication indicates that the forward bond price conforms with the “lumpy” income model developed in section 3.6.3, where the lumpy income is the coupon payment on the bond. If the bond does not pay any coupon during the forward period then we use the “no income” model from section 3.6.1.

Forward Bond Price Model—One Coupon Payment

The forward price per \$100 can be expressed as:

$$F = S \times (1 + r_1 \times f_1 / D) - c \times (1 + r_2 \times f_2 / D)$$

Where

- F = Forward price per \$100 face value including accrued interest (“dirty price”)
- S = Cash bond price including accrued interest
- r_1 = Interest rate to the forward expiry date
- r_2 = Interest rate between the coupon payment and forward expiry dates
- D = Day count basis (365 or 360)
- f_1 = Number of days to the forward expiry date
- f_2 = Number of days between the coupon payment and forward expiry dates
- c = Periodic coupon payment per \$100 of face value

It is interesting to note that the forward calculation is based purely on cashflows between today and the forward settlement date. Apart from calculating the initial cash price, S, there is no reference to the bond pricing formula. As with all forward calculations the aim of the model is reflect the cashflow consequence of entraining into a forward transaction.

This formula solves for the forward price. To determine the forward yield, enter the forward price into the bond price calculator and solve for the yield on the forward settlement date. This yield will reflect the cost of carry,

however, it is not just a function of the difference between the financing cost and coupon rate, it also reflects the timing and payment of coupons.

This model only allows the incorporation of one coupon payment. Including other coupon payments is simply a matter calculating the future value of each extra coupon using the same methodology as the first coupon. In practice the bulk of forward bond transactions have a forward term of three months or less, so encountering more than one coupon is uncommon.

The best way of building forward bond pricing models is to combine them with a cash bond price calculator. This allows you to automatically generate the current cash price, as well as the next coupon dates and coupon amounts. An example of a forward bond pricing spreadsheet is provided in *Exhibit 4.21*. This spreadsheet assumes that the two short-term interest rates, r_1 and r_2 , are the same.

Exhibit 4.21
Forward Bond Price and Yield Calculator

<i>Spreadsheet Example</i>		
Field	Cell	Cell Address: Formula (blank for input cells)
Inputs		
Trade date	20-Dec-95	=\$F\$11:
Forward settlement date	15-Jun-96	=\$F\$12:
Maturity date	15-Jul-99	=\$F\$13:
Coupon rate %	8.0000	=\$F\$14:
Number of periods/year (1, 2 or 4)	2	=\$F\$15:
Current yield to maturity	6.0000	=\$F\$16:
Repo rate till forward settlement	4.8500	=\$F\$17:
Repo rate day count basis (360 or 365)	365	=\$F\$18:
30/360 days count (y or n)	n	=\$F\$19:
Underlying Bond Details		
Settlement date	20-Dec-95	=\$O\$22: =F11
Forward date	15-Jun-96	=\$O\$23: =F12
Maturity date	15-Jul-99	=\$O\$24: =F13
Last coupon date	15-Jul-95	=\$O\$25: =COUPPCD(F22,F24,F28,F30)
Next coupon date 1	15-Jan-96	=\$O\$26: =COUPNCD(F22,F24,F28,F30)
Coupon rate %	8.0000	=\$O\$27: =F14
Number of periods/year (1, 2 or 4)	2.0000	=\$O\$29: =F15
Current yield to maturity % pa	6.0000	=\$O\$29: =F16
MS excel day count method	1	=\$O\$30: =IF(F19="n",1,0)
Clean price	106.3361	=\$O\$31: =PRICE(F22,F24,F27/100,F29/100,100,F28,F30)
Accrued interest at trade rate	3.4348	=\$O\$32: =IF(F22=F25,0,ACCRINT(F25,F26,F22,F27,1,F28,F30))

Exhibit 4.21 — continued

Field	Cell	Cell Address: Formula (blank for input cells)
Financing (or Repo) Details		
Current financing rate	4.85	\$Q\$35 : =F17
Repo rate day count	365	\$Q\$36 : =F18
Dirty bond price on trade date	109.7708	\$Q\$37 : =F32+F31
Number of coupons during repo	1	\$Q\$38 : =IF(F23>F26,1,0)
Number of repo days in forward period	178	\$Q\$39 : =F23-F22
Number of days from coupon date to fwd date	152	\$Q\$40 : =IF(F38=1,F23-F26,0)
Repo finance cost of bond	112.3671	\$Q\$41 : =F37*(1 + F35/(F36*100))*F39)
Forward Price Calculation		
Cumulative coupon 1 value at forward date	4.0808	\$Q\$44 : =F27/F28*(1 + F35/(100*F36))*F40)*F38
Dirty forward price	108.2864	\$Q\$45 : =F41-F44
Last coupon date at forward date	15-Jan-96	\$Q\$46 : =COUPPCD(F23,F24,F28,F30)
Next coupon date at forward date	15-Jul-96	\$Q\$47 : =COUPNCD(F23,F24,F28,F30)
Accrued interest at forward date	3.3407	\$Q\$48 : =ACCRINT(F46,F47,F23,F27,1,F28,F30)
Clean forward price	104.9457	\$Q\$49 : =F45-F48
Forward yield % pa	6.2093	\$Q\$50 : =YIELD(F12,F13,F14/100,F49,100,F15,F30)*100

To appreciate the impact of a different coupon, *Exhibit 4.22* takes the bond in the previous example and applies different coupon levels. While the yield to maturity and financing rate on each bond is the same, the forward yield at each coupon level is different. In general, as the coupon rate increases the absolute level of the cost of carry increases.

Exhibit 4.22		
The Impact of Coupons on Forward Bond Yields		
<i>Using the example in Exhibit 4.21, determine the impact on the forward yield of different coupon levels leaving all other inputs unchanged.</i>		
<i>Original bond</i>		
Trade date		20-Dec-95
Forward settlement date		15-Jun-96
Maturity date		15-Jul-99
Coupon rate %		8
Number of periods/year (1, 2 or 4)		2
Cash yield to maturity % pa		6
Financing/repo rate % pa		4.85
Forward price		104.9457
Forward yield — % pa		6.209315
Yield cost of carry		-0.20931
<i>Impact of Coupons</i>		
Coupon % pa	Forward Yield % pa	Cost of Carry % pa
4	6.1973	(0.1973)
8	6.2093	(0.2093)
12	6.2199	(0.2199)
16	6.2292	(0.2292)

4.4.4 A model for forward bond valuation

The forward value of a forward bond is the difference between the contract price in the forward bond agreement and the prevailing forward bond price:

$$\text{Forward Value} = \text{Forward Bond Price} - \text{Contract Price}$$

When determining the present value of the forward bond we need to be wary of the discounting interest rate used. While it is common practice to use the prevailing short-term money market rate to the forward term to determine the present value, it is not strictly correct. Determining the present value

involves converting a known future cashflow with specific characteristics into a known amount today. A forward bond forward value is obviously characterised by the difference between the current forward price and the contract price. However, these cashflows are also dependent on the characteristics of the underlying bond—most notably the credit quality of the bond. Consequently, the interest rate used to present value these cashflows should be a short-term interest rate on the bond. So, in simplistic terms the interest rate used for government bonds should be equivalent to a treasury bill rate, while the interest rate for bank bonds should be the same as bank-related money market instruments.

In summary, the present value interest rate should be the same as the financing interest rate, r_1 , used in the forward bond price formula. So, on a simple interest basis the formula is as follows:

$$\text{Present Value} = \text{Forward Value} / (1 + r_1 \times f / D)$$

In *Exhibit 4.23* an existing forward bond position is marked-to-market by calculating the current forward price and then determining the forward and present value of the forward bond. The forward value can be determined with reference to the “dirty” or “clean” price of the bond. Either method is acceptable as the difference in the two is simply accrued interest. In this example we use the clean price, which gives just the capital gain or loss on the position.

Exhibit 4.23

Forward Bond Valuation

Suppose you have purchased forward \$100 m face value of bonds on the following basis:

Original forward purchase

Trade date	17-Jun-96
Forward settlement date	02-Nov-96
Maturity date	15-Dec-05
Coupon rate %	6.5000
Number of periods/year (1, 2 or 4)	2
Current yield to maturity	7.2500
Repo rate till forward settlement	8.0000
Repo rate day count basis (360 or 365)	360
30/360 days count (y or n)	n
Dirty forward price	97.8629
Clean forward price	95.3765
Forward yield % pa	7.1988

Exhibit 4.23—continued

On 20 September interest rates have risen and you decide to mark the position to market (that is, calculate its present value). The forward price is now as follows:

Forward price on 20 September

Valuation date	20-Sep-96
Forward settlement date	02-Nov-96
Maturity date	15-Dec-05
Coupon rate %	6.5000
Number of periods/year (1, 2 or 4)	2
Current yield to maturity	7.7000
Repo rate till forward settlement	8.4000
Repo rate day count basis (360 or 365)	360
30/360 days count (y or n)	n
Dirty forward price	94.8213
Clean forward price	92.3350
Forward yield % pa	7.6829

In general the forward value is calculated excluding the effects of accrued interest, that is using the "clean" price.

$$\begin{aligned} \text{Forward value} &= \frac{(\text{forward price} - \text{contract price}) / 100 \times \text{face value}}{(92.335 - 95.3765) / 100 \times 100,000,000} \\ &= \frac{(3,041,500)}{(3,041,500)} \end{aligned}$$

$$\begin{aligned} \text{Present value} &= \frac{(\text{forward value} / 1 + \text{repo rate} \times f / D)}{-3,041,500 / (1 + .084 \times 43 / 360)} \\ &= \frac{(3,011,287)}{(3,011,287)} \end{aligned}$$

4.4.5 Repurchase agreements

In any forward bond market, a large proportion of forward transactions are linked to repurchase (repo) and reverse repurchase (reverse repo) agreements. In some markets it is estimated that the majority of bond dealer transactions is some form of repo or reverse repo. The bulk of bond repos are very short-term—in the order of one day to one week. Even in very liquid markets such as the US treasury bond market, repos for longer than six months are relatively rare.

A repo is the simultaneous execution, with one counterparty, of the sale of a cash bond and a forward bond purchase. That is, it is an agreement to sell a bond today and repurchase it at a date in the future at a fixed price.¹² A reverse repo is the simultaneous purchase of a cash bond and forward bond sale. A repo is arranged so that the sale of the bond today is at the prevailing cash price for the bond and the future repurchase of the bond is based on the forward bond price. Given that we know the cash price of the bond, the forward period and the coupon on the bond the only unknown in a repo is the financing interest rate. This makes the financing interest rate the key variable in any repurchase agreement and explains why this financing interest rate in the forward pricing formula is referred to as the “repo rate”.¹³

Exhibit 4.24 diagrammatically illustrates the mechanics of a repo.


12. Other names for repo transactions include buy-backs, reciprocal purchase agreements and bond lending. All have similar economic results, however, the mechanics can be different. For more detail see T Shanahan, “The Repo Market” (1991) (Summer) *Journal of International Securities Markets*.
13. In some markets this is referred to as the “cash” or “term” rate. Both of these terms can be confused with other interest rates, so they will be avoided in the remainder of this book.

Exhibit 4.24

Repurchase and Reverse Repurchase Agreements

A bond dealer decides to execute a Repurchase Agreement on a \$100 million bond holding with a bank. The current cash price of the bond is 110 per \$100 face value. The first leg of the transaction consists of the bond dealer selling the \$100 million face value of bonds and the bank paying \$110 million in cash for the bond. At the same time the bond dealer agrees to buy the bonds back from the bank in 1 month's time. Note that the bank is executing a reverse repo.

Trade Date — Cash/Spot Leg



In 1 Month - Forward Leg

At the end of 1 month the dealer buys the bond back. The price that they pay will be equivalent to the forward price — in this case a price of 111.7. The effect of this transaction is to allow the bond dealer to borrow money but retain a forward ownership of the bonds. The bank has invested cash with the dealer and holds the bonds as a form of security.

As we discussed in Section 3,¹⁴ the effect of a repo is to shift cashflows from one point of time to another—it does not actually change the participant's exposure to changes in the value of bonds. In *Exhibit 4.24*, the repo has provided a very efficient method of funding its bond positions, as it retains the same long exposure to bond price movements. However, it is able to finance this bond holding at a rate which is usually lower than its normal funding rate. Notice that by continually using repos, that is executing a new repo as soon as each repo matures, the dealer could fund its bond holding over long periods at an attractive rate. This explains why the judicious use of repos is a fundamental aspect of dealing in bonds.

Some of the reasons behind entering into repos and reverse repos are summarised as follows.

4.4.5.1 Repos

1. Obtaining funding at an attractive rate.
2. "Grossing up" investment returns—selling bonds already owned into a repo and then using the cash to invest in more securities.
3. Both repos and reverse repos offer a method of liquidity management to central banks. Repos are widely used to manage the cash position of the banking system. A repo allows the central bank to withdraw cash from the economy today and then re-inject that cash at a date in the future when it will be required.

4.4.5.2 Reverse repos

1. Investors wishing to invest cash and obtaining the underlying bond as security.
2. Often banks are required to hold government bonds for regulatory reasons. Under a reverse repo, the bank could obtain ownership of the bond. However, it is not exposed to the potential price volatility of the underlying bond.
3. An organisation which has created a short position in bonds (for example from the maturity of a forward sale or from an option transaction) could cover this position temporarily through a reverse repo.

A fascinating aspect of repos is that for every term to maturity there are multiple repo rates. Every bond series issued has its own characteristics such as the issuer, the maturity date, the coupon payment dates, and the coupon amounts. As we saw in *Exhibit 4.22*, just by varying the size of the coupon changes the forward yield to maturity on otherwise identical bonds. Given these different characteristics, the market can have greater or less interest in owning bonds. If a particular bond series is in demand, then the cost of financing it will be lower than less in demand bonds. This is because reverse repo counterparties will be willing to invest cash at a lower rate to obtain temporary ownership of highly favoured bonds.

The US treasury bond market is the most liquid cash and forward bond market in the world. While the issuer is constant and the terms similar the repo rate on different bond issues can vary substantially depending on the level of demand for specific issues. As an example, in late October 1995

14. See section 3.3.

while the overnight US\$ interest rate was 5.75% pa, the repo rate for treasury bonds varies from 5.50% pa down to 1.75% pa for the series, which are in heavy demand. Reasons for these very low repo rates reflects the existence of large short positions in these securities relative to their supply. These short positions are covered with reverse repos, and as the availability of bonds declines the short-position holders are willing to accept a lower return on the cash invested in a reverse repo.

We can solve for the repo rate implied in a forward bond transaction by re-arranging the formula from section 4.4.3 as follows (if we assume r_1 and r_2 are the same):

Calculating the Implied Repo Rate

The repo rate in the forward leg of a repo can be solved as follows:

$$r_1 = \frac{F - S + c}{S \times f_1 / D - c \times f_2 / D}$$

Where

- F = Forward price per \$100 face value including accrued interest at futures date (dirty price)
- S = Cash bond price including accrued interest
- r_1 = Repo rate to the forward expiry date
- D = Day count basis (365 or 360)
- f_1 = Number of days to the forward expiry date
- f_2 = Number of days between the coupon payment and forward expiry dates
- c = Periodic coupon payment per \$100 of face value

Exhibit 4.25 calculates the implied repo rate on a forward bond transaction using this formula.

Exhibit 4.25
Repo Rate Calculation

Calculate the implied repurchase rate in the following forward bond transaction.

Cash Bond Details

Trade date	20-Dec-95
Forward settlement date	15-Jun-96
Maturity date	15-Jul-99
Coupon rate %	8.0000
Number of periods/year (1, 2 or 4)	2
Current yield to maturity	6.0000
Dirty cash price	109.7708
Day count basis	Act/Act

Forward Details

Dirty forward price	108.2864
Repo rate day count basis (360 or 365)	365
Number of days in forward period	178
Coupon payment on 15-Jan-95	4.0000
Number of days from coupon date to fwd date	152

Repo calculation

Inputs:

S =	109.7708	$f_1 =$	178
F =	108.2864	$f_2 =$	152
c =	4.0000	D =	365

Formula:

$$r = \frac{F - S + c}{(S \times f_1 / D - c \times f_2 / D)}$$

$$= 4.8500\%$$

4.5 Bond futures

4.5.1 General description

Bond futures represent a standardised, exchange-traded forward bond contract. Like short-term interest rate futures contracts they have become an integral part of most financial markets, and typically represent a benchmark for long-term interest rate transactions.

The pricing and valuation of these instruments is derived from the forward bond calculations in section 4.4 above. As with all futures contracts, adjustments may need to be made for margin funding costs. The users of both bond futures and forward bonds are usually very similar and the reasons these products are used are closely related. In those countries where the underlying cash market is liquid, the volume in bond futures is substantial.

A list of the major bond futures contracts are listed in *Exhibit 4.26* along with the total volume for 1994.


Exhibit 4.26
List of Bond Futures Contracts and Volumes

Contract	Currency	Quote Method	Delivery Method	Exchange(s)	1994 Futures Volumes	
					No of contracts	Face Value (Bn)
A\$ 3 year bond	A\$	Yield	Cash settled	SFE	9,709,791	719
A\$ 10 year bond	A\$	Yield	Cash settled	SFE	800,263	59
Medium term notional bond (BOBL)	DEM	PPH	Physical delivery	DTB	5,647,859	402
German government bond (Bund)	DEM	PPH	Physical delivery	LIFFE	37,335,437	6,642
German government bond (Bund)	DEM	PPH	Physical delivery	DTB	14,160,460	2,519
Spanish government bond	ESP	PPH	Physical delivery	MEFF	13,191,895	548
10 year government French bond	FFR	PPH	Physical delivery	MATIF	50,153,150	5,171
Long gilt future	GBP	PPH	Physical delivery	LIFFE	19,048,097	1,481
Italian government bond (BTP)	ITL	PPH	Physical delivery	LIFFE	11,823,741	1,492
Japanese government bond (JGB)	JPY	PPH	Physical delivery	LIFFE	610,925	599
10 year Japanese government bond (JGB)	JPY	PPH	Cash settled	TSE	12,999,698	12,754
10 year Japanese government bond (JGB)	JPY	PPH	Cash settled	SIMEX	443,564	218
NZ\$ 3 year bond	JPY	PPH	Cash settled	NZFOE	101,229	7
NZ\$ 10 year bond	NZ\$	Yield	Cash settled	NZFOE	42,541	3
Swiss government bond	NZ\$	Yield	Cash settled	SOFFEX	949,657	14
2 Year treasury notes	SFR	PPH	Physical delivery	CBOT	939,043	188
5 Year treasury notes	US\$	PPH	Physical delivery	CBOT	12,462,838	1,246
10 Year treasury notes	US\$	PPH	Physical delivery	CBOT	24,077,828	2,408
US treasury bonds (T-bonds)	US\$	PPH	Physical delivery	CBOT	99,959,881	9,996
US treasury bonds	US\$	PPH	Physical delivery	MIDAM	1,385,904	69
					315,843,741	46,535

Notes: 1. "PPH" means the bond futures contracts are quoted in price terms as price per hundred units of face value.
2. "Yield" quotes mean the futures are quoted in terms of yield to maturity where the futures price is equivalent to 100 minus the yield.

Exhibit 4.26—continued

Composition of Bond Future Volume by Currency



Two important differences between these futures contracts is highlighted in the table in *Exhibit 4.26*:

1. *Quote method*: The price of most bond futures contracts is quoted as the current price per 100 units of face value (shown in the table as “PPH”). For these contracts, the futures price is essentially the same as the forward price previously calculated in section 4.4.3. The other alternative is the “yield” method. Futures prices are quoted as 100 minus the yield to maturity of the underlying forward bond. The futures quotation method is usually a reflection of the local bond market convention for quoting cash bond prices.
2. *Delivery method*: There are two alternative methods with which bond contracts are terminated: physical delivery and cash settlement. As its name implies, physical delivery requires that all open contracts at expiry must deliver (the futures contract seller), or take delivery of (the buyer), a defined amount and type of bonds. In the case of cash settlement, at expiry, all open contracts are reversed at the final settlement price of the contract (usually on the last day of trading). That is, all obligations under the futures contract are cancelled upon payment or receipt of the cash difference between the original traded price and the final settlement price of the futures contract.

In the case of some contracts, notably the A\$ and NZ\$ bond futures, they are quoted using the yield method and are cash settled against a basket of underlying bonds. This creates some additional pricing complexities for the forward bond formula, which are highlighted below.

4.5.2 Pricing and valuing bond futures

Conceptually, the pricing and valuation tools developed in section 4.4 can be applied directly to bond futures with the same sort of adjustments as were applied to short-term interest rate futures:

$$\text{Futures Price} = \text{Forward Price} + \text{Funding Adjustment} + \text{Convexity Adjustment}$$

As with all futures, there is no distinction between forward and present values and this should be incorporated into any hedging transaction using the PVBP in the same manner as the short-term futures contract.¹⁵

While the funding and convexity adjustments discussed in relation to short-term futures should strictly be applied, they are often ignored by market participants. The reason for this is twofold:

1. *Short forward period*: Most of the traded volume in bond futures across all contracts have a relatively short forward period (up to six months). As we have seen in earlier sections the funding and convexity adjustment calculations are usually extremely small for forward periods of under one year.
15. See section 4.3.3 for an example of dealing with the constant PVBP characteristics of futures contracts.

2. *Long-term instrument*: If the funding or convexity adjustment is calculated and spread over the fairly long life of the underlying bond, the impact of the adjustment tends to be fairly small.

Unfortunately, while it describes the conceptual relationship, deriving the final quoted futures price is not quite as simple as the formula above implies. In all bond futures additional adjustments are required depending on the nature of the delivery process. We can divide the futures price calculations into two general groups:¹⁶

1. *Delivery and conversion factors*—At the expiry of most bond futures contracts (for example, CBOT treasury notes and bonds) a physical delivery of bonds takes place. There is a specified list of approved bonds which can satisfy delivery and a “conversion factor” which is indented to convert each bond to the equivalent of the notional bond underlying the contract. This conversion process is not perfect and usually one of the approved bonds becomes the “cheapest” to deliver. Effectively, the price of the futures contract is based off the forward price per hundred (multiplied by the conversion factor) of the cheapest to deliver bond.
2. *Yield quotes and basket bonds*—These bonds are cash settled against the average yield on a basket of bonds on the last day of trading of the futures contract. The futures price is given by 100 minus the forward yield of the basket of bonds underlying the futures contract. Because each futures price point is one basis point in yield, and given bonds exhibit convexity, as the futures price changes so does the dollar, or “tick” value of each futures price point—the higher the futures price the higher the tick value. Both the SFE and NZFOE’s bond contracts trade on this basis.

5. FOREIGN EXCHANGE FORWARDS

5.1 Introduction

The interesting feature about foreign exchange transactions is that they are purely about cashflows. In interest rate markets there is typically an underlying security which has specific interest paying and maturity characteristics and it may be in limited supply. With a foreign exchange transaction there is no specific underlying instrument; any organisation can create its own tailor-made foreign exchange transaction.

Additionally, foreign exchange transactions are a homogenous product and completely fungible.¹⁷ That is, if a company enters into a foreign exchange transaction with a bank which settles in two days time, this can be offset by entering into an opposite transaction with another bank—the only difference will be the profit or loss due to differences in the exchange rate on the original and offsetting deals.

16. For more on this see Martin, op cit n 1, Ch 5.

17. If two financial instruments are fungible then they are perfect substitutes and can be used to replace one another.

Combining the underlying demand to execute foreign exchange deals for trade and capital transactions with this flexibility to create and manage foreign exchange positions, the enormous size and success of the OTC foreign exchange market can be understood. While currency futures exist, their volume is small relative to the OTC market.


The global foreign exchange market is the epitome of the “global financial village”—it is a huge marketplace spread across numerous financial centres and time zones. No matter what the time of day, it is possible to execute foreign exchange transactions involving major currencies. The BIS conducts a survey of global turnover every three years and the results from April 1995 are set out in *Exhibit 4.27*.

Exhibit 4.27
Global Foreign Exchange Turnover

Daily averages in billions of US Dollars (1)

Transaction	Apr-89 % share	Apr-92 % share	% change 1989-92	Apr-95 % share	% change 1992-95
Spot	350	400	14%	535	41%
Forward (2)	240	420	75%	695	65%
OTC Sub-Total	590	820	39%	1230	50%
Futures (3)	30	60	100%	72	20%
Total	620	880	42%	1,302	48%

Foreign Exchange Volume



Notes:

- (1) Figures have been adjusted for double counting.
- (2) Includes outright forwards and forwards which are one leg of an FX swap.
- (3) The Apr 95 volume was not provided with the BIS survey. Growth has been estimated from futures turnover growth of 20% from 1992 to 1995 (Source: FIA).

Source: BIS

In 1995 total daily volume for spot and forward FX instruments was US\$1.3 trillion. This represents an increase in volume of 48% from 1992, an increase on the 1989-1992 change of 42%. The proportion of transactions executed as forwards has steadily increased since the survey was commenced. In 1989 forwards and futures represented 44% of total volume, but by 1995 this has risen to 59%.

An important feature of the market is that the bulk of volume is made up of FX swap transactions—a simultaneous execution of spot and forward or forward and forward transactions. The BIS do not include hard data in the survey but they do note that only 14% of forwards are outright, the remaining 86% are part of an FX swap transaction. This implies FX swap volume of US\$ 598 billion per day in April 1995.

Reflecting the global nature of the FX market the survey also points to expanding turnover in most countries, as can be seen from *Exhibit 4.28*. Though, interestingly, the largest market, London, has grown at a more rapid rate since the survey started than its nearest rivals, the United States and Japan. In terms of geographical regions, the importance of Europe has grown in each survey, and it now represents more than half the total global volume.

Exhibit 4.28
Geographic Composition of Global FX Volume

Total OTC turnover by country

Transaction	Apr-89 % share	Apr-92 % share	% change 1989-92	Apr-95 % share	% change 1992-95
1 United Kingdom	184	290.5	58%	464.5	60%
2 United States	115.2	166.9	45%	244.4	46%
3 Japan	110.8	120.2	8%	161.3	34%
4 Singapore	55	73.6	34%	105.4	43%
5 Hong Kong	48.8	60.3	24%	90.2	50%
6 Switzerland	56	65.5	17%	86.5	32%
7 Germany		55	—	76.2	39%
8 France	23.2	33.3	44%	58	74%
9 Australia	28.9	29	0%	39.5	36%
10 Denmark	12.8	26.6	108%	30.5	15%
11 Canada	15	21.9	46%	29.8	36%
12 Belgium	10.4	15.7	51%	28.1	79%
13 Netherlands	12.9	19.6	52%	25.5	30%
14 Italy	10.3	15.5	50%	23.2	50%
15 Sweden	13	21.3	64%	19.9	-7%
Other	21.6	61.3	184%	89.2	46%
Total (1)	717.9	1076.2	50%	1572.2	46%

Note: (1) The total in this table is higher than the actual turnover in Exhibit 4.27 because this table has cross-border double counting. For example, a transaction executed by a bank in the UK and a bank in the USA would be recorded in both countries.

Exhibit 4.28—continued




Exhibit 4.29 demonstrates the declining importance of the US\$. In 1989 the US\$ was one of the legs in 90% of FX transactions but by 1995 this had fallen to 83%. Most of the US\$ share has been lost to non-DEM European Economic Union currencies.

Exhibit 4.29
Currency Composition of Global OTC FX Volume

Total OTC turnover by currency as a percentage

Currency	Apr-89 % share	Apr-92 % share	% change 1989-92	Apr-95 % share	% change 1992-95
1 US Dollar	531	672.4	27%	1020.9	52%
2 Deutschemark	159.3	328	106%	455.1	39%
3 Japanese Yen	159.3	188.6	18%	295.2	57%
4 Pound Sterling	88.5	114.8	30%	123	7%
5 French Franc	11.8	32.8	178%	98.4	200%
6 Swiss Franc	59	73.8	25%	86.1	17%
7 Canadian Dollar	5.9	24.6	—	36.9	50%
8 ECU	5.9	24.6	317%	24.6	0%
9 Australian Dollar	11.8	16.4	39%	36.9	125%
10 Other EMS currencies	17.7	73.8	317%	159.9	117%
11 Other	129.8	90.2	-31%	123	36%
Total (2)	590	820	200%	1230	200%


Note:

(1) The % share is expressed as a proportion of the global volume.

(2) As each FX deal involves two currencies the sum of the total volume in each currency is double the estimated global turnover.

Exhibit 4.29—continued

Turnover by currency — April 1995



5.2 A model for the forward foreign exchange price

In this section we will convert the generalised price formulae from Section 3 into a formula that generates a forward foreign exchange rate.

A forward FX transaction is an instrument in which a cash in one currency is exchanged for cash in another currency. Either currency can be thought of as the underlying asset. These assets provide constant and known incomes in the form of interest payments. We know that these assets can be exchanged on a spot basis at a rate of 1 unit of the base currency for S units of the terms currency.

The spot exchange rate is the current rate of exchange of two currencies. The synthetic replication above indicates that the forward exchange rate of two currencies is dependant on their respective interest rates. The forward FX pricing formula has a lot in common with the general forward models developed in Section 3, however, the formula needs to be adjusted for the fact that the interest amounts, r and q , are calculated on different currency amounts—1 unit of the base currency for every S units of the terms currency. We also need to incorporate the possibility that the calculation basis underlying each interest rate may differ. Generally for short-term forwards the main differences between money market interest rates are:

- the assumed number of days per annum (either 360 or 365); and
- whether the rate is based on a discount or yield calculation.¹⁸

The simplest approach is to look at the forward value of cash in each currency. As we have already discussed, cash by itself earns no income so we can apply the general “no-income” pricing model:

$$F = S \times (1 + r) \times f / D$$

This can be rewritten as follows for the forward cash value for the base (F_B) and the terms currency (F_T) as follows:

$$\text{Base Currency: } F_B = 1 \times (1 + r_B) \times f / D_B$$

$$\text{Terms Currency: } F_T = S \times (1 + r_T) \times f / D_T$$

where S is the exchange rate. The forward rate of exchange between these two currencies will be given by the ratio of these two forward amounts. This is summarised below:

18. See Martin, *op cit* n 1, Ch 3.

Short-term Forward Foreign Exchange Price

Using simple interest, the calculation is as follows:

$$F = \frac{S \times (1 + r_T) \times f / D_T}{(1 + r_B) \times f / D_B}$$

Where

F	=	Forward exchange rate
S	=	Spot exchange rate
r _T	=	Terms currency interest rate to forward date
r _B	=	Base currency interest rate to forward date
D _T	=	Terms currency day count basis (365 or 360)
D _B	=	Base currency day count basis (365 or 360)
f	=	Number of days to the forward expiry date from the spot settlement date.

An example of this calculation is set out in *Exhibit 4.30*.**Exhibit 4.30****Forward Pricing Example**

The current spot rate for US\$/CAD is 1.3513. Calculate the rate of a forward FX deal settling in 30 days from the spot date. The 1 month US\$ interest rate is 6.25% pa and the CAD rate is 8.2% pa. Calculate the implied forward FX rate.

$$S = 1.3513 \quad f = 30$$

$$r_T = 8.20\% \quad D_T = 365$$

$$r_B = 6.25\% \quad D_B = 360$$

$$\begin{aligned} F &= \frac{S \times (1 + r_T) \times f / D_T}{(1 + r_B) \times f / D_B} \\ &= \frac{1.3513 \times (1 + .082) \times 30 / 365}{(1 + .0625) \times 30 / 360} \\ &= \underline{\underline{1.3534}} \end{aligned}$$

$$\text{Forward Points} = \underline{\underline{0.0021 \text{ premium}}}$$

There are a number of assumptions underlying this calculation, which, if they do not hold, may require an adjustment to the formula:

1. *Simple interest*: There is assumed to be no compounding in the interest calculation. This can be adjusted for by applying the compound interest calculation.
2. *Zero-coupon*: The interest rates are assumed to be zero-coupon rates. This is generally a satisfactory assumption for forward FX deals of up to six

months, most interest rates longer than that contain re-investment risk. This assumption is relaxed in the section on long-term foreign exchange below.

3. *Interest rates are yields*: The formula assumes that interest rates are yields; that is, the interest amount is given by multiplying the principal value today. If an interest rate is from a discount security which calculates interest on a discount basis, then this will need to be converted to a yield basis.¹⁹
4. *Calculations from spot date*: The forward period is from the spot settlement date to the forward settlement date. Strictly speaking, the interest rates that should be used are the two day forward interest rates, however, this is usually ignored as the impact is minimal.

Bids and Offers

A forward FX rate is derived from three market rates: the spot FX rate and the interest rate in both currencies. In each case the appropriate bid and offer rate has to be identified. In FX markets the “bid” is the rate at which a dealer is willing to buy the base currency; and the “offer” is where the dealer will sell the base currency—the bid rate is lower than the offer rate. As its counterparty we do the opposite of the dealer so we will sell the base currency at the bid and buy it at the offer. As we noted earlier, from an end-user’s point of view a rule of thumb is that we will always lose money from the bid-offer spread, that is you “buy high and sell low”.

Unfortunately, in money markets bids and offers can have two meanings, depending on the underlying money market instrument. If the underlying instrument is a direct term deposit with a bank, then the bid will be where you can invest funds with the bank, and the offer is where you can borrow from the bank—the bid interest rate is lower than the offer interest rate. However, if the instrument is a tradeable security, such as a bank bill, where the underlying security is bought and sold according to its present value, the bids and offers are expressed in terms of interest rates. As a lower price implies a higher yield, the “bid” interest rate is higher than the “offer” interest rate. One way of avoiding confusion with these conventions is to apply the rule of thumb to interest rates: an end-user will invest at the lower interest rate and borrow at the higher interest rate.

Assuming that we are end-users rather than FX dealers, the simplest method of identifying the appropriate rates is to make use of the synthetic replication concept to identify whether to use the bid or offer. If we wish to buy the base currency forward in three months, the synthetic replication is to buy the base currency in the spot FX market, borrow in the terms currency and invest in the base currency. Using this, then, we use the appropriate rates as follows for a forward purchase or sale:

19. Ibid.

Calculating Forward Rates (End-User Perspective): Bids and Offers

Leg	Buy Base Currency Forward	Sell Base Currency Forward
Spot Foreign Exchange	offer	bid
Base Money Market ¹	bid (low rate)	offer
Terms Money Market ¹	offer (high rate)	bid

Note: The quote convention in this table assumes underlying instruments are bank deposits

An example of using the correct bids and offers is set out in the first part of Exhibit 4.31.

5.3 A model for forward FX valuation

The forward value of a forward FX transaction is the difference between the original contract price and the prevailing market forward price. A forward FX transaction is usually expressed in terms of a constant amount of the base currency, while the terms currency is left to vary and all gains and losses are generated in the commodity currency. The forward value in the terms currency of a forward FX where the base currency is purchased/terms currency sold, the deal can be expressed as follows:²⁰

$$\text{Forward Value}_{\text{TERMS}} = \text{Base Amount} \times F_M - \text{Base Amount} \times F_C$$

where F_M is the prevailing forward market rate and F_C is the contract rate.

If, instead, the terms currency is held constant then the situation is inverted as follows for a bought base currency position:

$$\text{Forward Value}_{\text{BASE}} = \text{Terms Amount} / F_C - \text{Terms Amount} / F_M$$

The mark-to-market on a forward FX position will be given by present valuing these forward value calculations. It is essential that we correctly identify the currency in which the forward value is generated. As it is usually the case that the interest rate differs between the two currencies, the present value of a forward cashflow in either currency will be different. So, if the forward value is calculated in the terms currency it should be present valued using the terms currency interest rate:

$$\text{Present Value}_{\text{TERMS}} = \text{Forward Value}_{\text{TERMS}} / (1 + r_T \times f / D)$$

otherwise if the forward value is in the base currency:

$$\text{Present Value}_{\text{BASE}} = \text{Forward Value}_{\text{BASE}} / (1 + r_B \times f / D)$$

Exhibit 4.31 provides a complete worked example for determining the current mark-to-market revaluation of a forward FX position.

20. A sold base currency/purchased terms currency has the same formula with the forward rates F_M and F_C reversed.

Exhibit 4.31
Forward FX Transaction

You currently hold a forward FX position where you buy GBP/sell USD 10 million at 1.5340. This contract will settle in 92 days time. Given the market rates below calculate the current revaluation of this position.

Market rates		
Spot GBP/USD:	bid	offer
	<u>1.5629</u>	1.5634
Three Month Money Market rates (Deposit rates)		
	Bid	Offer
GBP	6.55	<u>6.65</u>
USD	<u>5.53</u>	5.75

1. Calculate current forward FX price

We wish to revalue a bought GBP forward FX position. The current market value will be given by the forward FX price at which an offsetting position can be put in place. Consequently we need to generate a three month forward price to sell GBP/buy USD. The inputs to the forward price model will be the FX spot rate bid, the GBP money market offer and the USD money market bid.

$$\begin{aligned}
 S &= 1.5629 & f &= 90 \\
 r_T &= 5.53\% & DT &= 360 \\
 r_B &= 6.65\% & DB &= 365 \\
 F &= \frac{S \times (1 + r_T) \times f / DT}{(1 + r_B) \times f / DB} \\
 &= \frac{1.5629 \times (1 + .0553) \times 90 / 360}{(1 + .0665) \times 90 / 365} \\
 &= \mathbf{1.5589}
 \end{aligned}$$

2. Calculate the forward and present values

The principal value of the transaction is expressed in the terms currency, that is, USD 10 million so we use the formula:

$$\text{Forward Value} = \text{Terms Amount} / F_c - \text{Terms Amount} / F_m$$

Where

$$\text{Terms Amount} = 10,000,000 \text{ USD}$$

$$F_c = 1.5340$$

$$F_m = 1.5589$$

$$\text{Forward Value} = 10,000,000 / 1.5340 - 10,000,000 / 1.5589$$

$$= 104,125 \text{ USD}$$

$$= 66,794 \text{ GBP}$$

$$\text{Present Value} = \text{Forward Value} / (1 + r_T \times f / DT)$$

$$= 102,705 \text{ USD}$$

$$= 65,715 \text{ GBP}$$

The mark-to-market revaluation on this position is a profit of USD 102,705

5.4 Foreign exchange swaps

The BIS survey in section 5.1 identified that a very large percentage of global foreign exchange volume is in the form of FX swaps—the simultaneous execution of offsetting spot and forward foreign exchange transactions.²¹ These transactions need to be distinguished from cross-currency interest rate swaps (“currency swaps”). FX swaps are generally short-term in nature and consist of just a spot and forward leg. Currency swaps are longer-term instruments, typically in the range of 1 to 10 years, and involve a series of periodic interest exchanges. While they are different instruments, the overall economics of the two transactions have the same effect of switching one currency exposure to another for the life of the instrument.

The popularity of FX swaps is a reflection of the important tasks for which it can be used:

- the temporary conversion of one currency to another without creating an exposure to foreign exchange movements;
- extending existing forward FX positions;
- a vehicle for trading the pure interest differential of two currencies without an exposure to exchange rates; and
- arbitrage related activities such as covered interest arbitrage.

Like a bond repurchase agreement, an FX swap allows the two counterparties to temporarily exchange one asset for another—in this case one currency for another. By itself, the transaction does not create an outright

21. An FX swap can also be the simultaneous execution of offsetting forward transactions with different terms.

foreign exchange exposure, as any gains or losses on the initial spot currency exchange are offset by gains and losses on the final forward exchange.

A useful example of the impact of an FX swap is the extension of an existing forward position which expires on the current spot settlement date. The mechanics of an extension or rollover of an existing FX position consist of two legs:

1. *The spot leg*: a spot FX deal is executed which offsets the cashflows from the original forward deal. This will typically crystallise a gain or loss equivalent to the difference between the original forward rate and the current spot rate.²²
2. *The forward leg*: a forward FX deal is executed at the prevailing forward rate. The net cost of this transaction will be the spot and forward spreads and the forward points.

If these two legs are executed separately then there is a risk that the spot foreign exchange rate moves between executing the spot and forward legs. An FX swap reduces a rollover to one deal and removes the foreign exchange risk by simultaneously executing both legs. In an FX swap the transaction becomes insensitive to the spot exchange rate used, and the important variable in the transaction is the forward points—it is for this reason that forward points are also referred to as swap points or the swap rate.

Exhibit 4.32 demonstrates how a forward FX position expiring in two days can be extended for another three months with one FX swap transaction.

22. Some countries still allow “historic rate rollovers”. In these transactions the gain or loss is not crystallised on the rollover date and is carried till the forward transaction finally matures. This introduces another component to calculating the forward FX price—an interest adjustment for any gains or losses funded by the FX dealer. In many countries these transactions have been banned due to the possibility of concealing trading losses for long periods of time.

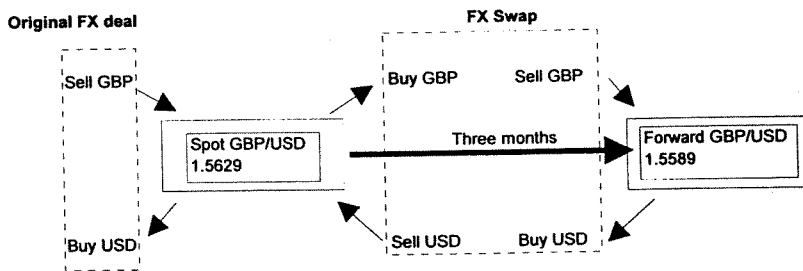
Exhibit 4.32

Extending a Forward FX Deal With an FX Swap

A company has a sold GBP 20 million/bought USD transaction expiring in two business days time. It wishes to extend this position for three months and decides to use an FX swap to do so where it buys GBP spot and sells GBP forward. The effect of the FX swap is to fund the settlement of the original transaction in two days time and create a new sold GBP/buy USD forward FX transaction in three months time. The net cost of this transaction is the forward or swap points over the next three months. Note that to simplify cashflows the original sold GBP position and the current spot FX rate are assumed to be same—this happy state of affairs rarely occurs in reality.

Current market rates:	Spot GBP/USD:	1.5629
	3 month Forward GBP/USD:	1.5589
	Forward or Swap Points:	(40)

Diagrammatic representation of FX swap



Cashflows of original deal and FX swap

Cashflows		GBP	USD
Day 0	Execute FX Swap		
Day 2	Settle Original FX deal	(20,000,000)	(31,258,000)
	Spot Leg of FX swap	<u>20,000,000</u>	<u>31,258,000</u>
	Net cashflows	—	—
Day 92	Forward leg of FX swap	(20,000,000)	31,178,000

Net Cost of FX swap =	(80,000)
------------------------------	-----------------

5.5 Long-term forward foreign exchange transactions

5.5.1 General description

As the name implies long-term forward foreign exchange (LTFX) transactions are a longer-term version of the forward FX transaction. It is an agreement between two parties who wish to agree on an exchange of currency cashflows at some date, possibly years in the future. For our purposes, we will consider an LTFX as any forward contract longer than six months.

LTFX contracts are relatively small proportion of total FX market volume. The BIS foreign exchange turnover statistics indicate that only 1% of FX volume is for longer than one year. While most FX dealers will quote LTFX transactions out for five years, given the lower liquidity of these instruments the bid-offer spread is wider than short-term forwards and reversing the position will not be as straightforward.

Typically, LTFX contracts are associated with hedging the FX exposures created by long-term borrowings or income streams created by assets in foreign currencies. Often LTFX transactions and currency swaps can be used interchangeably, the advantage of LTFX is that they can be more easily tailored to meet uneven future cashflows.

5.5.2 Pricing and valuing LTFX transactions

The synthetic replication of an LTFX is the same as a forward FX transaction—the borrowing and lending legs are just for a longer term. Similarly, the valuation procedure is identical. However, the valuation of LTFX is complicated by the following two effects:

1. *Zero-coupon yield*: The forward pricing and valuation models assume that there are not interest cashflows during the forward period—that is, the interest rates are zero coupon rates. This is a reasonable assumption when using money market interest rates, however, the quoted yields in most currencies which have a term to maturity of more than one year, usually are coupon-paying interest rates. The difficulty with coupon-paying yields is that there is a re-investment risk associated with each coupon payment. To price LTFX this risk has to be removed by deriving zero-coupon interest rates.
2. *Compounding*: Longer-term interest rates typically are expressed as compound interest rates—accordingly, compounding also needs to be incorporated into the model.

If we ignore either effect the LTFX price will be wrong, particularly if the yield curves in each currency have opposite shapes as this will exacerbate the difference between the zero-coupon and coupon interest differentials. *Exhibit 4.33* highlights the difference in forward price of a one year A\$/US\$ LTFX example when correctly calculated and when the short-term model is used.

Exhibit 4.33**LTFX Pricing: Using Zero Coupon Yields**

To demonstrate the impact of using zero-coupon versus coupon yield curves, calculate the LTFX using the two sets of interest rates provided below.

Current Market Rates

FX Rate	0.7202
1 year, US\$ coupon (sa)	5.50%
1 year, US\$ zero coupon (sa)	5.56%
1 year, A\$ coupon (ann)	7.60%
1 year, US\$ zero coupon (ann)	7.79%

Correct Forward Price

Using zero-coupon rate & compounding sa rate	0.7063	
<u>Incorrect Forward Prices (Short-term model)</u>		Difference
Using coupon rate ignoring compounding	0.7067	-0.0003
Using coupon rate & compounding sa rate	0.7072	-0.0008

The three components of a short-term forward FX model are a spot exchange rate and two money market instruments; whereas the LTFX model is comprised of a spot FX deal and two zero-coupon bond transactions. That is, a LTFX purchase of the base currency can be synthetically replicated by buying the base currency in the spot market, funding the spot settlement of the terms currency by borrowing using a zero-coupon bond and investing the base currency proceeds in a zero-coupon bond which expires on the forward settlement date.

To incorporate this into our forward pricing model, the interest rate legs will calculate the forward value of a single amount (compounded using zero coupon rates).²³ If we incorporate these concepts into the forward pricing model we can express the LTFX model as follows:

23. Zero coupon rates are often not observable as a market quote, so the rates will need to be generated from coupon paying, or par, yields such as prevailing swap rates. See Martin, op cit n 1, Ch 8 for more detail on calculating zero coupon interest rates from par rates.

Long-term Forward Foreign Exchange (LTFX) Price

Using simple interest, the calculation is as follows:

$$F = \frac{S \times (1 + r_T / m_T)^{n_T}}{(1 + r_B / m_B)^{n_B}}$$

Where

- F = Forward exchange rate
- S = Spot exchange rate
- r_T = Terms currency zero coupon interest rate to forward date
- r_B = Base currency zero coupon interest rate to forward date
- m_T = Terms currency payment frequency (that is, 1, 2, 4, 12)
- m_B = Base currency payment frequency
- n_T = Terms currency # of payment periods to the forward date
- n_B = Base currency # of payment periods to the forward date

An important characteristic of LTFX contracts is the impact of the interest differential on the forward price. Compared to a short-term forward FX contract, the interest rate legs of LTFX have considerably more impact on the forward price.

Exhibit 4.34

LTFX Pricing and Sensitivities

The graph below shows the sensitivity of a 5 year JPY/USD LTFX deal to changes in both the interest differential and the spot exchange rate. A feature of LTFX transactions is the increasing importance of the interest differential the longer the term to expiry. In this 5 year deal the impact of a move in the exchange rate of 1% is approximately equal to a change in the interest differential of 0.20% pa

Market Data	
Spot FX rate	101.00
USD 5 Year rate % pa (sa)	6.20
JPY 5 Year rate % pa (sa)	2.50
Interest Differential % pa	3.70

$$\begin{aligned}
 \text{LTFX Price} &= \frac{S \times (1 + rT / mT)^{nT}}{(1 + rB / mB)^{nB}} \\
 &= \frac{101 \times (1 + 0.025/2)^{10}}{(1 + 0.062/2)^{10}} \\
 &= \mathbf{84.27}
 \end{aligned}$$

The sensitivities of this position in foreign exchange points are as follows

$$\text{PVBP} = -0.0368$$

That is, a 1 bp rise in the interest differential will decrease the present value of the position by 0.0368 fx points.

$$\text{PVFP} = 0.0074$$

That is, a 0.01 change in the spot FX rate will alter the present value by 0.0074

$$\text{PVD} = 0.0075$$

Each day that passes increases the present value by 0.0075 fx points.

LTFX Price Sensitivity




Exhibit 4.34 calculates the LTFX price of a five year JPY/US\$ transaction and then graphs the sensitivity of the forward price to movements in the interest differential and the spot price. In this example a movement of 1% in the exchange rate has the same impact as a 0.20% pa change in the interest differential.

The forward value of a LTFX contract is the same as for a short-term forward FX contract: the difference between the contract value and the current market value. However, when calculating the present value we need to take account of the same yield considerations as the LTFX price. As a result we calculate the present value of the LTFX contract by using a zero coupon yield and applying the following compound interest, present value formula:²⁴


$$\text{Present Value}_{\text{BASE}} = \text{Forward Value}_{\text{BASE}} / (1 + r_B / m_B)^n_B$$

5.6 Par forwards

5.6.1 Instrument description

Another form of LTFX is the par forward. It is a series of LTFX contracts with regularly spaced settlement dates (for example, monthly or quarterly) at a constant exchange rate. *Exhibit 4.35* compares the exchange rate of a five year CHF/US\$ par forward transaction with a series of LTFX contracts which would achieve the same effect. The benefit from the point of view of an end user is that it allows the benefit or cost of a forward discount or premium to be spread out over the life of the transaction. In the example provided, from the point of view of a buyer of US\$, in the first two years of the transaction you buy US\$ at a substantially lower exchange rate than that implied by traditional forward instruments. The downside of course is that after two years the exchange rate is higher.

24. This is the base currency present value, to generate the terms currency PV just substitute terms currency variables for the base currency variables used.



In terms of the present value of these transactions the economics of a par forward and a series of LTFX are the same (as we will see below the pricing of par forwards ensures that this is the case). While in terms of the FX transaction alone there seems to be little added value in a par forward, the attraction of these instruments is that they can be very useful for cashflow management and also tax planning.

For example, suppose a Swiss-based distribution company is about to commence importing equipment from the United States. It has signed a five year contract which will require you to buy US\$10 million of equipment every quarter. The initial set-up costs associated with selling this equipment will be substantial and at current exchange rates you are likely to have negative CHF cashflows for the first two years, after which time cashflows will turn positive. You talk to your banker and he or she is concerned about financing the new project given its long lead time and currency exposures and suggests that you consider covering the exposures using the forward FX rates outlined in *Exhibit 4.35*. From the Swiss manufacturer's point of view a par forward would be more attractive than a series of LTFX contracts. Par forwards provide and immediately lower CHF cost for the equipment imports and possibly create a positive cashflow in the first two years. The downside is that the CHF cost of the equipment is relatively higher than a series of LTFX contracts. The par forward has allowed the Swiss company to obtain forward FX cover and also smooth out its cashflows.

A par forward is also interesting as it also shares some similarities with fixed-to-fixed cross currency interest rate swaps as it involves a constant exchange of currency cashflows as opposed to the variable cashflows of LTFX contracts. We will see that in our discussion of currency swaps that par forwards can be a useful tool in managing swap exposures.

5.6.2 Pricing and valuing par forwards

A par forward is a “smoothed” series of LTFX contracts. Pricing a par forward involves determining the cashflows from a similar series of LTFX transactions and then making an adjustment for the funding cost or benefit of evenly spreading out the currency cashflows. As is the case for any forward transaction the present value of executing a par forward should be zero.

The spreadsheet set out in *Exhibit 4.36* shows how to solve for the par forward rate for the CHF/US\$ example from the previous section. The first step is to calculate the LTFX rates for each periodic par forward date. We then need to determine what constant CHF delivery amount has an equivalent net present value to the CHF delivery amounts from the series of LTFX amounts. As can be seen the CHF par forward amounts will save the end user substantial CHF amounts—in effect the FX dealer is lending the difference between the LTFX rate and the par forward rate and having it repaid in the later delivery amounts. The dealer deserves a return on the money it has loaned to the end-user and so the par forward rate will include a funding cost. The solution of the par forward rate is to solve for a par forward rate which ensures that the total net present value of all of the funding differences is equal to zero. It is at this point that we know the money lent in the early par forward deliveries is re-couped, with interest, in the later deliveries.²⁵

25. In practice FX dealers will want to ensure that the interest rates used reflect market bids and offers and that they earn a positive funding NPV on the transaction. This can be achieved in this model by setting the target NPV at a rate which provides the dealers required return. For example, in *Exhibit 4.36* if the required NPV was CHF100,000 then the par forward rate would be 1.0476.

Exhibit 4.36
Pricing a Par Forward Transaction

Transaction Details

Spot FX Rate	1.13
Term (Yrs)	5
Delivery Frequency	Quarterly
Quarterly Amount	10,000,000

Solution

Unrounded Par Forward	1.04701178
Change	

A	B	C	D	E	F	G	H	I
Market Parameters		Forward FX rate		LTFX Cashflows		Par Forward CHF Amount	Net CHF Funding Difference	Net CHF Funding NPV
USD Rate	CHF Rate	USD Amount	CHF Amount	USD Amount	CHF Amount			
5.7500	2.0000	10,000,000	11,195,564	10,000,000	11,195,564	10,470,118	-725,446	721,837
5.7500	2.0000	10,000,000	11,092,093	10,000,000	11,092,093	10,470,118	-621,975	(615,802)
5.7817	2.0628	10,000,000	10,992,155	10,000,000	10,992,155	10,470,118	-522,037	(514,043)
5.8134	2.1257	10,000,000	10,894,820	10,000,000	10,894,820	10,470,118	-424,702	(415,793)
5.8358	2.2508	10,000,000	10,809,644	10,000,000	10,809,644	10,470,118	-339,526	(330,133)
5.8582	2.3759	10,000,000	10,730,621	10,000,000	10,730,621	10,470,118	-260,504	(251,409)
6.6166	2.8575	10,000,000	10,588,703	10,000,000	10,588,703	10,470,118	-118,585	(112,821)
7.3751	3.3391	10,000,000	10,434,820	10,000,000	10,434,820	10,470,118	35,297	33,026
7.2463	3.4347	10,000,000	10,383,014	10,000,000	10,383,014	10,470,118	87,104	80,653
7.1175	3.5304	10,000,000	10,342,908	10,000,000	10,342,908	10,470,118	127,210	116,508
7.0405	3.6254	10,000,000	10,298,796	10,000,000	10,298,796	10,470,118	170,322	154,230
6.9634	3.7205	10,000,000	10,265,579	10,000,000	10,265,579	10,470,118	204,539	183,032
6.9159	3.8062	10,000,000	10,227,428	10,000,000	10,227,428	10,470,118	242,690	214,577
6.8685	3.8920	10,000,000	10,196,121	10,000,000	10,196,121	10,470,118	273,997	239,261

Exhibit 4.36—continued

A	B	C	D	E	F	G	H	I	O
	Market Parameters		Forward FX rate	LTFX Cashflows		Par Forward CHF Amount	Net CHF Funding Difference	Net CHF Funding NPV	
	USD Rate	CHF Rate		USD Amount	CHF Amount				
15	6.8380	3.9668	1.0161	10,000,000	10,161,093	10,470,118	309,025	266,506	
16	6.8076	4.0416	1.0131	10,000,000	10,131,450	10,470,118	338,668	288,347	
17	6.7881	4.1057	1.0098	10,000,000	10,097,918	10,470,118	372,200	312,883	
18	6.7686	4.1697	1.0069	10,000,000	10,068,646	10,470,118	401,472	333,110	
19	6.7534	4.2142	1.0032	10,000,000	10,032,394	10,470,118	437,724	358,691	
20	6.7382	4.2587	0.9999	10,000,000	9,999,221	10,470,118	470,897	381,012	
							Net CHF NPV (Target)		0

Summary of Results

Average LTFX rate =	1.0477
Funding Cost =	0.0023
Par Forward rate =	1.0470

How this Spreadsheet works

1. Generate the Zero Coupon interest rates (on a quarterly basis for columns B and C).
2. Calculate the LTFX rates in Column C.
3. Calculate the USD and CHF cashflows for each quarterly roll for column E and F.
4. Enter a "guess" of the Par Forward Rate and enter into the cell labelled "Unrounded Par Forward".
5. Calculate the Par Forward CHF amount in column H by multiplying the USD amount by the Unrounded Par Forward Amount.
6. The Net CHF amount is simply the difference between columns F and G.
7. Calculate the NPV in column I by taking the present value of column H using the CHF zero interest rates and the compound interest present value formula.
8. From the Tools menu invoke the "Solver" function.
9. Make the Net NPV cell (column I) the target by changing the cell with the Unrounded Par Forward rate so that the target becomes Zero and then press solve, this will iteratively solve for the Unrounded Par Forward Rate which makes the NPV zero.

Once a par forward price can be generated then the valuation procedures are identical to an LTFX contract. However, while a series of LTFX positions can be valued individually, all of the par forward legs should be valued together, because of the interdependence between the series of deliveries.

6. EQUITY FORWARDS

6.1 Introduction

In their relatively short lifetime equity forwards have gained a reputation as a highly risky instrument. The October 1987 stock market crash, and 1989 “mini-crash”, prompted considerable conjecture that stock index futures exacerbated the market fall and, given the losses sustained by long position holders, were too risky an instrument to be used by the general public.²⁶ Then, in 1995, the collapse of British merchant bank Baring Brothers primarily due to unauthorised trading in share price index futures on the Japanese Nikkei index, prompted more regulatory “navel gazing” with respect to these instruments.

Despite the bad press, share price index futures have been an outstanding success if measured by volume growth since they were introduced in the United States in 1982. Stock index futures are a classic example of a derivative which “adds-value” to organisations and individuals with an exposure to share markets, as they provide a method of gaining an exposure to share market or hedging an existing exposure at considerably lower cost than transacting in the physical market. However, like their underlying market, the price volatility in share index futures is generally higher than most interest rate and currency markets—and as a result are a risky instrument in the hands of a novice user or uncontrolled trader.

An interesting feature of derivatives on equity index futures is the relatively high usage of options relative to both cash market volume and forward volume. Whereas in developed interest rate markets, option volume might be 10% of forward volume, in equity index markets the option percentage might be 20% or higher. This is partly explained by the high level of volatility, which encourages market participants to take insurance in the form of options.

Exhibit 4.37 summarises global exchange traded volume in equity forwards. Share price index futures dominate turnover, with very small volume in individual share futures. In the United States this is partly the result of regulatory restrictions, however, even where contracts have been listed on individual shares volume has been low.²⁷ The bulk of the volume is in US\$ denominated indexes, followed by Japan and then Germany.

26. Possibly the most spectacular case of the damage done by stock index futures to over-leveraged users was the Hang Seng Stock Index Futures contract traded on the Hong Kong Futures Exchange. The majority of long position holders at the time of the crash were individual speculators. After the crash, most of these “longs” defaulted on payment of their losses and the market was effectively closed.
27. Apart from Sweden, individual share futures contracts are relatively new, and this may also contribute to the low volume.

Exhibit 4.37
List of Share Index and Individual Share Futures Contracts

Contract	Currency	Contract Type	Delivery Method	Exchange	1994 Futures Volumes (1)(2)	
					No of contracts	Face Value (US\$ Bn)
Austrian Traded Index (ATX)	ATS	SFI	Cash settled	OTOB	348,291	3.455
All Ordinaries Index	AUD	SFI	Cash settled	SFE	2,552,546	102.708
Share Futures	AUD	ISF	Cash settled	SFE	32,441	0.002
Bel-20	BEF	SFI	Cash settled	BELFOX	154,574	0.795
San Paulo Stock Exchange Index (Ibovespa)	BRL	SFI	Cash settled	BM&F	10,583,594	92.218
Toronto 35 Index	CAD	SFI	Cash settled	TFE	104,209	9.543
Swiss Market Index (SMI)	CHF	SFI	Cash settled	SOFFEX	1,694,260	239.743
DAX Stock Index Future	DEM	SFI	Cash settled	DTB	5,140,803	804.121
Danish KFX Stock Index Future	DKK	SFI	Cash settled	CSE	429,466	8.162
IBEX 35	ESP	SFI	Cash settled	MEFF RV	27,020,886	39.362
CAC 40 Stock Index	FFR	SFI	Cash settled	MATIF	7,464,449	582.178
FTSE 100 Stock Index	GBP	SFI	Cash settled	LIFFE	4,227,490	595.662
Hang Seng Index	HKD	SFI	Cash settled	HKFE	4,192,574	257.324
Nikkei 225	JPY	SFI	Cash settled	OSE	6,208,821	1,109.523
Nikkei 300	JPY	SFI	Cash settled	OSE	4,684,480	124.362
Nikkei Index	JPY	SFI	Cash settled	SIMEX	5,801,098	518.331
Tokyo Stock Price Index (TOPIX)	JPY	SFI	Cash settled	TSE	2,623,067	371.105
Forty Index (3)	NZD	SFI	Cash settled	NZFOE	7,397	0.206
Swedish OMX Index	SEK	SFI	Cash settled	OM	1,706,984	44.518
Stock Futures	SEK	ISF	Physical Delivery	OM	234,250	0.002

Exhibit 4.37—continued

Contract	Currency	Contract Type	Delivery Method	Exchange	1994 Futures Volumes (1)(2)	
					No of contracts	Face Value (US\$ Bn)
Standard & Poors 400	USD	SFI	Cash settled	CME	285,962	79.640
Standard & Poors 500	USD	SFI	Cash settled	CME	18,708,599	5,612.299
Nikkei 225 (Settled in USD)	USD	SFI	Cash settled	CME	548,233	49.930
Russell 2000	USD	SFI	Cash settled	CME	36,239	—
Eurotop Index (Settled in USD)	USD	SFI	Cash settled	COMEX	62,231	—
Value Line Stock Index	USD	SFI	Cash settled	KCBT	50,259	14.123
Mini Value Line Stock Index	USD	SFI	Cash settled	KCBT	51,901	2.917
NYSE Composite	USD	SFI	Cash settled	NYFE	729,231	116.823
South African All Share Index	ZAR	SFI	Cash settled	SAFEX	2,185,672	—
Totals					107,870,007	10,779

Notes: 1. The Face Value has been estimated by multiplying the tick size by the /an estimated average value of the index and as such should be treated as a rough approximation only.


2. if the Face Value is a dash then there was insufficient information to make the estimate.

3. The NZFOE replaced the Forty Index with the NZSE-10 Contract in late 1995.

4. While the LEPO products traded on the ASX and SOFFEX are direct competitors of share futures they are strictly an option and have been included with equity options. This has little impact on any totals as LEPO volume is also very small.

Source: FIA

Exhibit 4.37 – continued
Composition of Equity Futures by Currency



Most statistics on derivatives show that equity derivative turnover is considerably smaller than for interest rate and currency derivatives. While this is in line with lower cash market volume in shares when compared with interest rate and foreign exchange markets, many market commentators suggest that the development of equity derivatives has lagged behind. As a result, it is widely perceived that equity derivatives will be the high growth derivatives areas in the remainder of the 1990s.²⁸

In the following sections we will review the pricing and valuation of:

- share price index futures (Section 6.2); and
- individual share futures (Section 6.3).

6.2 Share index futures

6.2.1 General description

A share price index (SPI) future is an exchange-traded contract based on a broad-based share price index. A buyer of an SPI future contract benefits from a rise in the value of the underlying index and loses from a fall in the index; the opposite applies for the seller. SPI futures are not deliverable; at expiry they are cash settled against the underlying index. Over the life of a contract the buyer should receive or pay the difference between the original purchase price and the final settlement price depending on whether the price has moved higher or lower.

For example, if a Standard and Poors (S&P) 500 futures contract is bought at a price of 600 and the price rises to 625, then the buyer will receive the gain equivalent to 25 index points. In this case each index point is worth US\$ 500, so the total gain is US\$ 2,500 per contract. The seller of this position is in the opposite situation, that is, it is facing a loss of US\$ 2,500.

An SPI futures contract is a more esoteric product than other forward contracts we have examined, in that the underlying is not a tradeable security and is not deliverable.²⁹ Many market observers often describe these products as purely a bet on whether the value of the underlying index will be higher or lower than the traded price on the expiry date, with a payoff linked to how right or wrong you are. In fact, when the S&P 500 contract was launched in 1982 it was widely perceived as a gimmick with little chance of success—by the end of the 1980s its notional daily volume exceeded all of the stocks traded on the New York Stock Exchange (NYSE). Despite its more esoteric nature the SPI future is a pure form of forward contract and its price is determined by the same cost of carry factors as any forward instrument. Further, the description of the SPI as a form of bet is just another way of defining a forward contract and can equally be applied to interest rate and equity futures.

28. See Francis, Toy and Whittaker (eds), *The Handbook of Equity Derivatives* (1995).

29. While the profit and loss behaviour of the SPI can be replicated by a portfolio of shares, this portfolio will not have exactly the same characteristics as the SPI future. For example, the SPI future price is based on an index whose value is related to its original base and how long it has been in existence.

As there is no underlying security, it is up to the futures contract specification to convert the “bet on a stock index” into a form which can easily be employed by equity market participants. The basic specification developed in the original US SPI futures contracts has been applied all over the world. The value of each contract is determined by multiplying the traded futures price by a fixed multiple. For example, the multiple in the FTSE 10 is GBP25, so if the futures price is 3,600 then the total contract value will be GBP90,000, and for every 1 index point change in price the value of this contract will change by GBP25. The Appendix to this chapter provides a summary of the contract specifications for listed SPI futures contracts.

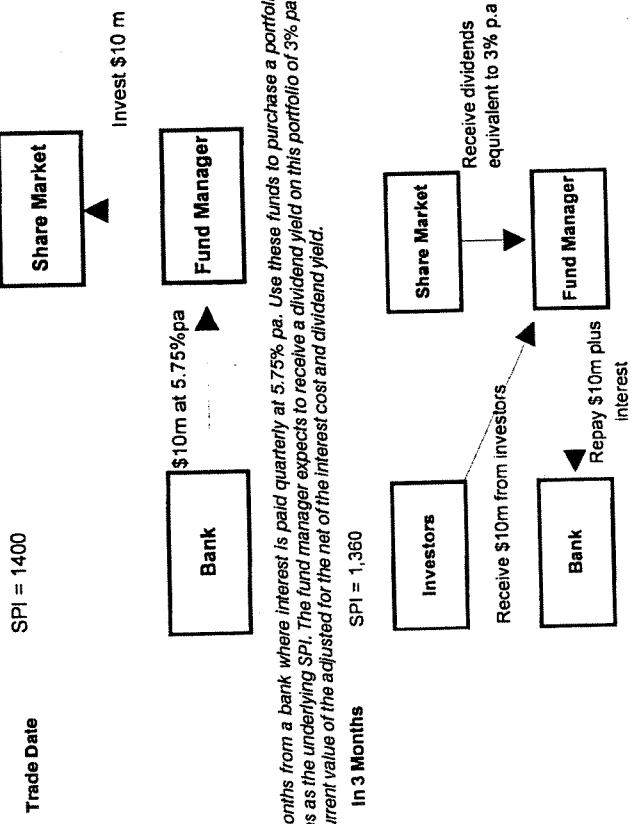
6.2.2 Synthetic replication of share price index futures

As with any forward contract, an SPI future can be replicated by using cash instruments. In this case the underlying asset is a portfolio of shares with the same weightings as the underlying index.

Suppose you are a fund manager operating in a share market without an SPI futures contract.³⁰ You intend to buy shares in three months time but wish to buy at prices prevailing today. The synthetic replication of an SPI futures contract would be to borrow funds to buy the portfolio of shares today and then repay the borrowing in three months time with the funds you intended to buy the shares with. The price of this forward purchase will be given by the current spot price minus the net financing cost of the portfolio. In this case the cost of carry is given by the borrowing cost minus dividends received on the portfolio.

30. This is a problem faced in a number of countries—while they may have active share markets, an SPI future does not exist and the only method of obtaining forward exposure to these markets is with a synthetic position.

Exhibit 4.38
Synthetic Replication of a Share Price Future Purchase



Borrow \$10m for three months from a bank where interest is paid quarterly at 5.75% pa. Use these funds to purchase a portfolio of shares containing the same shares and weightings of these shares as the underlying SPI. The fund manager expects to receive a dividend yield on this portfolio of 3% pa compounded quarterly. The Forward price on this portfolio will be the current value of the adjusted for the net of the interest cost and dividend yield.

After three months the fund manager receives funds from its investors and repays the borrowing. The effective price at which the portfolio was purchased in three months is the original value plus the cost of carry. From the buyers point of view the dividends received effectively reduce the forward purchase price while the interest cost increases the forward cost.

Exhibit 4.38—continued*Forward rate calculations—portfolio*

Original cash cost	=	10,000,000
Borrowing interest at three months	=	$10m \times (1 + 0.0575 \times 90 / 360)$ 143,750
Dividends received at three months	=	$10m \times (1 + 0.003 \times 90 / 365)$ 73,973
Net Cost of carry	=	69,777
Forward Price	=	10,069,777
Forward Premium	=	0.6978%

Forward rate calculations—index

Suppose we wish to express the forward price in terms of the SPI index. Essentially the index is another way of expressing the current value of a portfolio—in this case the value is 1,400 rather than \$10 million. The forward price calculation is then exactly the same as for the portfolio and is given by the net cost of carry. Using the premium as a summary of the net cost of carry then the three month forward SPI price will be:

Forward SPI	=	Cash SPI \times cost of carry
	=	$(1 + 0.006978)$
	=	1,409.77
Cost of carry	=	9.77

Gain or loss on forward purchase

The actual value of the index in three months time is 1,360—the value of the original portfolio has declined by 2.85%. The loss on your forward position is even greater because the cost of carry has increased your effective purchase price.

Loss	=	Index value – Forward purchase value
	=	1360 – 1409.77
	=	49.77 index points or –3.53%

Exhibit 4.38 provides an example of a synthetic purchase of SPI futures. As this example demonstrates, the index can be considered the current value of a portfolio of shares and the forward price of that index is given by the interest cost of financing that portfolio minus the expected dividends to be received on that portfolio.

6.2.2.1 Imperfections: slippage, transaction costs and short selling


Most SPI futures have “broad-based” underlying indices—that is, the index is made up of a large number of stocks, which are intended to replicate the performance of the whole stock market or the stock market leaders.³¹ Even in the case of the “leader” type indices the number of shares involved can create logistical execution problems for synthetic replication. For example, executing an order for 20 shares simultaneously, even on an electronic trading network, may be difficult. It is likely that the portfolio will suffer from “slippage” as you wait for the underlying shares to be purchased—so the index value you achieve may differ from the index value at the time you commenced the strategy.

As synthetic replication is the cornerstone of arbitrage activities, this requirement to buy or sell a large number of stocks can lead to arbitrage problems. This is one of the reasons why SPI futures can trade away from “fair market value”, or the forward price implied by synthetic replication. If an arbitrage exists it has to be of significant magnitude to overcome the risks of slippage and any other imperfections.

This deviation from fair market value is a characteristic that is most pronounced in less liquid futures markets or underlying share markets which have not developed methods of effectively purchasing portfolios. A common complaint by fund managers in all markets is that unless they hold futures contracts till expiry (and the futures price converges to the cash price), then the hedge provided by the futures contract can over- or under-perform the underlying index due to this mis-pricing.³² *Exhibit 4.39* shows the deviation from fair value of the world’s largest SPI futures contract, the S&P 500, from its inception to 1994. Not surprisingly the mis-pricing was greatest in the contract’s early days and during the stock market crash. Since 1992 the discrepancy has fallen significantly, as mechanisms for executing the cash market leg of an arbitrage improved and the market became less one-sided.

31. For example, the all ordinaries and all share indexes basically includes all shares with their market weightings while the BEL-20 and Toronto 35 Index are based on the top 20 and 35 shares respectively. The S&P 500, Nikkei 225 and FTSE 100 are somewhere in between.

32. In fact, a motivating factor behind the development of equity swaps was fund managers wishing to obtain a tailor made method of linking their performance to the underlying index without the mis-pricing risk with futures.



There are various methods that have been developed for decreasing the risk of slippage including:


1. *Imperfect portfolios*: This involves creating a portfolio with fewer stocks than the index but with a close correlation to the index. The obvious risk in this strategy is that any differences in price movements between the portfolio and the index undermine the arbitrage gains.
2. *Index portfolios*: In some share markets index portfolios are a traded instrument. Typically these are arranged by market makers and transaction costs are higher than for shares.
3. *Portfolio trades in electronic trading networks*: The development of electronic trading networks at futures and stock exchanges has improved the possibility of executing multiple stock buy and sell orders simultaneously—providing there is reasonable liquidity in each market. Many of the activities of so-called “program” traders involves a computer program monitoring share price levels, and then if an arbitrage appears automatically generating buy or sell orders for an electronic trading network. The crucial element in this type of operation is liquidity; even if a buy order can be executed instantaneously, unless there is a corresponding sell order at the same price the transaction will not be executed.

Another imperfection in synthetic replication is the transaction cost. As already noted, the transaction costs in stock markets tend to be higher than equity derivatives markets. While this tends to encourage the use of futures and options it also creates an additional cost in synthetic replication and hence arbitrage opportunities. In fact, even in well-developed markets which have introduced the measures above to reduce slippage, some mis-pricing

still occurs because of transaction costs. For example, if the share brokerage is 1%, then the futures price must move by more than 1% from fair value to make an arbitrage worthwhile.³³

Exhibit 4.40 compares the mis-pricing of the Australian all ordinaries index against the transaction costs of executing an arbitrage. As the graph demonstrates, the mis-pricing generally stays with the transaction cost hurdle.

33. It is important to note that the mis-pricing tends to be less than the arbitrage gap, because the major arbitrage players are often associated with share market-makers and their cost of execution is usually lower than customer executions.



Another important regulatory constraint on the synthetic replication of a sold forward position is any limitation on short selling and stock lending. A synthetic sold position requires the ability to sell a portfolio of stock without owning them and then funding the settlement of this transaction by borrowing the underlying shares till the forward expiry date. Short-selling is viewed suspiciously in many share markets and is occasionally sighted as a factor which leads to share price volatility and manipulation. Accordingly,

there can be considerable regulation around short-selling and share lending. In some markets, including Asian “tiger” countries such as Thailand, short-selling is not permitted in any form—meaning that a forward sale cannot be replicated. In these markets the fact that one side of the market is “missing” will undermine the development of an equity derivatives market.

From the point of view of synthetic replication and forward pricing, the greater the restrictions on short-selling and stock lending the more likely is mis-pricing of forward contracts. A market maker which buys a futures contract cannot hedge itself using a synthetic forward purchase and must sell another futures contract. If this tends to force down the price relative to fair value then this mis-pricing is likely to persist, as arbitragers will not be able to take advantage of the discrepancy.

In most markets the restrictions that apply to short-selling are considerably less onerous and consist of the following types of requirements:

1. *Eligible security*: Typically, stock exchanges only wish to allow short-selling in stock where manipulation is difficult and there is considerable liquidity. As a result, for a share to be eligible for short-selling it needs to meet minimum market capitalisation and turnover requirements.³⁴
2. *Limit on short-sales*: To avoid manipulation, the cumulative short sales of a particular stock by a group of related companies is limited to some percentage of the outstanding shares on issue (for example, 10% is used by a number of exchanges).
3. *Uptick and downtick rules*: In a number of exchanges a short sale must be identified and it cannot be made at a price lower than the last sale price. This can be referred to as both the uptick and downtick rule. Its effect is to constrain short-selling in a falling market.

These types of restrictions allow short-selling to occur with a minimum of constraints, and they will generally not be a cause of a divergence between futures prices and the fair forward price.³⁵ If the restrictions are more onerous for the index you wish to price, attempt to quantify the cost to the short-seller and incorporate this in another form of transaction cost.

6.2.3 *A model for the share price index futures price*

As the synthetic replication discussion in the section above indicates the forward price of an SPI future is based on the generalised price formulae from section 3. In this case the asset is the share and the asset return is the dividend. The funding cost will typically reflect the prevailing wholesale interest rate at which market makers and arbitragers can borrow for the forward term—typically a bank credit money market interest rate.

34. For example, in Australia the shares must have a market capitalisation in excess of A\$100 million, 50 million shares on issue and a ratio of share turnover to shares on issue in excess of 7 to 8%.
35. The downtick rule may cause some discrepancies where the price is free-falling (such as in a market crash), as the short-sellers will not be permitted to execute transactions. Though, in these conditions, forward mis-pricing is generally the least of anyone’s worries!! This can be seen in Exhibit 4.39 by the 10% discount in the futures price to fair value in October 1987.

Our generalised models can specify dividends as a constant percentage per annum or else as a lump sum. While dividends are generally paid at a fixed amount per share regardless of the price of that share, calculating this could become extremely complicated in an SPI future as there are potentially hundreds of dividends paid on different dates. However, the fact that there can be hundreds of stocks means that dividend payments throughout the year and in fact this stream of dividends is more like a continuous yield than a series of lumpy payments. For this reason, broad-based index models typically express dividends in terms rather than cents or pfennig per share.

Fortunately, this task is made somewhat easier by the fact that actual historical dividend yields are provided for most share indexes, and these yields tend to be fairly stable. For SPI futures pricing purposes we need to be able to estimate the expected dividends over the forward period and express it as a yield on the prevailing index price. Many market practitioners simplify this by taking the historical dividend yield and using this in the forward pricing model.

If we incorporate dividends as a per annum yield, the SPI futures pricing model is an extension of the constant income model as follows:

Share Price Index Futures Price (for Broad-Based Index)

Using simple interest, the calculation is as follows

$$F = S \times (1 + (r - q) \times f / D)$$

Where

- F = Forward SPI price
- S = Cash or spot price of the share price index
- r = Interest rate to the forward expiry date
- D = Day count basis (365 or 360)
- f = Number of days to the forward expiry date
- q = Dividend yield expressed as a % per annum on the same day count basis as the interest rate



This same model could be applied to an OTC share price index forward contract.

Exhibit 4.41 provides a spreadsheet model for calculating short-term SPI futures prices as well as any implied arbitrage opportunities. The example shown is for the FTSE-100, while a mis-pricing of 1.81 index points is calculated this represents a tiny proportion of the index (0.04%) and cannot be considered an arbitrage opportunity. In fact, in this example we have set the “arbitrage gate”, or minimum mis-pricing requirement at 10 index points—reflecting our anticipated transaction costs.

Exhibit 4.41 Share Price Index Futures Price Calculator		
Field	Cell	Cell: Formula
<i>Spreadsheet example—broad-based index only</i>		
Inputs		
Underlying SPI index	FTSE-100	
Current Cash Index Price	3680.4000	\$E\$9 :
Today's Date	01-Dec-95	\$E\$10 :
Forward Delivery Date	15-Mar-96	\$E\$11 :
Current Interest Rate (% pa)	6.4000	\$E\$12 :
Annual Dividend yield (% pa)	4.0000	\$E\$13 :
Current Futures Price	3704.0000	\$E\$14 :
Arbitrage "gate" (fut price equiv)	10.0000	\$E\$15 :
Contract "Tick size"	25.0000	\$E\$16 :
Outputs		
Implied Futures Price	3705.81	\$K\$18 : =E9*(1+(E12-E13)*(E11-E10)/365000)
Contract Value	92,600	\$K\$19 : =E16*E14
Cost of Carry (index points)	25.41	\$K\$20 : =E18-E9
Implied Arbitrage Opportunity (Indication only)	25.41	\$K\$21 : =ABS(E18-E14)
Potential Arbitrage	NO ARBITRAGE	\$K\$22 : =IF(E21>E15,IF(E18>E14, "BUY FUTURES, SELL PHYSICAL", "SELL FUTURES, BUY PHYSICAL"), "NO ARBITRAGE")

Exhibit 4.41—continued

Pricing sensitivity graphs



While this model provides a simple and generally accurate model there are some features of the model that need explanation:

- funding of margins;
- continuous dividends;
- forward price contango and dividend yield;
- dividend seasonality; and
- leader indexes.

These are discussed in the remainder of this section.

As with all futures contracts there will be a funding impact arising from the initial margins and daily mark-to-market gains and losses. Given most SPI futures' volume is concentrated in the first delivery month, these funding issues can generally be ignored. If, however, you do wish to include the funding effects, then follow the same interest adjustment steps as outlined in Section 4.9.

It is important to realise that this is a simple interest model, so the interest rate has no compounding and is a zero coupon interest rate. As we have noted elsewhere, this means most money market interest rates can be directly entered into the model. However, a feature of most calculated dividend yields is that they are equivalent to a continuous yield. While the difference on the forward price is usually insignificant, it is sometimes necessary to convert the dividend yield from a continuous rate to a simple interest rate appropriate for the period.

The implication of this model is that as long as dividend yields are lower than prevailing interest rates, which is the "usual" situation, then the futures prices will be higher than the cash index price—referred to as futures price "contango". As we noted in the synthetic replication discussion the futures price can mis-price relative to the "fair value" of the forward index and, as the cash value of the index is known as well as the interest rate, then this mis-pricing is often expressed as the expected dividend yield. It is a useful exercise to use this model to solve for the implied dividend yield in futures prices and look for any pricing discrepancies. *Exhibit 4.42* shows the SPI futures price curve for the S&P 500 in December 1995 and calculates the implied dividend yield.³⁶ The implied dividend yield can be solved by rearranging the SPI future price formula as follows:

$$q = \frac{(1 + r \times f / D - F / S)}{f / D}$$


36. This futures model and the implied "dividend yields" exercise can also be applied directly to commodity futures where the shape of the forward price can often be highly volatile. Obviously commodities do not pay dividends, so in the case of a commodity the asset return is relabelled the "convenience yield", which reflects the convenience of owning the asset today as opposed to some date in the future.

Exhibit 4.42**SPI Futures and Implied Dividend Yields****CME S&P 500**

Futures Month	Days to Expiry	Futures Price	Interest Rate A/365	Implied Dividend Yield % pa
Cash Index		613.70		
Dec-95	11	614.80	5.893	-0.05%
Mar-96	102	620.40	5.727	1.82%
Jun-96	194	625.45	5.640	2.04%
Sep-96	291	630.20	5.636	2.26%

382

Historical 1 year dividend yield 2%

S&P Futures Curve and Dividend Yield

Note: 1. LIBOR rates used and multiplied by 365/360 to make them comparable to the dividend yield result.

We have assumed that the dividend yield is a constant rate throughout the year. While a convenient assumption, there is often a "seasonality" in the payment of dividends. So, while the dividend yield for a whole year may be 4% pa, this may be made up of an effective yield of 4.5% pa in the first half of the year and 3.5% pa in the second half of the year. Depending on where the forward period occurs in these dividend seasons, the dividend yield should be altered. More sophisticated users will actually make forecasts of the dividend yield and take account of this seasonality. The dividend yield calculated in *Exhibit 4.42* tends to show very rough seasonality.

The first contract month has a dividend yield of effectively zero³⁷—the market expects no dividend effect on the index in the 11 days remaining to expiry. However, the remaining three delivery months have a dividend yield of approximately 2%, which closely resembles the dividend yield over the

37. The exhibit actually calculated in Exhibit 4.42 is -0.05%. The prospect of a negative dividend is highly unlikely—it is purely rounding error if the futures price is reduced by the minimum tradeable unit of 0.05 then the dividend yield is 0.22%, which is the closest the market can get to a zero dividend yield.

past year calculated as a percentage of the prevailing index price. This suggests that there is little seasonality in the S&P 500 dividend yield and/or the market as a whole ignores it.

The assumption of the dividend as yield is reasonable for broad-based indexes, however, for leader style indexes with a small number of shares (for example, 20 or less) this assumption may lead to errors as the dividend stream will start to take on a “lumpy” appearance. In fact, it is possible that for short-term SPI futures on a narrow-based index, there may be no dividend payments at all during the forward period. This model can be used as an estimate, however, the more appropriate solution is to use the individual share price model in section 6.3 to calculate the futures price of each share and then weight these prices using the underlying index weightings.

6.2.4 Valuation of share price index futures

The forward valuation of an SPI future is straightforward, reflecting the difference between the original contract futures price and the prevailing futures price. As with all futures contracts, all profits and losses are paid as they occur—accordingly, there is no difference between present and future values.³⁸ As a result the present value sensitivities of an SPI future are greater than an equivalent OTC forward by the size of the present value discount over future values.

Determining the risk characteristics of an SPI future follows the same PVBP concepts used previously: calculate the present value impact of a 1 point change in each of the pricing variables—cash SPI, interest rates and dividend yields—as well as the impact of 1 day passing.

6.3 Individual share futures

6.3.1 General description

While *Exhibit 4.37* shows that the volume in individual share futures (ISF) contracts is very small, there has been considerable interest by market participants and exchanges globally in this product over recent years. It is also a product which is worth understanding, as it represents the basic building block for pricing equity options and for narrowly-based index options.

An ISF is an agreement to buy or sell the underlying shares at an agreed date in the future. The buyer of a share future is not entitled to any dividends (that is, the ISF trades “ex-dividend”) over the forward period however the share futures will take part in any corporate events such as bonuses, stock splits and rights issues.³⁹

At present, ISFs are listed on the Swedish Exchange OM and the Sydney Futures Exchange. The OM contracts have physical delivery while the SFE contracts are currently cash settled. ISFs are directly comparable to the low exercise price option (LEPO) contracts listed on the Australian Stock

38. To review the valuation of a futures contract and differences with forwards see the discussion on short-term interest rate futures in section 4.

39. See the chapter on equity options in Martin, op cit n 1, to see the effect of these adjustments.

Exchange Derivatives and SOFFEX markets. LEPOs are call options on a stock with an exercise price close to zero. While LEPOs are strictly an option, the low exercise price makes it very likely that the option will be exercised and the LEPO behaves like a futures contract.

As with SPI futures the advantage of an ISF over cash market transactions is that it provides an exposure to share price movements at a lower transaction cost. *Exhibit 4.43* provides an example comparing the cost of trading shares versus ISFs.

Exhibit 4.43

Comparing ISF and Cash Market Transaction Costs

How could you have used share futures in the March quarter

Undertake a spread trade to capitalise on Newscorp's outperformance

One of the most outstanding events that has occurred in the sharemarket in the first three months of this year has been the outperformance of News Corporation shares against the sharemarket (as illustrated in the chart below).

To benefit from this, a trader could have undertaken a spread trade between NCP Share Futures and SPI futures contracts.

Scenario:

In January, a share trader believed that News Corporation shares would perform the sharemarket during the March quarter. To benefit from this expected scenario, the trader undertakes a spread transaction, by buying News Corporation (NCP) Share Futures contracts and simultaneously selling Share Price Index (SPI) futures and later closing not the futures position by undertaking a reserve spread transaction.

Spread implementation:

On 3 January, NCP Share Futures were trading at \$5.12 and the SPI was trading at 1927.

As illustrated in the table, this produces a contract value of \$7,270 ($\$5.12 \times 1,420^*$) for one SPI contract. This means that 6.6 NCP Share Futures contracts will be required for every one SPI contract sold (ie $\$48,175/\$7,270$).


The trader must however, allow for the volatility of NCP shares versus the overall market ie, the beta of NCP. At the time of implementing the spread transaction, NCP had a beta of 1.02. Therefore the number of Share Futures contracts required would be less and is calculated as 6.5 ($6.6/1.02 = 6.5$). The trader therefore buys 6 NCP Share Futures contracts at \$5.12 with a value of \$43,622 ($\$5.12 \times 1,420 \times 6$) and simultaneously sells 1 SPI contract at 1927 with a value of \$48,175.

Closing the spread:

On 21 March, the trader's expectation comes to fruition, with NCP Share Futures rising \$1.49 to \$6.61 (an increase of 0.3%). The share trader closes out the futures position by undertaking a reserve spread by selling 6 NCP Shares contracts and buying 1 SPI contract.

Exhibit 4.43—continued

News corporation outperforms the market in March quarter



Calculations required		NCP share futures	SPI
Actual prices	3/1/95	\$5.12	1927
	21/3/95	\$6.61	1932
	% change	+29.1%	+0.3%
Contract value:	@ 3/1/95	\$7,270	\$48,175
		(\$5.12 × 1,420)	(1927 × \$25)
No of contracts required		6.6	
		(\$48,175/\$7,270)	
Adjusted for NCP beta		6.5	
@ 1.02		(6.2/1.02)	
Nearest whole contracts required		6	

Result of the spread

The trader realises a profit of 149 cents on each NCP Share Futures contract and given that each 1 cent movement = \$14.20, this represents a profit of \$12,695 (ie 149 cents × \$14.20 × 6 contracts). This more than offsets the 5 point loss incurred on the SPI contract, which given each 1 point movement = \$25, translates into a loss of \$125 (ie 5 points × \$25 × 1 contract).

Overall the trader realised a net profit on this spread transaction of \$12,570, (before transaction costs). The trader has thus benefited from trading relative performance which has less risk than trading absolute performance.

Source: J S Martin, *Derivative Maths* (IFR, 1996).

6.3.2 A pricing model for individual share futures

The synthetic replication of an ISF is the same as for an SPI future, however, instead of the underlying being a portfolio of shares it is a single stock. The major impact from a pricing point of view is that the dividend payment is "lumpy". Depending on the dividend payment policy of the underlying company, the owner of shares will receive a dividend amount on a quarterly, semi-annual or annual basis. As a result the ISF model requires the magnitude and payment date of any dividends in the forward period.

While companies generally attempt to maintain a fairly stable dividend payment policy, it is dependant on the profitability of the company and its own internal cashflow requirements. In very general terms the dividend

payment of an individual company is not as stable as for a broad-based stock index. As a result, more effort is required in estimating the expected dividend on the underlying share in the forward period—a potentially difficult task if the ISF has an expiry date longer than one year.

The ISF pricing behaviour has something in common with a forward bond contract and the same type of “lumpy” asset income model can be applied:

Individual Share Future Pricing Model

Using simple interest and one income payment, the calculation is as follows

$$F = S \times (1 + r_1 \times f_1 / D) - c - (1 + r_2 \times f_2 / D)$$

Where

F	=	Forward price
S	=	Cash share price
r_1	=	Interest rate to the forward expiry date
r_2	=	Interest rate between the dividend “ex” and forward expiry dates
D	=	Day count basis (365 or 360)
f_1	=	Number of days to the forward expiry date
f_2	=	Number of days between the dividend “ex” and forward expiry dates
c	=	Dividend expressed in the same units as the cash price

Note: This model is also applicable to OTC share forwards.

In this model the dividend adjustment takes place on the date the underlying share is deemed “ex-dividend”, and shares purchased on that date are no longer entitled to that dividend. The dividend received then earns interest between the “ex” date and the futures expiry date.

As with SPI futures, longer term ISFs should take account of potential funding costs associated with holding the futures position.

Another interesting feature of ISFs is that they have to reflect the pricing behaviour of the underlying share in the case of a dividend payment. A feature of the Australian share market is that for a company which pays the full company tax rate a tax credit, known as a “franking credit”, is attached to the dividend payment.

Effectively, the company has pre-paid some tax owing on the share for the share holder. An example provided by the SFE on calculating the effective cost of physically hedging an ISF, incorporating the “franking” credit, is shown in *Exhibit 4.44*.

Exhibit 4.44**Cost of Hedging an SFE ISF****How to calculate the hedge price of SFEs share futures**

On 1 May, 1995 assume a futures broker, who is also a market maker received a call from a client interested in buying June BHP Share Futures (going long). As a market maker (which means they also take principal positions), the broker will take the other side of the deal and thus will sell (go short) BHP Share Futures. To protect his/her futures position, the market maker will simultaneously go long on physical BHP shares.

To quote the BHP Share Futures price, the market maker needs to calculate the fair price of the BHP Share Futures contract. There are two pricing formulas used to calculate the fair price of Share Futures, which are based on two standard arbitrage positions that a market maker can establish:

- (1) long the underlying shares and short the Share Futures
- (2) short the underlying shares and or long the Share Futures

In this particular example, the market maker has established position (1) and therefore will calculate the effective hedge price on the basis of being long shares and short Share Futures using the following formula.

$$\text{Share futures effective price} = (\text{SP} + (\text{SP} \cdot \text{CC} \cdot \text{Days} / 365)) - \text{DIV} + \text{Stamp duty Y. F.C.} + \text{T}$$


Whereby

SP	=	PURCHASE PRICE OF THE UNDERLYING SHARES
	=	ON 1 MAY BHP SHARES WERE TRADING AT \$18.18
CC	=	THE COST TO CARRY THE POSITION FOR THE DURATION OF THE TRANSACTION EXPRESSED AS AN ANNUAL PERCENTAGE RATE (SIMPLISTICALLY THIS IS THE RISK FREE RATE OF RETURN AVAILABLE IN THE MONEY MARKET FOR AN INVESTMENT WITH A SIMILAR DURATION)
	=	60%
DAYS	=	THE NUMBER OF DAYS THE CAPITAL IS INVESTED
	=	59 DAYS [1 MAY TO 29 JUNE (DAY OF EXPIRY OF JUNE BHP SHARE FUTURES CONTRACT)]
DIV	=	THE AMOUNT OF CASH DIVIDEND WAS PAID ON THE LAST DAY OF THE CONTRACT
STAMP DUTY	=	THIS IS THE 0.3% STATE GOVERNMENT LEVY ON THE PRINCIPAL SHARE VALUE
	=	\$18.18 * 0/003 * 2 (ROUND TURN)
	=	\$0.11 PER SHARE
F.C.	=	THE AMOUNT OF FRANKING CREDITS ATTACHED TO THE CASH DIVIDEND
	=	GIVEN THE 26 CENT DIVIDEND IS FULLY FRANKED AND THE TAX RATE IS 33%, THE AMOUNT OF FRANKING CREDIT IS:
	=	\$0.26 * 0.33 / (1 - 0.33) * 100%
	=	\$0.1281 PER SHARE
T	=	TOTAL TRANSACTION COST PER SHARE
	=	ARE APPROXIMATELY \$0.023 ROUND TURN

Based on the above factors, the "all-up" price for a hedge for the June 1995 BHP Share Futures contract as at 1 May 1995 would thus be:

$$\$18.18 + 6\% * 59/365 - 0.26 + 0.11 - 0.1281 + 0.023 = \$18.10$$

As expected the fair value of the June 1995 BHP Share Futures is priced at a discount to the physical shares due to the fact that the cash settlement process means Share Futures do not attract dividends and thus are quoted on an ex-dividend basis.



The valuation requirements of ISFs are the same as for SPI futures.


Chapter 5 Pricing Options


by Satyajit Das

1. OVERVIEW

Option contracts provide the purchaser with the right to either buy from (a call option) or sell to (a put option) the writer of the option at an agreed price (the strike price) a fixed amount of an underlying asset on (European style exercise) or any time before (American style exercise) a pre-nominated date (the expiry date). The purchaser pays the writer a fee (the premium) in return for effectively guaranteeing a maximum purchase price (call option) or minimum sales price (put option). Option valuation is concerned with the determination of the fair value of the premium.

The distinguishing feature of option valuation is the asymmetric nature of the payoff of the instrument. *Exhibit 5.1* sets out the pay-offs of option and forward contracts, highlighting this feature of option contracts.





The structure of the payoff dictates that the purchaser has unlimited potential profit with the loss limited to the premium paid. The writer in contrast has a limited gain (the premium received) but an unlimited loss. The structure of the payoffs means the actual value of the option is contingent on the future asset price either at maturity (for a European option) or the *path* or

sequence of asset prices between the entry into the option and the option expiry (for an American option because of the risk of early exercise).

This position complicates the process of risk neutral valuation for option instruments. In the case of a forward, the forward instrument payoffs can be replicated by a spot position in the asset, which is then financed or invested. For example, a forward purchase is hedged by a spot purchase of the asset which is then financed through till the forward maturity. This replication is allowed by the *symmetric* nature of the payoffs to the parties entering into the forward. In the case of an option, the asymmetric nature of the payoff dictates both a more complex method of replication of the option position and the determination of the fair value of the instrument.

In this chapter, the process of option valuation is examined. The process of option replication through trading in the spot asset is considered in Chapter 11. Two other aspects of option valuation, the sensitivities of the value of the option to each of the parameters used to value the option and the estimation of volatility, are considered in Chapters 10 and 8 respectively.

The structure of this chapter is as follows:

- the basic concepts of option valuation, including the key determinants of option value, are considered first;
- the concept of risk neutral option value is developed, using a simple probability based framework;
- the approach of risk neutral mathematical option valuation is introduced;
- two approaches to option valuation—the Black-Scholes model and the binomial model—are examined;
- the assumptions underlying the model and their violation in financial markets are then analysed; and
- the adaptations of the basic model for different problems of valuation—early exercise, dividend paying stock as well as the different asset classes—are then detailed.

2. OPTION VALUATION CONCEPTS

2.1 Option pricing nomenclature

The value of an option is known as at maturity because one of the key determinants of option value, namely, the asset price at maturity, is known. As at maturity, the price of an option is usually given as:

Call option

$$P_C = \text{Maximum } [0; S_m - K]$$

Put option

$$P_P = \text{Maximum } [0; K - S_m]$$

Where

P_C = Premium/value of call option

P_P = Premium/value of put option

S_m = Spot price of asset at maturity

K = Strike price of option

The maximum function is necessitated by the choice of option exercise, which rests with the purchaser who will not exercise the option unless it is economically advantageous to do so.

In determining the value of an option, it is usual to distinguish between:

- the *intrinsic value*; and
- the *time value* of an option.

An option's intrinsic value is based on the difference between its exercise price and the current price of the underlying debt instrument. If the option is currently profitable to exercise, it is said to have intrinsic value, that is, a call (put) option has intrinsic value if the current price of the instrument is above (below) the option's exercise price. The intrinsic value is given by:

For a call option

$$S_t - K$$

For a put option

$$K - S_t$$

Where

S_t = Spot price at time t

The concept of intrinsic value requires additional clarification in the context of the applicable exercise rules. In the case of a European option, as exercise is only permitted as at maturity, intrinsic value prior to that date requires the strike price to be discounted to the relevant date. This is usually given as Ke^{-rft} (that is, the strike price discounted back to the date of valuation at the continuously compounded risk free rate). The discounted strike price is then compared to the spot price of the asset. To a degree the concept of an intrinsic value of a European option prior to maturity is redundant as the option is incapable of being exercised prior to maturity. The computation of the *theoretical* intrinsic value of the European option prior to maturity is only relevant as a means for identifying the *sources* of value for the option.

In the case of an American option where the possibility of early exercise is present and permissible, the exercise price does not need to be discounted as it will be paid in full at the date of exercise.

Whether or not the option has intrinsic value, it may have time value. Time value is defined as:

$$\text{Option Premium} = \text{Intrinsic Value} + \text{Time Value}$$

Therefore

$$\text{Time Value} = \text{Option Premium} - \text{Intrinsic Value}$$

The time value of the option reflects the amount buyers are willing to pay for the possibility that, at some time prior to expiration, the option may become profitable to exercise.

The values identified are subject to a number of value constraints:

$$\text{Option Premium} \geq 0$$

$$\text{Intrinsic Value} \geq 0$$

$$\text{Time Value} \geq 0$$

Three other option valuation terms merit comment:

- in-the-money;
- at-the-money; and
- out-of-the-money.

It is customary for market participants to refer to particular options as belonging to one of the three groups. An option with an exercise price at or close to the current market price of the underlying security is said to be at-the-money. An option with intrinsic value is referred to as being in-the-money, while an out-of-the-money option is one with no intrinsic value, but presumably with some time value.

2.2 Factors affecting option values

The fundamental direct determinants of option value include:

1. the current price of the underlying asset (S);
2. the exercise price of the option (K);
3. interest rates (R_f);
4. the time to expiry (T); and
5. the volatility of prices on the underlying asset (σ).

Other factors affecting option valuation include the type of option (that is, whether the option is American or European) as well as payouts from holding the underlying instrument.

The general effect of each of the five major relevant variables on the value of an option (where all other variables are held constant) is summarised in *Exhibit 5.2*.

Factor	Effect of increase in factor on value of	
	Call	Put
Strike price	Decrease	Increase
Spot price	Increase	Decrease
Interest rate	Increase	Decrease
Time to expiry	Increase	Increase
Volatility	Increase	Increase

The effect of changes in the spot price of the instrument, option strike prices, and time to expiry on the pricing of options are relatively easily understood.

In the case of a call option, the higher the price of the underlying instrument, the higher the intrinsic value of the option if it is in-the-money

and hence the higher the premium. If the call is out-of-the-money, then the higher the underlying instrument's price the greater the probability that it will be possible to exercise the call at a profit and hence the higher the time value, or premium, of the option. In the case of put options, the reverse will apply.

The impact of changes in exercise price is somewhat similar. For an in-the-money call option the lower the exercise price, the higher the intrinsic value, while for an out-of-the-money call the lower the exercise price the greater the probability of profitable exercise and hence the higher the time value. A similar but opposite logic applies in the case of put options.

The impact of time to expiration and option valuation is predicated on the fact that the longer an option has to run, the greater the probability that it will be possible to exercise the option profitably, hence the greater the time value of the option.

The impact of volatility derives from the fact that the greater the expected movement in the price of the underlying instrument, the greater the probability that the option can be exercised at a profit and hence the more valuable the option or its time value. In essence, the higher the volatility, the greater the likelihood that the asset will either do very well or very poorly, which is reflected in the price of the option.

The impact of interest rates is less clear intuitively. The role of interest rates in the determination of option premiums is complex and varies from one type of option to another. In general, however, the higher the interest rate, the lower the *present value* of the exercise price the call buyer has contracted to pay in the event of exercise. In essence, a call option can be thought of as the right to buy the underlying asset at the discounted value of the exercise price. Consequently, the greater the degree of discount the more valuable is the right, hence, as interest rates increase and the degree of discount increases commensurately, the corresponding option value increases. In fact, a higher interest rate has a similar influence to that of a lower exercise price. A similar but opposite logic applies in the case of a put option. The higher interest rate decreases the value of the put option as it reduces the current (present valued or discounted) value of the exercise price that the buyer has contracted to receive.

An alternative way of looking at the impact of changes in interest rates is to view the option as a means for replicating the exposure to the asset. For example, a call option provides exposure to the asset in a manner analogous to a purchase of the asset. If interest rates increase, then the call option increases in attractiveness as a means of replicating the asset exposure. This is because the lower cash outlay entailed by the premium means that the cost (in forgone interest) is lower if the option is utilised. This means that the price of the call option is bid up. In the case of a put option, such a transaction can be viewed as an alternative to selling the asset. In this case an increase in interest rates reduces the premium of a put option as the proceeds of the sale are not received until the option is exercised and entails higher opportunity costs where interest rates increase. Decreases in interest rates have a similar impact but in reverse.

3. RISK NEUTRAL OPTION VALUATION

The development of a formal model for the valuation of options requires an understanding of the concept of risk neutral valuation techniques. The application of risk neutral valuation arguments is identical to that used in the context of the valuation of forward contracts.

The risk neutral approach is predicated on two central premises:

1. The value of any option must be equal to the expected payoffs under the instrument.
2. The value of the option can be determined by replicating the payoff or economic profile of an option using a position in the underlying asset and cash. This means that a short (long) position in the option can be hedged by the offsetting position in the asset and cash to create a riskless portfolio which should return the risk-free rate of interest.

The risk neutral valuation approach allows determination of the value of an option by arbitrage arguments which, very importantly, implies that the value of the option can be determined independent of risk preferences of the parties to the transaction as well as any expectations about the *direction* of underlying asset prices.

As in the case of pricing of forward contracts, the creation of this risk neutral portfolio to replicate the payoff profile of the derivative allows the value of the derivative instrument to be determined, because both the asset and the derivative contract have the same value driver—the changes in the price of the underlying asset.

The use of risk neutrality is central to both the valuation and trading or synthetic creation of options. In this chapter, the concept of risk neutrality is used to derive the value of the option contract. The process of using the concept of risk neutrality to synthesise option instruments is discussed in Chapter 11.

4. A BASIC OPTION PRICING MODEL

Within the framework of risk neutrality outlined, a very simple intuitive option pricing model can be derived. This model, in fact, exhibits all the characteristics and dimensions of option valuation generally.

The central unknown in valuing an option is the *forward* asset price as of the expiry date (this assumes a European style exercise). As with any other unknown variable in financial valuation, determination of value of the underlying contract is feasible by estimating the range of possible values and the probabilities attaching to individual possible outcomes.

Exhibit 5.3 sets out a basic option pricing model utilising this fundamental approach.

Exhibit 5.3
Simple Option Pricing Model

Assume the following scenario:

Current Spot Price of Asset (S) = 100
1 year Risk Free Interest Rate (Rf) = 10.00% pa

This implies an arbitrage free forward of 110 assuming that the asset does not pay any income. The *actual* spot price of the asset in one year's time is not known but is expected to be in the range set out below and the probability of the asset price being at any particular level is also specified:

Expected Asset Price In 1 Year	Probability (%)
90	10
100	20
110	40
120	20
130	10

The above table assigns a designated probability to all possible forward asset price states (effectively, the assumed distribution of forward asset prices).

This data can now be utilised to price the following call option:

K = 110
T = 1 year
Rf = 10.00% pa

The fair value of the option today should consistent with the risk neutral argument be the expected payoff of the option contract which is as follows:

Expected Asset Price In 1 Year	Probability (%)	Value of Call Option	Probability Adjusted Expected Value of Call Option
90	10	0	0
100	20	0	0
110	40	0	0
120	20	10	2
130	10	20	2
		Total	4

The expected value of the call option in this case is 4 at maturity of the option which must be discounted back at 10.00% pa for 1 year to provide the present value of the option. The current value of the option is 3.64.

The basic model requires, consistent with all option valuation models, the following input parameters:

1. the possible values or prices that can be assumed by the underlying asset as at option expiry;
2. the probabilities attaching to the possible values that can be assumed by the asset as at option expiry; and
3. the risk free interest rate to allow discounting of the expected values of the option.

In the basic option pricing model described, the asset price as at the forward date is restricted to five possible values. In reality, for a real world transaction, the range of possible values is large and theoretically infinite, although in reality many of the value states are unlikely or have very low probabilities. However, the requirement dictated by this model for both identification of each of the possible asset price states and the probability attaching to each possible asset price state is tedious.

In practice, the process of generating the possible forward asset prices and their probabilities is simplified by the introduction of an assumed asset price distribution.

5. ASSET PRICE DISTRIBUTIONS

The concept of a distribution of forward prices is central to option pricing theory. The seminal work of Black and Scholes in their path-breaking option pricing model is as much about their seminal work in simplifying the generation of the forward asset price as it is about the valuation of options.

Black and Scholes introduces a breathtakingly simple but robust assumption in generating the forward asset price distribution. In effect, they assumed that if the spot price of the asset and the *distribution of asset price changes (the asset returns)* are known, then it would be possible to determine the complete distribution of forward asset prices. To generate the distribution of asset price changes Black and Scholes assumed that the stock prices follow a process which is termed a continuous *random walk*. The introduction of this assumption means that the changes in the price of or continuously compounded returns on the underlying asset are normally distributed and that the forward asset price is lognormally distributed.

The introduction of these assumptions has a number of very significant implications. The forward asset prices and the probabilities of a particular asset price can now be calculated utilising the characteristics of a normal distribution. The major features of a normal distribution is that the *complete* distribution of asset prices and probabilities can be expressed in terms of two variables:

- the mean expected return (μ); and
- the standard deviation of the expected return (σ).


In the Black-Scholes approach, the mean expected return is taken to be the risk free rate of interest. In effect, the price of the underlying asset is expected to drift towards the forward price. The standard deviation of the expected return which equates to the volatility of asset prices (identified above as a key determinant of option values) is defined to be the standard deviation of the logarithmic returns on the underlying asset to the maturity of the option. The estimation of volatility is a particularly vexed issue and is dealt with in detail separately in Chapter 8. In the remainder of this chapter, volatility is assumed to be a known term.

The advantage of utilising a normal distribution is evident in that the two variables identified can also be used to infer the probability of the asset price

taking on particular values. For example, once the mean and standard deviation are known, the following probabilities are also known:

- 67% probability that the forward asset price will be between the mean (the forward price) ± 1 s;
- 95% probability that the forward asset price will be between the mean (the forward price) ± 2 s; and
- 99% probability that the forward asset price will be between the mean (the forward price) ± 3 s.

The introduction of the concept of log normality should also be explained. *Exhibit 5.4* sets out the comparative shape of a normal as against lognormal distribution.



The lognormal distribution differs from the symmetrical normal distribution in that it exhibits a skew with its mean, median and mode all differing from that in a normal distribution. The major advantage of using a lognormal distribution is that a lognormally distributed variable can only take positive values (between zero and infinity) in contrast to a normal distribution which allows variables to take on both positive and negative values.

The introduction of the assumptions allows the generation of a *complete distribution* of forward asset prices and their probabilities. This in turn allows *all* possible positive intrinsic values of option as at maturity to be generated. The intrinsic values are then weighted by their respective probabilities to determine the expected value of the option, which is then discounted to the present to calculate the option's fair value.

The introduction of the concept of the distribution of forward asset prices transforms the problem of option valuation in several respects:

- it makes the valuation of the option independent of the *direction* of asset price movements; and
- it converts the problem of option valuation into a problem of estimation of the *volatility* of the returns on the asset (more particularly, the annualised standard deviation of the log of the price changes).

The introduction of the asset price distribution assumption allows the development of mathematical option pricing models.

6. OPTION PRICING THEORY

6.1 Approaches to mathematical option pricing

Mathematical option pricing models seek to calculate the price of particular options, utilising the identified fundamental determinants of option value and incorporating these within a defined formalised mathematical framework. The technological approach of formal valuation is the same regardless of the type of option being evaluated, whether it be an option on commodities, equity stock or market index currencies, or a debt instrument as well as futures contracts on each of these assets. The fundamental approach also extends to non-standard or exotic options. The approach to valuation of these types of instruments is discussed in Chapter 7.

The development of mathematical option pricing models requires the following distinct steps:

1. specific assumptions as to market structure and behaviour of the underlying instrument;
2. definition of certain arbitrage boundaries on the potential value of the option; and
3. derivation of a pricing solution.

6.2 Market assumptions

Mathematical option pricing models can be developed to synthesise the many factors which affect the option premium within identified arbitrage boundaries. In order to develop such mathematical option pricing models, it is necessary to make a number of restrictive assumptions, including:

- asset trading is continuous, with all asset prices following continuous and stationary stochastic processes;
- asset has no intermediate cash flows (for example, dividends, interest et cetera);
- the asset price moves around in a continuously random manner;
- the distribution of the asset's return is log normal;
- the variance of the return distribution is constant over the asset's life;
- the risk free rate of interest is constant over the option's life;
- the option is European;
- no restrictions or costs of short selling; and
- no taxes or transaction costs.

Most of these assumptions are self-explanatory and are consistent with efficient capital market theory. Stochastic process refers to the evolution of asset prices through time modelled as random, characterised by continuous series of price changes governed by the laws of probability as prescribed. By continuous processes, it is usually implied that the price of the underlying asset can vary over time but does not have discontinuities or jumps, that is, the price movement over time of the asset could be graphed without lifting the pen from the paper.

The stochastic process, as assumed, is one that is determined the same way for all time periods of equal length. Specifically, the traditional approach to option valuation assumes that the price of the underlying asset has a particular type of probability distribution, assumed to be a log normal distribution. It is also assumed that the standard deviation of this distribution is constant over time.

6.3 Boundary conditions

Mathematical option pricing requires identification of certain boundaries that can be placed on the values of options based on arbitrage considerations. The concept of arbitrage in this context relies upon dominance whereby a portfolio of assets is said to dominate another portfolio if, for the same cost, it offers a return that will, at least, be the same. The underlying assumption, in this context, is that if these boundary conditions are breached, arbitrage activity would force the prices of the underlying assets and options within the arbitrage boundaries as arbitragers would enter into transactions designed to take advantage of riskless profit opportunities.

Exhibit 5.5 sets out the major boundary conditions on the price of an option. In the interest of clarity, in this section, the boundary conditions to option value are stated with reference to generalised assets. There are nine major arbitrage boundaries discussed in *Exhibit 5.5*. Put-call parity for options is a special arbitrage condition, which is discussed below.

Exhibit 5.5**Boundary Restrictions on the Value of an Option****Notation*

The following notation is used in outlining the boundary conditions:

- S = asset price
 S_m = asset price at maturity
 K = strike price
 T = time to maturity
 R_f = risk free interest rate
 σ = volatility of returns from asset
 P_{ca} = price of American call
 P_{ce} = price of European call
 P_{pa} = price of American put
 P_{pe} = price of European put
 $PV(K)$ = present value of amount K (that is, $Ke^{-R_f T}$)
 C = Intermediate cash flow on asset

Arbitrage Boundaries

$$P_{ce} \text{ or } P_{ca} \geq 0$$

$$P_{pe} \text{ or } P_{pa} \geq 0$$

This states that the value of an option is greater than or equal to 0. Option exercise is voluntary, consequently, purchasers will never exercise an option if the value of the option entails a loss and therefore option prices cannot take on negative values.

At maturity of the option:

$$P_{ce} \text{ or } P_{ca} = 0 \text{ or } S - K$$

$$P_{pe} \text{ or } P_{pa} = 0 \text{ or } K - S$$

The value of a call or put option will be either 0 or its intrinsic value at maturity. If this condition is not satisfied, arbitrage opportunities exist. For example, if a call at maturity sells for less than $S - K$, arbitrageurs could lock in a profit by borrowing enough to purchase the call and exercising it immediately making a riskless profit after paying back the loan.

Prior to maturity of the option:

$$P_{ca} \geq 0 \text{ or } S - K$$

$$P_{pa} \geq 0 \text{ or } K - S$$

If at any time prior to maturity, an American option contract sells for less than its intrinsic value, an arbitrage opportunity exists to purchase the option and exercise immediately while buying or selling the physical asset to lock in a riskless profit.

Exhibit 5.5—continued

$$P_{ca} \geq P_{ce}$$

$$P_{pa} \geq P_{pe}$$

An American option cannot sell for a premium value less than an identical European option. The American style of option confers all the benefits of the European contract plus the capacity of early exercise.

For an asset which has no intermediate cash flow, it can be economically demonstrated that it is superior to sell the option prior to maturity rather than being exercised early.

The validity for this arbitrage condition can be established by constructing the following two portfolios:

1. Buy a European call for P_{ce} and invest PV of strike price $PV(K)$.
2. Buy S .

The payoff from these two portfolios as at maturity of the options is:

Portfolio Value	Out-of-the-money ($S_m < K$)	In-the-money ($S_m \geq K$)
$P_{ce} + PV(K)$	$0 + K$	$(S_m - K) + K$
S	S_m	S_m

At maturity, the portfolio consisting of the European call and the present value of the strike price is never worth less than the asset, so the current cost of the first portfolio can never be less than that of the second. This implies that an American call option will usually never be exercised prior to maturity as the investor would only receive the intrinsic value of the option ($S - K$), which is less than $S - Ke^{-Rt}$ for any positive interest rate. Consequently, a rational investor will always sell a call option to somebody else rather than exercise the option.

For an asset which has intermediate cash flows, it can be economically demonstrated that early exercise is possible.

The proof of this strategy can be established by constructing portfolios which are similar to above. The two portfolios are as follows:

1. Buy a European call for P_{ce} and invest PV of strike price $PV(K + C)$ where C is the intermediate cash flow from holding the asset.
2. Buy S .

The payoff from these two portfolios is set out below as at maturity of the options is:

Portfolio Value	Out-of-the-money ($S_m < K$)	In-the-money ($S_m \geq K$)
$P_{ce} + PV(K + C)$	$0 + K + C$	$(S_m - K) + K + C$
S	$S_m + C$	$S_m + C$

The payout table indicates that at maturity the portfolio of the call and cash never pays less than the second portfolio consisting of the stock. The first portfolio gives a higher return when the option expires out-of-the-money. Consequently, the first portfolio cannot sell for less than the second portfolio. This means that whenever a European call is in-the-money prior to maturity the lowest value it can trade for will be equal to the stock price minus the investment required to receive an amount equal to the strike price plus any intermediate cash flow at maturity; that is, $S - PV(K + C)$. The lower limit on the value of the European call must be the lower limit on an American option's value.

This implies an optimal exercise policy for an American option on the asset with intermediate cash flows. If the lower limit on the call's value [$S - PV(K + C)$] is greater than the amount received by exercising ($S - K$), then it is better to sell than exercise. However, if [$S - PV(K + C)$] is less than ($S - K$), then the American call should be exercised rather than sold prior to maturity.

Exhibit 5.5—continued

An American call should therefore only be exercised early where the discounted value of the strike price and intermediate cash flows are greater than the strike price. In essence, it is $PV(C)$ that determines whether the option is sold or exercised. This implies that an American option where the underlying asset has high intermediate cash flows is more likely to be exercised early.

$$P_{ce} \text{ or } P_{ca} \leq S$$

$$P_{pe} \text{ or } P_{pa} \leq S$$

This can be illustrated with an example. A call option cannot be worth more than the underlying asset because if the option were worth more than the asset, then a riskless arbitrage profit could be made by writing a call and using the proceeds to buy the asset. If the call is exercised, the asset can be delivered and the strike price received in return, while if the call is unexercised at maturity the asset which has a positive value will be held, thereby allowing the arbitrageur to make a positive profit without incurring any risk.

P_{ce} or P_{ca} must be worth less than an identical option with a lower exercise price; and P_{pe} or P_{pa} must be worth less than an identical option with a higher exercise price.

In this case, the call with the low exercise price offers a greater chance of being in-the-money; consequently, it cannot sell for a price which is lower than an option which has less chance of being in-the-money. The reverse is true for put options.

P_{ce} or P_{ca} or P_{pe} or P_{pa} cannot be worth less than an identical option with a shorter time to maturity.

Intuitively, the longer the maturity on the option, the greater the opportunity for there to be a sufficiently large change in the asset price to push the option into the money. Consequently, an option with a longer maturity cannot sell for less than an equivalent option with a shorter maturity. If this condition is violated, an arbitrage can be set up whereby the arbitrageur writes the shorter-dated option while purchasing the option with the longer maturity to lock in a riskless arbitrage profit.

* The description of option boundary conditions draws on Ian Rowley, "Pricing Options Using the Black-Scholes Model" (1987) (May) *Euromoney Corporate Finance* 108.

6.4 Put-call parity

Put-call parity which defines the relationship between the price of a European call option and a European put option with the same exercise price and time to expiration is an additional arbitrage boundary on option values. Utilising the same notation as that used previously, put-call parity can be stated as follows:

$$P_{ce} + PV(K) = P_{pe} + S$$

This implies that buying a call and investing PV of K is identical to buying a put and buying the asset.

The proof of this relationship can be established by setting up two portfolios:

1. buy a European call for premium P_{ce} and invest $PV(K)$; and
2. buy a European put for premium P_{pe} and buy the asset for S .

The payoffs on the two portfolios at maturity are as follows:

Portfolio value	Out-of-the-money ($S_m < K$)	In-the-money ($S_m > K$)
$P_{ce} + PV(K)$ $P_{pe} + S$	$0 + K$ $(K - S_m) + S_m$	$(S_m - K) + K$ $0 + S_m$

At maturity, the two portfolios are equal irrespective of whether the option expires in or out-of-the-money.

The put-call parity condition can be restated as follows:

Synthetic call (reversal)

$$P_{ce} = P_{pe} + S - PV(K)$$

Synthetic put (conversion)

$$P_{pe} = P_{ce} - S + PV(K)$$

Long asset/forward

$$S = P_{ce} - P_{pe} + PV(K)$$

Short asset/forward

$$-S = P_{pe} - P_{ce} - PV(K)$$

For European options, arbitrage possibilities will exist if the put-call parity conditions are not fulfilled. An example of put-call parity arbitrage is set out in *Exhibit 5.6*.

Exhibit 5.6
Put-Call Parity Arbitrage

Assume that the a forward on an asset is trading for 1 month forward at 86.14/86.15. Call options on the contract with strike price 86.00 on the contract are trading at 0.28/0.33. Put options with an identical strike price are trading at 0.02/0.07. Both options are on the forward contract. In these circumstances, it is possible to create a synthetic 86.00 call at less than 0.28 as follows:

- buy 86.00 put at 0.07;
- buy forward at 86.15; and
- sell 86.00 call at 0.28.

This transaction effectively creates a call at less than the 0.28 received. This can be proved as follows: the sold call and the bought put are equivalent to a synthetic short forward position at a price of 86.00. The position creates a net cash flow to the grantor of 0.21. Of the 0.21, 0.15 is lost through the bought forward position at 86.15 which is above the synthetic short price of 86.00. However, the forward loss of 0.15 is more than offset by the 0.21 gain on the option.

Where the underlying asset pays out a cash flow of C , put-call parity can be restated as:

$$P_{ce} + PV(K + C) = P_{pe} + S$$

It is important to note that the put-call parity theorem is only valid for European options. Synthetic positions for American options are not always pure. For example, if S decreases and an American put is exercised, then you

would lose the difference between K and S immediately, not at the forward date.

This means that put-call parity for American options can be stated as follows:¹

$$P_{ca} - S + PV(K) < P_{pa} < P_{ca} - S + K$$

6.5 The concept of a riskless hedge

The derivation of the mathematical option pricing model also requires understanding of the concept of a riskless hedge. A riskless portfolio by definition consists of an asset and a corresponding option held in proportion to the prescribed hedge ratio, continuously adjusted, whereby the portfolio is perfectly hedged against movements in asset prices as changes in the call price and the asset price are mutually offsetting.

Certain riskless portfolio positions are set out below:

Position	Hedge
long position in calls	short Δ assets for each call held
short position in calls	long Δ assets for each call sold
long position in puts	long Δ assets for each put held
short position in puts	short Δ assets for each put sold

Delta (Δ)² refers to the sensitivity of the option premium to changes in the asset price.

Utilising the constructs of portfolio theory or the capital asset pricing model, it can be predicted that a riskless portfolio should earn no more than a risk free rate of return. It is important to note that outside the context of this riskless hedge construct, the values derived by mathematical option pricing models are not meaningful.

Mathematical option pricing models, such as Black-Scholes, utilise the concept of the riskless hedge to set up portfolios of the asset and cash which are managed dynamically over time to replicate the payoff of an option. It is then possible to utilise the techniques of stochastic calculus to derive a mathematical solution to the valuation problem.

7. THE BLACK-SCHOLES OPTION PRICING MODEL

Black and Scholes³ were the first to provide a close form solution for the valuation of European call options. The mathematical derivation of the Black-Scholes Option Pricing Model is beyond the mathematical capabilities assumed for this Chapter. A detailed derivation of Black-Scholes is set out in Chapter 19.

1. For a mathematical proof see John C Hull, *Introduction to Futures and Options Markets* (Second edition, Prentice Hall, Englewood Cliffs, NJ, 1995), pp 210-218.
2. See detailed discussion in Chapter 10.
3. Fischer Black and Myron Scholes, "The Pricing of Options and Corporate Liabilities" (1973) 81 *Journal of Political Economy* 637.

Black-Scholes option pricing model is usually specified as follows:

$$P_{ce} = S \cdot N(d1) - K e^{-R_f T} \cdot N(d2)$$

Where


$$d1 = [\ln(S/K) + (R_f + \sigma^2/2) T] / \sigma \sqrt{T}$$

$$d2 = [\ln(S/K) + (R_f - \sigma^2/2) T] / \sigma \sqrt{T}$$

$$= d1 - \sigma \sqrt{T}$$

A number of aspects of the formula require explanation:

- $N(d1)$ and $N(d2)$ are cumulative normal distribution functions for $d1$ and $d2$. *Exhibit 5.7* sets out the concept of the cumulative normal distribution function diagrammatically. In the diagram, the area given is the probability that a variable with a normal distribution will be less than $d1$;
- \ln is the logarithm of the relevant number; and
- $K e^{-R_f T}$ is the amount of cash needed to be invested over period or time T at a continuously compounded interest rate of R_f in order to receive K at maturity.



The price of a European put option can be derived by utilising put-call parity:

$$P_{pe} = K e^{-R_f T} \cdot N(-d2) - S \cdot N(-d1)$$

Two aspects of the Black-Scholes' Option Pricing Model require comment:

1. The calculation of the cumulative normal distribution function $[N(d)]$ is undertaken either utilising a $N(d)$ table or directly utilising numerical procedures. *Appendix A* to this chapter sets out the methodology for calculating cumulative normal distribution function.
2. The model requires specification of the volatility of prices on the underlying instrument (parameter σ in the above equation). Techniques

for the estimation of volatility, including discussion of the various issues thereto, is detailed in Chapters 8 and 9.

Exhibit 5.8 sets out an example of utilising the Black and Scholes Option Pricing Model to calculate the price of an option.

Exhibit 5.8

Using Black-Scholes Option Pricing Model

Calculate the price for a call and put option on an asset based on the following information:

S = 105.00
 K = 100.00
 T = six months (0.5 yrs)
 RF = 10% pa (0.10)
 $\sigma = 20\%$ (0.20)

Using the above inputs, we can compute the call option price as follows:

$$d1 = [\ln(105/100) + (0.1 + 0.20^2/2) 0.5]$$

$$= 0.769$$

$$d2 = 0.769 - 0.141 = 0.628$$

Using the normal cumulative distribution table:

$$N(0.769) = 0.7791$$

$$N(0.628) = 0.7349$$

Therefore, the call option value is:

$$P_{ce} = 105 \times 0.7791 - 100 \times e^{-0.10 \times 0.50} \times 0.7349$$

$$= 11.900$$

The value of the call option is \$11.90 or 11.33% of Asset Price.

The value of the call can be dissected as follows:

$$\text{Intrinsic value} = 105 - 100 \times e^{-0.10 \times 0.50} = 9.877$$

$$\text{Time value} = 11.90 - 9.877 = 2.023$$

The corresponding put option value is:

$$N(-0.769) = 0.2209$$

$$N(-0.628) = 0.2651$$

$$P_{pe} = 100 \times e^{-0.10 \times 0.50} \times 0.2651 - 105 \times 0.2209$$

$$= 2.023$$

The value of the put option is \$2.023 or 1.9261% of asset price. The put value is all time value as the option is out-of-the money.

The actual components of the formula show how the value of the option is determined by the combination of asset and borrowing that replicates the payoff profile of an option.

The formula actually represents the replication of the call option through investment in the asset and borrowing to finance the position with the position being adjusted over time in line with asset price movements.

- The term $S \cdot N(d1)$ represents the amount of asset which must be held and financed. The term $N(d1)$ is effectively the delta of the option.

- The term $Ke^{-R_f T}$ represents the amount that must be borrowed to finance the holding of the asset.
- The actual premium represents the difference between the terms which ensures cash neutrality of the portfolio.

Intuitively, assuming physical settlement of the option, if the option is exercised, then the seller will have to transfer to the buyer assets valued at S_m . In effect, the strike price K will need to be satisfied through delivery of assets valued at S_m . This means that the S.N (d1) is, intuitively, the present value of receiving the asset in the event of exercise. The second term— $Ke^{-R_f T}$ would under this approach represent the present value cost of the strike price in the event of exercise. In a risk neutral world, if the option is likely to be exercised, then the difference between the two terms would represent the expected payout of the option which in turn would equate to the premium to render the transaction a zero return transaction.

This is evidenced by the fact that if the option is deep in-the-money then both $N(d1)$ and $N(d2)$ approach 1. This means that the call value approaches $S - Ke^{-R_f T}$. Similarly, as the option approaches expiry, that is, T approaches 0, both $N(d1)$ and $N(d2)$ approach 1 and $e^{-R_f T}$ also approaches 1. This means that the call value approaches $S - K$.

8. BINOMIAL OPTION PRICING MODEL

8.1 Concept

The binomial option pricing model utilises an identical logical approach to Black-Scholes. However, in contrast to Black-Scholes, the binomial approach assumes that the security price obeys a binomial generating process. The binomial approach also assumes that the option cannot or will not be exercised prior to expiration (that is, the option is European).

The valuation process begins by considering the possibility that the price can move up or down over a given period by a given amount. This enables calculation of the value of the call option at expiration of the relevant period (which is always the greater of zero or the price of the instrument minus the exercise price). The riskless hedge technique starts at expiration and works backwards in time to the current period for a portfolio consisting of the physical security sold short or one sold futures contract on the relevant asset and one bought call option on the relevant asset.

Since the portfolio is riskless, it, consistent with Black-Scholes, must return the risk free rate of return over the relevant period. The derivation of the value of the call option using this approach is predicated on the fact that the call option must be priced so that the risk free hedge earns exactly the risk free rate of return.

Exhibit 5.9 sets out a simple example of a 1 step binomial model.

Exhibit 5.9**One Step Binomial Option Pricing Model**

In order to illustrate the logic of a binomial option pricing model, consider the following example:

$$S = 100$$


The asset price is expected to increase or decrease by 10% to 110 or 90 respectively over the next year. Assume a call option utilising the following parameters:

$$S = 105$$

$$T = 1 \text{ year}$$

$$R_f = 10\% \text{ pa}$$

The value of the call option can be ascertained based on the expected increase or decrease in the asset price as follows:



To determine the fair value of the option it is necessary to create a riskless hedge, entailing investment in the asset to offset the position in the call, such that the portfolio value is known with certainty *irrespective of whether the asset price increases or decreases*.

The construction of the riskless portfolio requires the following steps:

1. Assume the position consists of 1 sold call which is offset by holding Δ units of the asset.
2. The value of Δ is determined on the basis that it will be equate to a value which makes the portfolio riskless. This entails that the value of portfolio will be the same for both an increase and a decrease in the asset price. Therefore:

$$110\Delta - 5 = 90\Delta$$

$$\Delta = 0.25$$

This means that to hedge or replicate 1 sold call it would be necessary to hold .25 units of the asset. Irrespective of whether the asset price moves up or down, the portfolio has a value of 22.5 as at the expiry of the option.

Based on the intuition that a riskless portfolio should earn the risk free rate of interest, it is now possible to derive the fair value of the option based on the following steps:

3. The value of the riskless portfolio constructed must in present value terms be:

$$e^{-R_f T} 22.5 = e^{-.10 \cdot 1} 22.5 = 20.36$$

4. The fair value of the option can now be calculated using the known value of the portfolio of assets at commencement of the transaction ($.25 \times 100 = 25$) as follows:

$$\text{Value of Asset Portfolio} - \text{Option premium} = \text{Value of Riskless Portfolio}$$

$$25 - \text{Option Premium} = 20.36$$

$$\text{Option premium} = 4.64$$

Exhibit 5.9—continued

The fair value of the call option is 4.64. If the option was trading at a value higher (lower) than the fair value, then the riskless portfolio would cost less (more) than the option premium to create allow the creation of a portfolio which yield more than the risk free rate of return (a mechanism for borrowing at less than the risk free rate). In either case, the value difference would attract arbitrage to realign the value of the option to eliminate the possibilities for arbitrage.

8.2 Generalised binomial option

In the binomial model, the time period to option expiry is divided into a series of discrete intervals. This contrasts with the continuous time model of Black-Scholes. However, as the number of intervals to maturity increases, the resulting increase in final stock prices begins to approximate the continuous log normal distribution. This allows a more generalised version of the model to be created. *Exhibit 5.10* sets out the generalised version of the binomial model.

Exhibit 5.10**Generalised Binomial Option Pricing Model**

In order to generalise the Binomial Option Pricing Model it is necessary to define: the factor amount the asset price can go up or down as at each step of the binomial tree and the probability of an up or a down move.

It can be shown that:*

$$\begin{aligned} u &= e^{\sigma \sqrt{t/n}} \\ d &= e^{-\sigma \sqrt{t/n}} \\ &= 1/u \\ p &= (e^{rf.t/n} - d) / (u - d) \end{aligned}$$

Where:

- u = the factor amount the stock price can go up
- d = the factor amount the stock price can go down
- e = the exponential term
- σ = the volatility of logs of the returns of the asset price
- t = time to expiry of the option
- n = number of steps
- p = the probability of an upward move in the asset price
- rf = the risk free rate of interest

The value of an option in a 1 step tree can therefore be stated as:

$$P = e^{-rf.t/n} [p \cdot P_u + (1 - p) P_d]$$

Where:

- P_u = the option value on an increase in the asset price
- P_d = the option value on a decrease in the asset price

The model can be extended using a similar logic for multiple steps.

* See Hull, op cit n1, Chs 10 and 15 for a full proof and derivation of these relationships.

Exhibit 5.11 sets out a multi-step binomial option model using the generalised model created. *Exhibit 5.12* sets out a multi-step binomial model for a put option.

Exhibit 5.11**Multi-Step Binomial Model For European Call Option**

Assume the following parameters for a European call option:

$S = 100$

$K = 100$

$T = 1$ year

$R_f = 10\%$ pa

$\sigma = 20\%$ pa

Assume also the number of steps (n) to be used is 2.

Exhibit 5.11—continued


In order to use the binomial model, it is necessary to calculate u, d, and p.

$$u = e^{.20 \sqrt{1/2}} = 1.151910$$

$$d = 1/u = 0.868123$$

$$p = (e^{-.10 \cdot 1/2} - 0.868123) / (1.151910 - 0.868123) = .645371$$


This allows the construction of the binomial tree as follows:



In order to now derive the value of the option, it is necessary to work back through the tree solving for the price of the option as at each node of the tree. In the above case the major node which is relevant is node B where the value of the option can be given as:

$$P = e^{-.10 \cdot 1/2} [.645371(32.69) + (1 - .645371) 0] = 20.07$$

This allows restatement of the tree as follows:



The value of the option at commencement (node A) can now be calculated as follows:

$$P = e^{-.10 \cdot 1/2} [.645371(20.07) + (1 - .645371) 0] = 12.32$$

Exhibit 5.11—continued

The complete tree therefore is as follows:

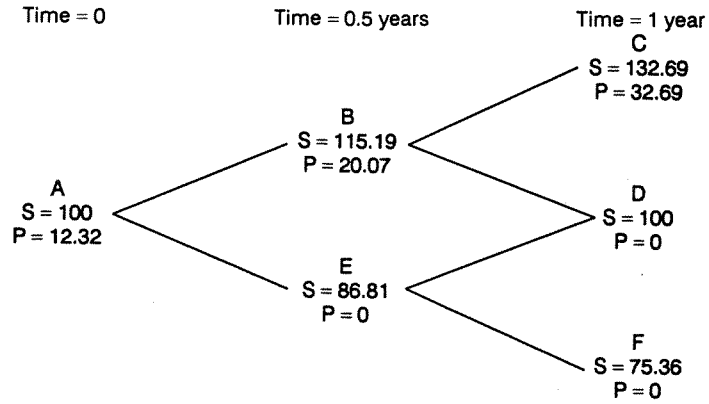


Exhibit 5.12**Multi-Step Binomial Model For European Put Option**

The binomial model can be used to value a European put option in an entirely similar way. Assume the same parameters as for the example in *Exhibit 5.12* with the exception that the option is a European put option.

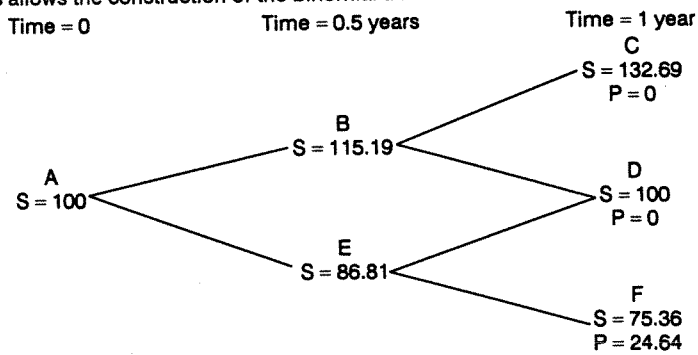
The inputs u , d , and p are as before:

$$u = e^{.20 \sqrt{1/2}} = 1.151910$$

$$d = 1/u = 0.868123$$

$$p = (e^{-.10 \cdot 1/2} - 0.868123) / (1.151910 - 0.868123) = .645371$$

This allows the construction of the binomial tree as follows:

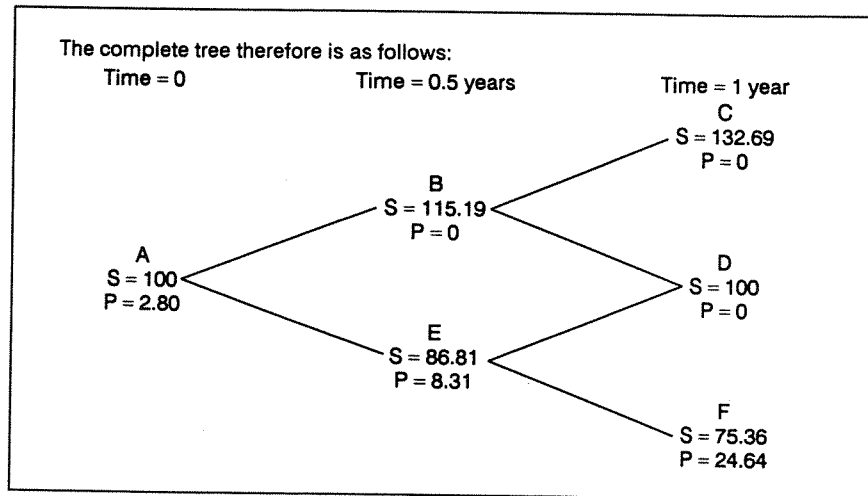


In order to now derive the value of the option, it is necessary to work back through the tree solving for the price of the option as at each node of the tree. In the above case the major node which is relevant is node E where the value of the option can be given as:

$$P = e^{-.10 \cdot 1/2} [.645371(0) + (1 - .645371) 24.64] = 8.31$$

This in turn allows calculation of the value of the put at commencement (node A) as:

$$P = e^{-.10 \cdot 1/2} [.645371(0) + (1 - .645371) 8.31] = 2.80$$



8.3 Numerical solution issues

The simple binomial models with small numbers of steps are relatively easy to solve. However, it is necessary to increase the number of steps significantly to increase the accuracy of the estimate of the value of the option.

In general, the minimum number of steps utilised in valuing options is around 50 to 100. This means that there are:

- $n + 1$ terminal stock prices ($n =$ the number of steps) or 51 ($50 + 1$) terminal stock prices; and
- 2^n possible paths or 2^{50} possible price paths.

The large number of paths places substantial demands on numerical techniques to solve through the tree to estimate the price of the option.

8.4 Features of binomial models

The binomial option pricing model contains the Black-Scholes formula as a limiting case. If, for the binomial option pricing model, the number of sub-periods are allowed to tend to infinity, the binomial option pricing model tends to the option pricing formula derived by Black and Scholes.

The major advantage of the binomial model is that as the time to maturity is segmented into a series of discrete variables the model can take into account specific option values prior to maturity. This allows the binomial approach to be used to provide a solution not only for the closed form European option pricing problem but also for the more difficult American option pricing problems when numerical simulation approaches must be employed (see more detailed discussion below). In essence, the binomial pricing approach is useful as it can accommodate more complex option pricing problems, such as non-constant interest rates and volatility, debt options and exotic options such as path dependent structures.

9. OPTION PRICING MODELS—ISSUES IN APPLICATION

The major attraction of option pricing models, such as Black-Scholes and its binomial variations include:

- the fact that all input variables, other than volatility, are directly observable; and
- the models do not make any reference to the investor's attitudes to risk.

While the model plays a central role in option valuation trading, the underlying assumptions do not necessarily hold true in practice. In particular, violations of the model's assumptions exist in the following areas:

- asset price behaviour;
- constant and measurable volatility;
- constancy of interest rates;
- no intermediate cash flows; and
- the issue of early exercise.

Some of these violations are significant (like assumptions about asset price behaviour), while others are relatively minor (constancy of interest rate except for debt options, intermediate cash flows and early exercise) in terms of their impact on the validity of these models.

The key assumption, that price changes are continuous through time (that is, the assumption that there are no jumps or discontinuities between successive asset prices), independent and log normally distributed over time with constant variance, may be violated in practice. The assumption of independence of asset price changes, as required by efficient market theory, is not wholly convincing. The empirical evidence and support for the log normal distribution of asset prices and its constancy over time is also not completely convincing.⁴ It is clear that option prices are sensitive to the stochastic processes assumed and changes in the assumptions produce significant, large percentage changes in option prices.

Empirical research highlights that *true* distributions differ from theoretical normal distribution in two respects:

1. The distributions of actual asset price changes are characterised by *fat tails*, that is these distributions display larger extreme price changes (both positive and negative) than implied by a theoretical normal distribution. This means that the theoretical models would *underprice* out-of-the-money and in-the-money option. This reflects the fact that the theoretical distribution allocates a lower probability to very high intrinsic values than is the case in reality.
2. Asset price behaviour appears to be characterised by discontinuities in the asset price changes or jumps. This contributes to the fat tails of the distribution.

The violation of the asset price behaviour assumptions underlying Black and Scholes has prompted the development of variations on the basic model

4. For a discussion of the log normality of asset price changes see Chapter 16 where the evidence is considered in the context of value at risk calculations which also rely on the assumption of log normality.

which make use of alternative stochastic processes, including absolute diffusion, displaced diffusion, jump processes and diffusion-jump processes. Empirical tests have tended to show that these alternative models are not able to provide better predictions of actual prices than the Black-Scholes type of model on a consistent basis. The price differences resulting from differing assumptions as to the underlying asset price movements in fact are *no* greater than the price differences that result from different assumptions of volatility.

The asset price volatility factor required as an input to option pricing models must be forward looking, that is, a forecast of the probable size (although not necessarily the direction) of asset price changes between the present and the maturity of the option. The problem in volatility estimation (the determination of the true constant volatility of the asset price) is, in practice, sought to be overcome by utilising two types of volatility: historical and implied.

Historical volatility is based on past prices of the underlying asset computed as the standard deviation of log relatives of daily price returns (usually annualised) over a period of time. Utilising historical volatility requires the selection of the period over which price data is to be sampled. It is possible to utilise price information over long periods (up to five years or longer) to derive the volatility estimate. This assumes that volatility is constant over long periods. It is also possible to use a much shorter period (less than 30 days) to get a good estimate of the current level of volatility. It is necessary to adjust the volatility input into the option pricing formula on a regular basis where short-term volatility is used on the basis that the volatility actually varies significantly.

Implied volatility is determined by solving an option pricing model (such as Black-Scholes) in reverse, that is, calculating the volatility which would be needed in the formula to make the market price equal to fair value as calculated by the model. Where this method is used, the implied volatility equates the model premium to the actual premium observed in the option market. An interesting problem with implied volatility measures is that options with different strike prices but with the same maturity often have different implied volatility. This is the phenomenon of the volatility smile, which is dealt with in Chapter 10. In addition, volatility appears to illustrate a term structure.

Historical volatility is a measure of past, already experienced, price behaviour. To the extent the option pricing model is validated, implied volatility reflects market expectations of future price behaviour during the life of the option. Both measures are important, and comparing the two can reveal interesting insights into the market in the underlying asset. However, no normative rule for derivation of the volatility estimate is available. This means that in reviewing option premiums, particularly where the value of the option in question is sensitive to the volatility estimate utilised, any option price suggested by an option model must be regarded with caution.

In addition, the models usually assume constant volatility. This assumption is clearly breached in practice, as volatility changes over time often very significantly.

Some attempted solutions to the volatility measurement model have sought to explicitly take into account the stochastic nature of volatility itself by

using multi-factor numerical techniques which utilise two stochastic variables, namely, the asset and the volatility. A detailed discussion of issues pertaining to estimation of volatility is set out in Chapters 8 and 9.

The assumption that interest rates are constant is particularly problematic in the case of options on some debt instruments. This is because interest rate changes drive asset price changes where the asset itself is an interest bearing security. In addition, the volatility of asset prices in the case of debt instruments is a function of remaining maturity and, in turn, interest rates, which reflect the shape of the yield curve. As maturity diminishes, the volatility of the asset also diminishes and constant variance cannot be assumed.

The impact of intermediate cash flows depends on the pattern of payments and the certainty with which the cash flows can be predicted. The Black-Scholes model does not appear to be very sensitive to assumptions about intermediate payouts, which are certain. Where the intermediate cash flows are uncertain, however, the closed form Black-Scholes approach appears to break down and it is particularly difficult to compute American call prices.

The Black-Scholes model sets a lower limit for the price of the American call, but the model does not encompass the additional problem of determining the optimal time to exercise the option where the possibility of early exercise is not excluded. However, the model provides a reasonable approximation for the prices of American options.

The problems of both intermediate cash flows and early exercise can be solved with some adjustments to the standard models. These adjustments are examined below in detail.

Empirical tests of the Black-Scholes model indicate that the model is remarkably robust and provides accurate pricing for at-the-money options with medium to long maturity. The model appears to systematically misprice out-of-the-money and in-the-money options and also options where volatility increases or time to maturity is very short. In general, however, the model appears to successfully capture the essential determinants of option prices and United States studies show that traders cannot make consistent above normal returns on an after tax, post commission basis by setting up hedged portfolios.

Different models, which seek to overcome some deficiencies of Black-Scholes, essentially introduce new assumptions and do not necessarily produce improvements in pricing predictions.

The increased effort in improving and developing variations on available theoretical option pricing models creates the added problem of model selection. Clearly, there is no simple basis for selecting between the various techniques as the actual benefit from a particular model will depend on the user's objective. The selection of any model practically depends on the user's assumptions concerning the underlying asset price process. As there is no universal or true underlying process of asset prices, there can be no universal option pricing model and therefore no definitive fair value price for options.

In practice, models (such as Black-Scholes and the binomial price approaches) have remained successful because of their logical simplicity,

computational efficiency and their robustness. Market participants have sought to deal with the failure of model assumptions in real markets by a series of adjustments to the models, in practice, including:

- adjusting the volatility utilised for options with different maturities or different strikes to adjust for the undervaluation of in or out-of-the-money options; and
- increasing volatilities for shorter dates options to adjust for the potential impact of large jumps and consequent changes in the price of the underlying asset.

The strength of the option pricing models identified may ultimately lie in their capacity to compress the four observable variables into one other variable, implied volatility, which can then be interpreted. However, the problem of model pricing performance has led to option traders utilising a range of pricing techniques and risk management techniques to manage the risk of writing options.

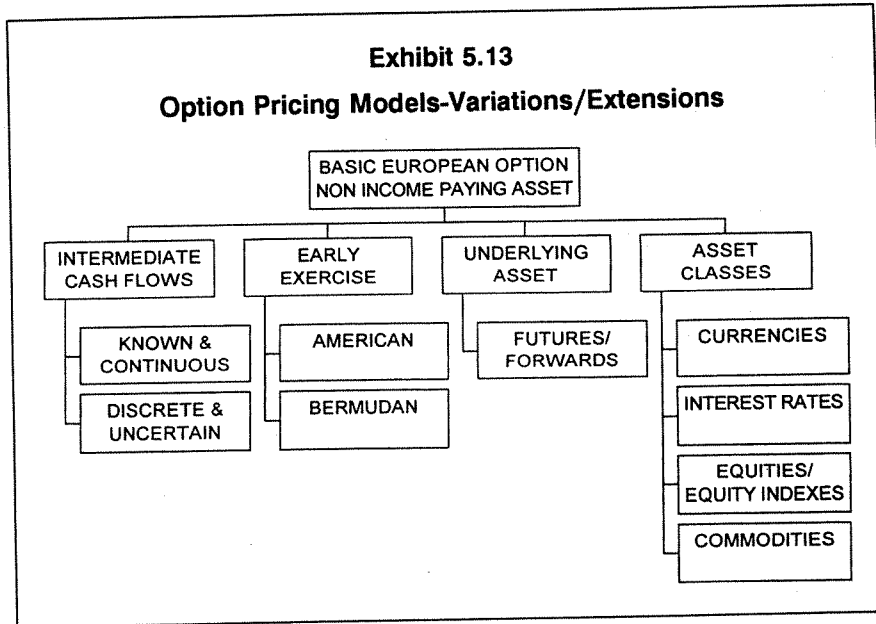
10. OPTIONS PRICING MODELS—EXTENSIONS TO THE STANDARD MODELS

10.1 Overview

In this section, the amended versions of the basic option pricing model are considered. The basic models are designed to deal with European options on non-income-producing assets. The extensions described extend the basic model in a number of specific areas:

- adjustment for intermediate cash flows;
- the possibility of early exercise;
- a change in the underlying asset to *a forward or futures contract* on the asset; and
- coverage of specific asset classes, such as equity, currency, debt or commodities.

Exhibit 5.13 sets out the hierarchy of option pricing models.



10.2 Intermediate cash flows

In practice, most assets pay out intermediate cash flows. These take the form of:

Asset Class	Income/Intermediate Cash Flow
Equities/equity market indexes	Dividends
Debt/interest rates	Coupons
Currencies	Interest rate on the foreign currency
Commodities/commodity Price Indexes	Asset lease rates or convenience yields

In practice, there are a number of mechanisms for incorporating the impact of these intermediate cash flows into the option pricing models:

- adjust the holding cost by the yield on the asset, that is, reducing the risk free rate (R_f) by the yield (Y), thereby using the holding cost of $R_f - Y$ in the model.
- adjusting the spot price of the asset by the income expected on the asset. This is done in one of two ways:
 - (a) Where the income is assumed to be continuous, by adjusting the asset by replacing S by the term $S e^{-Y.T}$ where Y = the continuously compounded expected rate of return on the asset. *Exhibit 5.14* sets out a valuation model adapting the standard Black-Scholes model for an asset which pays continuous income. An example of applying this type of model is set out below using a currency option where the risk free interest rate in the foreign currency (R_{f_f}) is used instead of Y in the formula in *Exhibits 5.19* and *5.20*.

- (b) Where the income is discrete and known, the spot price of the asset may be adjusted by discounting the known income to the commencement of the transaction and subtracting the discounted income from the spot price to derive an ex-income asset price which is then used as S in the model. *Exhibit 5.18* sets out an example of this type of adjustment with regard to an equity option.

Exhibit 5.14

Valuation Model for a European Option Asset Paying Continuous Income

For a call option:

$$P_{ce} = S e^{-Y.T} . N(d1) - K e^{-Rf.T} . N(d2)$$

Where

$$d1 = [\ln(S/K) + (Rf - Y + \sigma^2/2) T] / \sigma \sqrt{T}$$

$$d2 = [\ln(S/K) + (Rf - Y - \sigma^2/2) T] / \sigma \sqrt{T}$$

$$= d1 - \sigma \sqrt{T}$$

Where all terms are as defined previously and Y = the continuously compounded expected rate of return on the asset.

For a put option:

$$P_{pe} = K e^{-Rf.T} . N(-d2) - S e^{-Y.T} . N(-d1)$$

Where a binomial option pricing model is utilised to value the option, the value of the asset must be adjusted at the node at which the income is paid (or, in practice, the date at which the entitlement to the income is lost, such as an ex-dividend or ex-coupon date). The value of the asset is reduced by the amount of the income flow.

This creates a number of problems. The tree of asset prices becomes non-recombining, that is, an up move followed by a down move is no longer the same as a down move followed by an up move in the asset price. The tree also becomes larger as a result of the non-recombining nature of the tree. This is exacerbated where there are several cash flows. A common approach to improve the numerical efficiency of the solution of the binomial tree under these circumstances is to treat the asset as the asset price (ex-dividend), which is modelled through the tree and the present value of the future dividends which is added to the modelled asset price at each node.⁵

10.3 Early exercise/American options

The market for options trades both European and American options. As noted above, an American option will generally not be exercised early as the economic rationale favours the sale of the option. However, as noted above under certain circumstances, early exercise is possible.

5. See Hull, op cit n1, pp 365-368.

In practice, the risk of early exercise is particularly evident in the following cases:

- For in-the-money call options where the asset pays a high yield or income payment. This is because the benefit of receiving the high yield or the capture of the cash flow of the income payment on the asset may yield a superior return to the uncertain value of the asset and call at maturity. This is particularly relevant for equity options with a large dividend payable prior to option expiry where payment of the dividend would substantially reduce the value of the asset reducing the value of the call option. A similar logic applies to currency call options on a currency which has high interest rates whereby the option is likely to be exercised early.
- In the case of put options, a deep in-the-money put can economically be exercised early with the proceeds received invested at the risk free rate to yield a superior return to the uncertain intrinsic value of the put at maturity.

The valuation of American options is usually undertaken in two ways:

1. using a modified version of Black-Scholes option pricing model; and
2. using the binomial approach to option pricing.

The Modified Black-Scholes European Option Pricing Formula relies on the intuition that the standard Black-Scholes model provided a lower estimate of the value of the American option. It is identical to Black Scholes except that the formula checks to see if the value it is returning is below the intrinsic value of the option. Where the Black-Scholes European Option value is below the intrinsic price of the option then the Modified Black-Scholes American Formula returns the intrinsic value of the option. That is:

Black-Scholes American Option Value = Maximum (Black-Scholes European Value; Intrinsic Value)


The binomial option pricing model is well suited to estimating the fair value of an American exercise option. This reflects the fact that this approach incorporates all possible paths taken by the asset price as well as the distribution of asset prices as at the expiry of the option. This allows American options to be priced through a process whereby it is possible to calculate option values at each node of the tree and to test for the feasibility of early exercise. If the option at any node has a higher intrinsic value (that is, the value on early exercise) than the theoretical value of the option, the higher intrinsic value is used in the solution back through the tree, effectively incorporating the risk of early exercise.

Exhibit 5.15 sets out an example of using a binomial option pricing model to value an American put option.

Exhibit 5.15**Multi-Step Binomial Model for an American Put Option**

The use of a multi-step binomial option pricing model for an American put option can be illustrated with the example given in *Exhibit 5.12*. Assume all the factors stated in that example with the exception that the option is now an American put.

The binomial tree in that case was as follows:



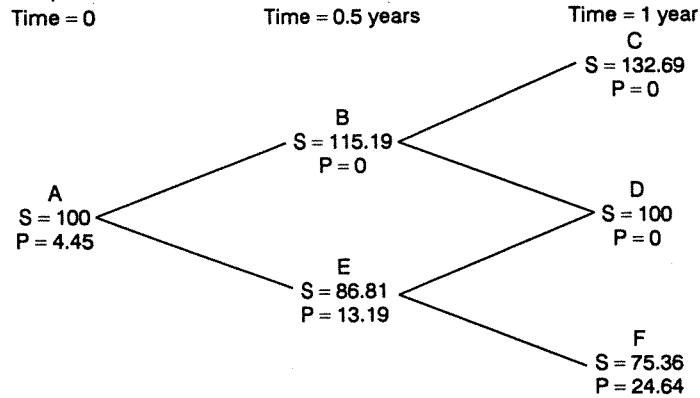
In solving back through the tree it is feasible to compare the value of a European put option with *the intrinsic value of the option* if the option is exercised at that node.

In this example, as at node E, the value of the European option is 8.31. However, if the option is exercised, it will have an intrinsic value of 13.19 ($100 - 86.81$). This means that the holder of the option would at this point rationally exercise the put option early.

In order to value the American put, the intrinsic value of the option at node E is substituted for the theoretical European value. This in turn allows calculation of the value of the put at commencement (node A) as:

$$P = e^{-10 \cdot 1/2} [0.645371(0) + (1 - 0.645371) 13.19] = 4.45$$

The complete tree therefore is as follows:



The difference between the value of the American put (4.45) and the value of the European put (2.80) of 1.65 can be attributed to the value of the right of early exercise.

An alternative approach is to use quadratic approximation methods to value American options. Under this approach, it is assumed that an American option is equal to a European option plus a separate early exercise option. The quadratic approximation method determines the early exercise option value and then adds it to the value calculated by the Modified Black-Scholes European Formula. The early exercise option value is determined by an iterative process. An example of this approach is that utilised by Barone-Adesi and Whalley.⁶

A variation on the American option is the Bermudan option, which is only capable of exercise on a nominated number of discrete dates prior to expiry. In effect, it is some way between an American and a European option. These types of options are priced using a binomial option pricing model which tests for the risk of early exercise at the relevant nodes of the tree.

10.4 Options on forward/futures contracts

Where the option is on a forward or futures contract, the basic Black-Scholes model can be altered to adjust for the changed nature of the underlying asset. *Exhibit 5.16* sets out Black's version of the original basic Black-Scholes model for a premium paid European option on a futures contract.

6. See G Barone-Adesi and R E Whaley, "Efficient Analytic Approximation Of American Option Values" (1987) 42 (June) *Journal of Finance* 301.

Exhibit 5.16**Black Option Pricing Model for Forward/Futures Contracts**

For a call option:

$$P_{ce} = e^{-Rf.T} [F \cdot N(d1) - K \cdot N(d2)]$$

Where

$$d1 = [\ln (F/K) + (\sigma^2/2) T] / \sigma \sqrt{T}$$

$$d2 = [\ln (F/K) - (\sigma^2/2) T] / \sigma \sqrt{T}$$

$$= d1 - \sigma \sqrt{T}$$

Where all terms are as defined previously and F = forward or future prices of the underlying asset.

For a put option:

$$P_{pe} = e^{-Rf.T} [K \cdot N(-d2) - S \cdot N(-d1)]$$

The intuition behind the Black reformulation of the Black-Scholes option pricing model in the context of futures is that the investment in a futures contract requires no commitment of funds (deposits, margins et cetera are ignored), whereas investment in the physical asset (for example, a share in the case of an equity option) imposes a cost. Consequently, nothing is paid or received (up-front) in setting up the hedge which entails buying or selling the futures contract.

The value of a call option on a futures contract should be lower than the value of a call option on the physical asset, as the futures price should already impound the carrying costs associated with the physical commodity. That is, the futures price is in essence the forward price, which will naturally reflect any carry costs over the relevant period. The usual qualifications concerning early exercise of American options will, of course, apply.

In general, the mean reverting property of interest rates causes these implied volatilities to decrease with option maturity. Relationships between volatility and maturity of caps, floors and collars et cetera are difficult to observe since most of these over-the-counter instruments have maturities beyond the longest maturity of traded Eurodollar futures. However, in practice, this relationship can be extrapolated.

Exhibit 5.17 sets out an example of utilising the Black option pricing model to derive the price of an option on a forward contract.

Exhibit 5.16**Black Option Pricing Model for Forward/Futures Contracts**

For a call option:

$$P_{ce} = e^{-R_f T} [F \cdot N(d1) - K \cdot N(d2)]$$

Where

$$d1 = [\ln (F / K) + (\sigma^2 / 2) T] / \sigma \sqrt{T}$$

$$d2 = [\ln (F / K) - (\sigma^2 / 2) T] / \sigma \sqrt{T}$$

$$= d1 - \sigma \sqrt{T}$$

Where all terms are as defined previously and F = forward or future prices of the underlying asset.

For a put option:

$$P_{pe} = e^{-R_f T} [K \cdot N(-d2) - S \cdot N(-d1)]$$

The intuition behind the Black reformulation of the Black-Scholes option pricing model in the context of futures is that the investment in a futures contract requires no commitment of funds (deposits, margins et cetera are ignored), whereas investment in the physical asset (for example, a share in the case of an equity option) imposes a cost. Consequently, nothing is paid or received (up-front) in setting up the hedge which entails buying or selling the futures contract.

The value of a call option on a futures contract should be lower than the value of a call option on the physical asset, as the futures price should already impound the carrying costs associated with the physical commodity. That is, the futures price is in essence the forward price, which will naturally reflect any carry costs over the relevant period. The usual qualifications concerning early exercise of American options will, of course, apply.

In general, the mean reverting property of interest rates causes these implied volatilities to decrease with option maturity. Relationships between volatility and maturity of caps, floors and collars et cetera are difficult to observe since most of these over-the-counter instruments have maturities beyond the longest maturity of traded Eurodollar futures. However, in practice, this relationship can be extrapolated.

Exhibit 5.17 sets out an example of utilising the Black option pricing model to derive the price of an option on a forward contract.

Exhibit 5.16**Black Option Pricing Model for Forward/Futures Contracts**

For a call option:

$$P_{ce} = e^{-Rf.T} [F \cdot N(d1) - K \cdot N(d2)]$$

Where

$$d1 = [\ln (F/ K) + (\sigma^2/ 2) T] / \sigma \sqrt{T}$$

$$d2 = [\ln (F/ K) - (\sigma^2/ 2) T] / \sigma \sqrt{T}$$

$$= d1 - \sigma \sqrt{T}$$

Where all terms are as defined previously and F = forward or future prices of the underlying asset.

For a put option:

$$P_{pe} = e^{-Rf.T} [K \cdot N(-d2) - S N(-d1)]$$

The intuition behind the Black reformulation of the Black-Scholes option pricing model in the context of futures is that the investment in a futures contract requires no commitment of funds (deposits, margins et cetera are ignored), whereas investment in the physical asset (for example, a share in the case of an equity option) imposes a cost. Consequently, nothing is paid or received (up-front) in setting up the hedge which entails buying or selling the futures contract.

The value of a call option on a futures contract should be lower than the value of a call option on the physical asset, as the futures price should already impound the carrying costs associated with the physical commodity. That is, the futures price is in essence the forward price, which will naturally reflect any carry costs over the relevant period. The usual qualifications concerning early exercise of American options will, of course, apply.

In general, the mean reverting property of interest rates causes these implied volatilities to decrease with option maturity. Relationships between volatility and maturity of caps, floors and collars et cetera are difficult to observe since most of these over-the-counter instruments have maturities beyond the longest maturity of traded Eurodollar futures. However, in practice, this relationship can be extrapolated.

Exhibit 5.17 sets out an example of utilising the Black option pricing model to derive the price of an option on a forward contract.

Exhibit 5.17**Example of Using the Black Option Pricing Model for Forward/Futures Contracts**

Assume the following information:

Forward Price = $F = 100$
 Strike Price = $K = 100$
 Time to Maturity = $T = 12$ months (1.00 yrs)
 Risk Free Rate = $R_f = 10\%$ pa (0.10)
 Volatility = $\sigma = 20\%$ pa (0.20)

Using the above inputs, we can compute the call option price as follows:

$d_1 = 0.10$
 $d_2 = -0.10$

Using the $N(x)$ table:

$N(d_1) = N(0.10) = 0.5398$
 $N(d_2) = N(-0.10) = 0.4602$
 $N(-d_1) = N(-0.10) = 0.4602$
 $N(-d_2) = N(0.10) = 0.5398$

Therefore, the call and put prices are as follows:

$$P_{ce} = e^{-0.10 \times 1.00} [100 \times 0.5398 - 100 \times 0.4602]$$

$$= 7.20 \text{ or } 7.20\% \text{ of Future Asset Value}$$

$$P_{pe} = e^{-0.10 \times 1.00} [100 \times 0.5398 - 100 \times 0.4602]$$

$$= 7.20 \text{ or } 7.20\% \text{ of Future Asset Value}$$

The computational method where the Black model is used to price options on futures contracts for the price of the relevant options will differ depending on the type of margining system applicable. When the margining system dictates that proceeds are not paid up-front to option writers prices have to be higher to compensate the seller for the fact that the premium is not received at the beginning and consequently it is not available for investment. If it is assumed initially that the premium is paid over to the seller of the option only at maturity, the premium would have to be increased by the additional interest that could have been earned over the life of the option if the premium was available for investment. Consequently, the value of the call option will become:

$$P_{ce} = F.N(d_1) - K.N(d_2)$$

In addition, the put-call parity relationship, where proceeds are not paid up-front for open futures contracts, is different:

$$P_{pe} = P_{ce} - F + K$$

In practice, the adjustment process is not simple, because if nothing changes, part of the premium will be paid over to the writer of the option as the time value decays to zero over the life of the option.

11. OPTIONS ON DIFFERENT ASSET CLASSES

11.1 Overview

To date, the focus on option valuation has been general, rather than focused on specific types of assets. In this section, adjustments to the option pricing model dictated by the *type* of underlying asset are examined. The amendments are specifically designed to encompass the particular characteristics of each asset class. The adjustments required for the possibility of early exercise, discussed above, apply uniformly to these cases where the option is American style, allowing the possibility of exercise prior to maturity. Similarly, the use of an amended model where the underlying asset is a futures contract on the asset, as described above, is also applicable to each of the asset classes described.

11.2 Options on equity/equity market index

Options where the underlying asset is an individual equity stock or equity market index require the model to be adjusted for the dividends paid on the underlying asset. The major issue with the income stream attaching to the asset is the uncertainty relating to future dividends both in terms of amounts and their exact timing.

In order to incorporate the potential impact of the income stream the two approaches described above are utilised:

1. Where the dividend income is assumed to be continuous, by adjusting the asset by replacing S by the term $S e^{-Y.T}$ where Y = the continuously compounded expected dividend rate on the asset. *Exhibit 5.14* sets out a valuation model adapting the standard Black-Scholes model for an asset which pays continuous income.
2. Where the dividend income is discrete and known, the spot price of the asset may be adjusted by discounting the known dividend to the commencement of the transaction and subtracting the discounted income from the spot price to derive an ex-income asset price which is then used as S in the model. *Exhibit 5.18* sets out an example of this type of adjustment with regard to an equity option.

Exhibit 5.18
Valuation of Equity Option

Assume a \$4.00 call option on ABC Ltd (ABC) with an expiry in four months' time is required to be priced. The underlying share price of is \$5.00. The risk free return to option expiry is 13.10%. The volatility of the underlying shares' rate of return is 40%. The last dividend paid by ABC Bank Ltd was \$0.30 per share. Assume that ABC is expected to pay a dividend of \$0.30 per share just prior to option maturity.

In order to value the option, it is necessary to adjust the spot price of the asset for the expected dividend. The present value of this dividend discounted at the risk-free rate of return is \$0.287. Therefore, the ex-dividend share price on \$4.713 (calculated \$5.00 - \$0.287).

The inputs into the option pricing model are:

$$S = 4.713$$

$$K = 4.00$$

$$\sigma = 0.40$$

$$R_f = .1310$$

$$T = 0.33$$

Using the standard Black-Scholes formula:

$$d_1 = 1.017$$

$$d_2 = 0.787$$

Therefore:

$$N(1.0177) = 0.846$$

$$N(.787) = 0.785$$

The call option value is determined as follows:

$$P_{ce} = 4.713 (0.846) - (0.9577) (4.00) (0.785) \\ = \$0.98$$

The Black-Scholes option-pricing model values the ABC option at 98 cents.

In practice, the first approach is used with longer dated equity options while the second approach is used with shorter dated equity options.

A second consideration in relation to equity options is the potential dilutionary impact of conversion. This exists as a problem in the case of equity options or warrants issued by the company. It is not a problem in exchange traded or over-the-counter options *on existing equity securities*. This is because it is only in the first case that the exercise of the option results in the issue of *additional equity*.

As the standard option pricing model assumes that the exercise of the option does not impact the value of the underlying asset, it is necessary to adjust the model for the following effects:

- the exercise of these options increases the number of shares which are on issue; and
- the payment of the exercise price of the option creates an additional cash inflow into the issuer of the options.

The standard model must usually be adjusted in two ways to reflect this impact:

1. The spot price of the asset must be adjusted for dilution as follows:⁷

$$S_d = (S \cdot Q_s + W \cdot Q_w) / Q_s$$

Where

S_d = the dilution adjusted spot price of the equity

S = the spot price of the equity

W = Market value of warrants

Q_s = Number of shares currently on issue

Q_w = Number of warrants on issue

2. The value of the theoretical call is adjusted to incorporate the dilution value as follows:

$$P_{ce}(\text{adj}) = P_{ce} \cdot [Q_s / (Q_s + Q_w)]$$

Where

$P_{ce}(\text{adj})$ = Dilution adjusted value of the call

P_{ce} = Call option premium

In practice, the expected dilutionary effect of exercise will make the premium on the warrants (call options) lower than for a standard call option.

11.3 Options on currency/foreign exchange

In theory, the Black model for pricing options on forward contracts is capable of being utilised to value currency options. However, in practice, the way several interest rates are involved in ways differing from the assumption of the Black-Scholes model dictate the use of different models.

Garman and Kohlhagen⁸ argue that it is the interest rate *differential between domestic and foreign risk free rates* that reflects the expected price drift of the underlying asset rather than a single interest rate that is appropriate in a single asset option.⁹ The Garman-Kohlhagen model is set out in *Exhibit 5.19*. *Exhibit 5.20* sets out an example of pricing a currency option using the model.

7. See Aswath Damodaran, *Damodaran on Valuation* (John Wiley, New York, 1995), pp 336-339.
8. Mark B Garman and Steven W Kohlhagen, "Foreign Currency Option Values" (1983) 2 *Journal of International Money and Finance* 231.
9. The Black-Scholes model employs interest rates in two different contexts: first, to discount future values; and, secondly, as an arbitrage-based surrogate for the drift of the deliverable instrument. The first use takes place outside the $N(\cdot)$ while the second takes place inside, reflecting the distribution of maturation values. Garman and Kohlhagen (*ibid*) show that it is only the interest rate differential that controls the distribution features while the interest rates control the discounting of future values.

Exhibit 5.19**Garman-Kohlhagen Model for Valuation of Currency Option**

For a call option:

$$P_{ce} = S e^{-R_f T} N(d1) - K e^{-R_d T} N(d2)$$

Where

$$d1 = [\ln(S/K) + (R_d - R_f + \sigma^2/2) T] / \sigma \sqrt{T}$$

$$d2 = d1 - \sigma \sqrt{T}$$

Where all terms are as defined previously and

R_d = the risk free interest rate in the domestic currency

R_f = the risk free interest rate in the foreign currency

For a put option:

$$P_{pe} = K e^{R_d T} N(-d2) - S e^{R_f T} N(-d1)$$

Exhibit 5.20
Valuation of a Currency Option

Assume the following facts relating to US\$/Yen:

$$S = 115.00$$

$$K = 110.00$$

$$T = 6 \text{ months (0.5 years)}$$

$$R_f = 3\% \text{ pa (0.03) (the yen interest rate)}$$

$$R_{fd} = 6\% \text{ pa (0.06) (the US\$ interest rate)}$$

$$\sigma = 12.00\% \text{ pa (0.12)}$$

The calculations are as follows:

$$d1 = 0.743$$

$$d2 = 0.658$$

Therefore:

$$N(0.743) = 0.771$$

$$N(0.658) = 0.745$$

$$N(-0.743) = 0.229$$

$$N(-0.658) = 0.255$$

The value of the call option is:

$$\begin{aligned} P_{ce} &= S e^{-R_f T} N(d1) - K e^{-R_{fd} T} N(d2) \\ &= 115. e^{-0.03 \times 0.5} \times 0.771 - 110. e^{-0.06 \times 0.5} \times 0.745 \\ &= 7.92 \end{aligned}$$

The fair value of the call in Yen is 7.92.

The value of the put option is:

$$\begin{aligned} P_{pe} &= K e^{-R_{fd} T} N(-d2) - S e^{-R_f T} N(-d1) \\ &= 110. e^{-0.06 \times 0.5} \times 0.255 - 115. e^{-0.03 \times 0.5} \times 0.229 \\ &= 1.28 \end{aligned}$$

The fair value of the put in Yen is 1.28.

The difference between the Black model and the Garman-Kohlhagen model¹⁰ gains importance where the difference between the two rates is very small or when one or other rate is large. The latter is particularly relevant in the case of American options because the high interest rate may influence the early exercise of the option.

10. Garman and Kohlhagen (ibid) also show that by substituting the forward currency price as at the option expiry (relying on the interest parity condition that the fully arbitrated forward rate should equal to the spot rate adjusted by the interest differential between the currencies) into the model it is possible to derive the Black option pricing model on forward/ futures contracts, thereby showing that currency options can be treated on the same basis as options on forwards generally.

11.4 Options on commodities

Options on commodities create issues in pricing primarily through difficulties in the estimation of the convenience yield/asset payout rate. This reflects the substantial difficulties in the estimation of this parameter. In practice, options on commodities may be priced in one of the following ways:

- Options on physical commodities using the continuous income version of the Black-Scholes model (see *Exhibit 5.14*) with the asset convenience yield being used as the income term Y .
- Options on forward commodities using the Black model for options on forward contracts (see *Exhibit 5.16*).

The latter approach has the advantage of already impounding the convenience yield or asset payout rate in the forward commodity price used in the model.

11.5 Options on debt instruments

11.5.1 Distinctive features of debt options

The pricing of options on interest rates and debt instruments are particularly complex and several distinctive features of debt instruments must be incorporated into the pricing of debt options.

The key features which require incorporation in the pricing mechanism include:

- Debt instruments, typically, have a defined maturity and their limited and declining life represents special problems in option pricing.
- The underlying security in the case of debt instruments usually involves payouts in the form of interest during the life of the option.
- The rate of interest cannot be assumed to be constant as, first, interest rate changes drive price changes in the underlying asset, and, secondly, most interest rate security values do not depend on a single random variable but on a number of random interest rates.
- Volatility of the underlying debt instrument cannot be assumed to be constant.

Most debt securities have a defined maturity. This is in contrast to other assets, such as equities, currencies and commodities, which do not have fixed lives.

It is important to distinguish between two classes of options on debt instruments: options on the cash market debt instrument (that is, the actual physical debt security); and futures on the relevant debt instrument. In practice, both types of options coexist and are available. This is despite the fact that in any market, a cash market, a futures market and one options market (either on the cash market instrument or the futures contract) would usually be sufficient to fulfil all risk transfer possibilities since the option on the cash market and the option on the futures market will, generally, serve similar functions.

In practice, options on cash market or physical debt instruments takes one of two forms: *fixed deliverable*, whereby a debt instrument *with specified characteristics* is required to be delivered; or *variable deliverable*, whereby a *specified existing debt issue* is required to be delivered. For example, a six month call option on a 90 day Eurodollar deposit (90 days commencing from the expiry date of the option) is a fixed deliverable option. In contrast, a three month put option on the 8.5% August 2002 US treasury bond, which requires delivery on that specific security—irrespective of remaining term to maturity—is an example of a variable deliverable option.

Variable deliverable options create complex pricing issues. This reflects the fact that unlike other cash market assets which have infinite lives or futures contracts which are not based on a particular, wasting debt security (futures contracts have particular characteristics which are specified and constant), actual physical debt securities are affected by the passage of time in two respects:

1. The underlying debt instrument itself has a shorter tenor or period to maturity as the option itself approaches expiration.
2. At maturity, the value of the interest rate security converges to a known constant value (usually, par or face value) and the volatility of the security approaches zero.

The impact of intermediate cash flows on the underlying debt instrument will depend on whether the underlying asset for the debt option is a cash market debt security or a futures contract on the relevant instrument. Where the option is on a futures contract of the relevant debt instrument, the fact that there are no coupon interest payments and that the maturity of the particular debt securities is fixed, means that the general pricing technology applicable to options on forward/futures contracts (see discussion above) can be utilised. However, where the option is on the physical debt security, the presence of intermediate cash flows can be problematic.

The assumption made by models such as Black and Scholes that there are no intermediate cash payouts can be relaxed using a modification of the formula which allows for payments that are proportional to the price of the underlying security. However, the normal type of adjustment utilised may not be appropriate in the case of debt options. Where the option is on an underlying security which bears a coupon, the accrued interest is continuously added to the full price of the bond representing a continuous payout to the holder of the debt security. As the coupons are fixed in dollar amount not proportional to the price of the underlying debt security, this type of modification proposed would be inappropriate.

As noted above, basic option pricing models assume that only one interest rate, the risk-free rate, is relevant. However, at any given point in time, a variety of risk-free interest rates for different maturities are observable. Each of these interest rates and, consequently, the shape of the yield curve as a whole is subject to change over time.

A major difficulty in relation to the pricing of debt options relates to the fact that the price of the underlying asset (the debt security) itself is a function of interest rates. Moreover, it is unlikely depending on the type of option, that it is a function of the risk-free interest rate utilised to present

value the exercise price of the option. An additional complication arises from the fact that where the option is a variable deliverable option (as defined above), the exact interest rate required to value the underlying debt instrument itself is subject to change with the passage of time.

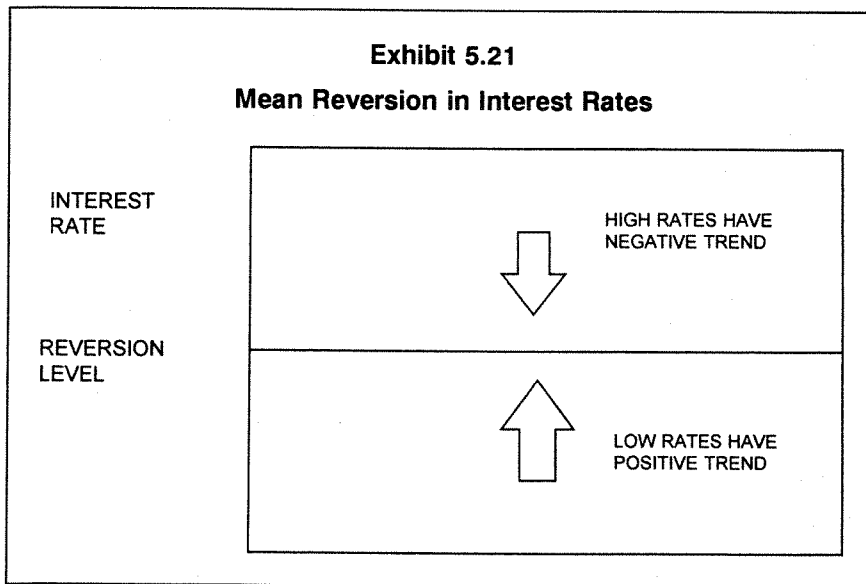
These difficulties mean that the value of options on debt instruments or interest rates do not depend on a single random interest rate variable but may depend on a *number* of different random interest rates (in effect, the complete term structure of interest rates).

The effect of changes in interest rates and the time to expiration are particularly complex. For options on assets, such as shares, as the risk-free rate increases, the value of the call option increases as the present value of the exercise price in the event of exercise declines; that is, if the call option and the security itself are regarded as different ways for an investor to capture any gain on the security price, as rates rise the increased cost of carry on the underlying security will make the call more attractive, leading to an increase in its value.


However, in the case of debt options, it is unreasonable to assume (as is usually done in the case of equity options) that the price of the underlying debt security is independent of the level of interest rates. Significant movements in the prices of the asset will occur as a result of changes in interest rates and, in general, any cost of carry consideration would be minor relative to the change in the value of the underlying security. For example, it would be reasonable to assume that rate increases will usually have a negative impact on the price of call options on debt instruments, as a rise in interest rates will most likely cause a fall in the price of the underlying instrument or futures contract.

The constancy of the variance of volatility of the underlying debt instrument also cannot be assumed. This results from two factors:

- The stochastic process followed by interest rates appears to have a mean reversion quality—that is, there is an inbuilt drift that pulls them back to some long run average level (see *Exhibit 5.21*).



- Volatility of debt securities (in the case of a variable delivery option) is likely to tend to zero, reflecting the fact that at maturity, the value of debt instrument itself must converge to a known value (usually, the par value of the security). An additional factor in this regard is that the price volatility of a security is itself a complicated function of the actual volatility of interest rates of varying maturity and the time to maturity at the security itself (see *Exhibit 5.22*).



Consequently, the volatility of debt instruments will generally be a function of the assumed stochastic process of interest rate movement, assumptions about the shape and future movement of interest rate across the whole yield curve, and the remaining life of the security at a given point in time. These complexities dictate that constant variance cannot be assumed and it is probable that the volatility, itself, may also be stochastic variable.

The complexity of these interactions can be illustrated with reference to the effect of changes in time to expiration on such options. For options on assets with unlimited lives, an option with a longer time expiration will, generally, be worth more than a comparable option with a short term to expiry on the basis that it has all the attributes of the shorter dated option plus more benefits for the holder, that is, there is greater probability that the option can be profitably exercised. This property need not necessarily hold for debt options, particularly variable deliverable options, as depending on the relative magnitude of the time value and the intrinsic value, it is conceivable that under certain circumstances, an option with a longer time to expiration may be worth less than one with a shorter term. This would reflect the fact that securities usually begin to trade closer to par as the instrument approaches maturity. The greater price stability may affect the value of the option.

11.5.2 Approaches to pricing debt and interest rate options

In practice, the pricing of debt and interest rate options fall into two categories:

1. Options on futures/forwards where the underlying asset is a standardised debt instrument—these types of options are valued in a manner consistent with the types of theoretical option pricing models outlined, in particular, the Black option pricing model, as some of the problems identified can be minimised.
2. Options on physical debt instruments (particularly, on physical bonds)—are more problematic and usually entail the use of various numerical, usually binomial or lattice, option pricing model.

The first type of approach is the one commonly used with pricing caps, floors and collars as well as options on swaps/swaptions (although Bermudan style swaptions or those where the final maturity of the underlying swap is fixed at the time of entry may be more akin to the problems of pricing options on physical debt instruments).

For example, pricing caps, floors and collars utilising the first approach (using the Black model for options on forward/futures rates) entails the following steps:

- The cap, floor or collar agreement is analytically separated into a series of option contracts. For example, an interest rate cap agreement may be split up into a series of put option contracts on the prices of short term debt securities pricing off the relevant interest rate index.
- Each separate option is then valued utilising the identified model. In determining the price of each option, it is important to note that the input for the current spot price of the index is not the physical market

price at the time the agreement is entered but the then current futures or forward price on the relevant index.

- The option premium for each contract is calculated and then summed to give the actual price for the overall contract.

Exhibit 5.23 sets out an example of using this approach to price a single period cap.

Exhibit 5.23

Example of Calculating Cap/Floor Prices Using Black's Option Pricing Model

Yield approach

Calculate the premium for a \$1m 15% cap on three months LIBOR for one period of three months commencing in three months' time where the forward rate for three month LIBOR in three months' time is 15.016% pa and the three month risk free rate is 15.00% pa. This information can be reformulated for input into the model as follows:

$$F = 0.15016$$

$$K = 0.1500$$

$$T = 0.25$$

$$R_f = 0.15$$

$$\sigma = 0.17$$

The price of the cap (call option on yield) can be calculated as follows:

$$d_1 = 0.055$$

$$d_2 = -0.030$$

Therefore:

$$N(0.055) = 0.5219$$

$$N(-0.03) = 0.4880$$

The price of the call is calculated as follows:

$$P_{ce} = e^{-0.15 \times .25} [0.15016 \times 0.5219 - .15 \times 0.4880] \\ = 0.004978$$

As the asset and strike price were specified in yield terms, it is necessary to restate the option premium as follows:

$$P_{ce} = [t.Fv / (1 + F.t)] \times P_{ce}$$

Where

t = the interest rate period (.25 years)

FV = face value of the option (1,000,000)

Therefore, the value of the cap is as follows:

$$P_{ce} = [.25 \times 1,000,000 / (1 + .15016 \times .25)] \times .004978 \\ = \$1,199.54 \text{ or } 0.1199\% \text{ of face value}$$

The price of the equivalent floor (put on interest rates) is as follows:

$$N(-0.055) = 0.4781$$

$$N(-0.03) = 0.5120$$

$$P_{pe} = 240,954.57 \times e^{-0.15 \times 0.25} [0.15 \times 0.5120 - 0.15016 \times 0.4781] \\ = \$1,162.40 \text{ or } 0.1162\% \text{ of FV}$$

Exhibit 5.23—continued*Price approach*

Calculate the premium for an 8% cap on three month LIBOR for one year period of three months in one years' time where the forward rate for three month LIBOR in one year is 7% pa. This information can be reformulated for input into the model as follows:

F = 982,800.98 (the value of US\$1,000,000 face value 3 month security discounted at 7.00% pa)

K = 980,392.16 (the value of US\$1,000,000 face value 3 month security discounted at 8.00% pa)

T = one year (1.0)

RF = 6.5% (0.065)

$\sigma = 0.00344$

The price of the cap is calculated as follows:

d1 = 0.715

d2 = 0.712

Therefore:

$N(0.715) = 0.7627$

$N(0.712) = 0.7616$

$N(-0.715) = 0.2373$

$N(-0.712) = 0.2384$

Therefore, the cap premium (put option on price) is:

$$P_{pe} = e^{-0.065 \times 1} [980,392.16 (0.2384) - 982,800.98 (0.2373)] \\ = \$474.92 \text{ or } 0.0475\% \text{ of FV}$$

The Black option pricing model involves the assumption that σ —the volatility of the forward/futures contract—is constant. As noted above, this assumption is unlikely to hold in practice because of the mean reverting process. This dictates that when the period to option expiration is large, the price of the forward or futures contract is not greatly sensitive to current interest rates, but as the time to option expiry decreases, the current level of interest rates becomes progressively more important in determining the forward or futures price with the result that the volatility of the forward or futures price may increase with the effluxion of time.

In practice, the Black option pricing model can be applied if some adjustments designed to minimise impact of this phenomenon are adopted.

Utilising this approach, applied volatility for forward interest rates are calculated usually from traded futures options or from traded caps, floors, et cetera. The debt option being valued is then priced utilising the implied volatilities generated. This multiple use of the Black model, first, to calculate implied volatility and, secondly, to price the option, allows errors that may be caused by the use of the inexact model to be reduced. More importantly, they ensure that the calculated option prices are reasonably consistent with traded option prices.

In general, the mean reverting property of interest rates causes these implied volatilities to decrease with option maturity. Relationship between volatility and maturity of caps, floors and collars, et cetera, are difficult to observe since most of these over-the-counter instruments have maturities

beyond the longest maturity of traded Eurodollar futures. However, in practice, this relationship can be extrapolated.

The major identified problems with pricing options on interest rates or debt instruments are most evident in pricing options on physical debt instruments, particularly medium to long-term bonds. A variety of approaches have emerged towards pricing these types of options usually incorporating an interest rate term structure model. Chapter 6 examines interest rate option pricing using these types of approaches.

12. SUMMARY

Option contracts, because of their asymmetric payoff profiles, present a particular challenge in pricing. However, using the standard assumption of risk neutral valuation, it is possible to estimate the fair value of the option contract by determining the expected value of a portfolio consisting of the underlying asset and cash which is adjusted dynamically through time to replicate the payoff of the option. The basic model thus derived can then be adjusted in a number of ways to estimate the fair value of different types of options as well as options on different asset classes.

Appendix A

Cumulative Normal Distribution Function¹¹

Tables for the cumulative normal distribution function (N) are attached.

Alternatively, a polynomial approximation can be used:

Where $x \geq 0$

$$N(x) = 1 - N'(x) (a_1 k^1 + a_2 k^2 + a_3 k^3)$$

Where $x < 0$

$$N(-x) = 1 - N(x)$$

Where

$$k = 1 / (1 + \alpha x)$$

$$\alpha = 0.33267$$

$$a_1 = 0.4361836$$

$$a_2 = -0.1201676$$

$$a_3 = 0.9372980$$

and

$$N'(x) = (1 / \sqrt{2\pi}) \cdot e^{-x^2/2}$$

This provides values for N(x) that are usually accurate to about four decimal places and are always accurate to within 0.0002.

11. See M Abramowitz and I Stegun, *Handbook of Mathematical Functions* (9th ed, Dover Publications, New York, 1972).

Chapter 6

Interest Rate Option Pricing Models

by Tim Rowlands

1. INTRODUCTION

With the publication of their benchmark paper, Black and Scholes (BS) provided a means for valuing options on a wide range of financial instruments. Their work gave us an analytic framework that has since been put to many and varied uses. Many extensions and enhancements of their original model have been made, but the basic approach has remained intact. However, when the underlying instrument on which the option is granted or purchased is an interest rate or interest rate product, there is considerable strain put on some of the key BS assumptions.

First, there is the assumption that the rate used to discount the forward option price back to today is constant whilst the interest rates, upon which the underlying depends for its value, are evolving randomly. This represents an inconsistency which is magnified when the life of the option is a significant proportion of the length of maturity of the underlying instrument, for example a three year option on a one year swap. In effect, using this example, this assumption is saying that the forward rate from three to four years is a stochastic (random) quantity yet the spot three year rate is fixed. In other words the four year rate is stochastic while the three year rate isn't. Clearly as the option life grows as a proportion of the overall combined option plus underlying maturity this becomes less palatable. It is possible to imagine circumstances where a three month rate remained relatively constant whilst a ten year rate evolved randomly (as in a three month option on a ten year bond) but a little harder to imagine a three year rate remaining constant whilst a four year rate is expressed as random with its full normal volatility. Stochastic interest rate enhancements to the basic BS model have been used to overcome this in the case of stock options but they become much more difficult for interest rate options.

Secondly, there is the assumption that the annualised volatility of the underlying interest rate is constant for the life of the option. This is a difficult assumption to work with, even in simple options, but in the case of interest rate options there is a well-defined volatility term structure observed in the markets that clearly contravenes this constant volatility assumption. Stochastic volatility models and models involving deterministic, but not constant, volatility have been developed in an attempt to overcome this deficiency.

Thirdly, there is the assumption of an interest rate or bond price process that behaves in a Brownian motion fashion with the consequent growth of the total standard deviation of the underlying rate to be related to the square root of time. This implies that the distribution of rates or prices will spread to infinity as time goes to infinity, thus if you wait long enough an interest rate

of any size (for example, 100%, 1000%, even 1000000%) is possible. This is contrary to experience, as there are political and economic forces at work to keep interest rates at “reasonable” levels. This is referred to as “mean reversion”—the process whereby there is a tendency for rates and prices to return to their long-term average levels or average rates of growth that is stronger the further from this long-term average they have strayed. This is a particularly complicated addition to the interest rate modelling process since it links the change of the interest rate over a period of time explicitly with its level, and this is determined by the dynamics of the recent past. In other words, we have to keep track of the path of interest rates because we need to know where it is to determine where it is going. This causes problems in that it makes analytic option pricing solutions tough and numerical solutions typically become much more computationally intensive (for example, “bushy” trees and lattices instead of clean connected ones). This dependence of movements in future rates or prices on the past is described as the process being non-Markovian. When the rates or prices are independent of the past (we don’t need to know where we are to predict how we will move in the future), this is described as Markovian.

Numerous attempts to overcome some or all of these problems have been made, and are continuing to be made, so that the modelling of interest rate processes and interest rate option pricing remains a rich and interesting field of study. In the sections below, we will quickly review the types of models in use, we will look at interest rate processes at a conceptual level and then use these as a starting block to unravel the various different modelling approaches used for interest rate options. Using a binomial lattice for illustration, we will explore what is required to implement an interest rate option pricing model.

2. FAMILIES OF INTEREST RATE MODELS

Although there are many different interest rate models, they can be grouped into various families of models that allow easy comparison at a macro level. The main split in the early chronology of models was into equilibrium models and arbitrage free models. This broad classification is still relevant, but there has been a shift towards arbitrage free models in recent years so that most models these days fit into this category. The fundamental difference between the two categories is that equilibrium models begin with a model for the economy as a whole and derive a suitable process for the time value of money based on estimated economic utility functions for the participants in the economy. From this, interest rates can be modelled and the value of future fixed, floating and contingent cashflows can be determined. The best known of these models is the Cox, Ingersoll, Ross model.

Arbitrage free models begin with the choice of an interest rate process and then development of a description of rates into the future, based on the need to conform with a no arbitrage condition, that is, there should be no self-financing strategy which generates a risk-free profit. This arbitrage-free state is the equivalent of the risk-neutral valuation of the traditional BS model. Consequently, the probability distribution for rates associated with these models is artificially chosen to fit the arbitrage-free condition and

current interest rates and is thus not related to that which would be forecast by the market and is generally known as the risk-neutral probability distribution.

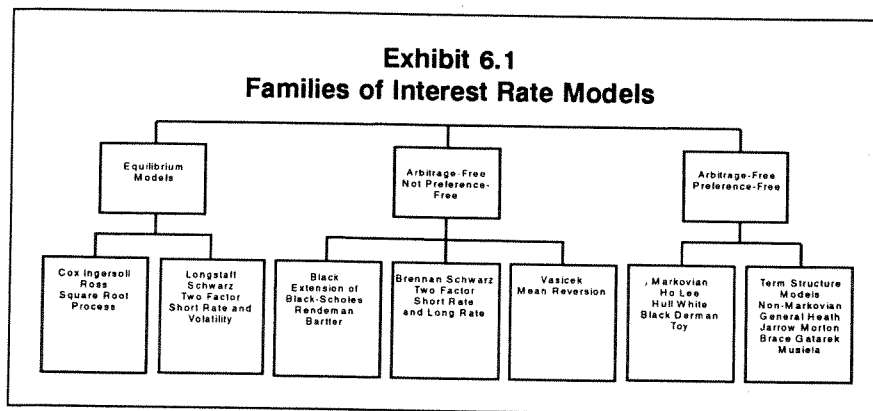
The next major division in the family of arbitrage free models is into models which are and are not preference free. Construction of an arbitrage free model (beginning with a simple expression of the interest rate process in a manner similar to the equilibrium models) enables the building of a term structure of interest rates but this will be unlikely to be consistent with the observed market-term structure. This is because some sort of trend will be assumed for the short-term rate and this will prove to be too simplistic to describe the subtleties seen across the whole term structure. In addition, the market observed short rate will not match the "risk-free" rate of the kind used in the BS stock option since this is a value representing the return on a money-market account rather than a bond return. In order to establish consistency with the observed term structure, an estimate must be made for the market price of risk. This represents the premium over the "risk-free" rate that must be returned to match the observed term structure. Unfortunately, this is not a market observable depending basically on the preferences or risk appetite of investors in the economy. In the BS model, this is removed by building a replicating portfolio of a combination of the money market account (growing at the risk-free rate) and the underlying stock. This can't be done for interest rates because the risk-free rate is a tradeable part of the term structure of interest rates. To avoid the problems associated with trying to estimate the market price of risk, there have been developed models which specifically fit the observed term structure of interest rates in an arbitrage-free manner. These models are described as "preference free" since they don't require an estimate of the market price of risk. Unfortunately, this fitting process is done numerically and so there is a loss of analytic tractability. This has led to the preponderance of numerical models such as lattices to accomplish this.

There are also divisions between the models on the basis of the nature of the underlying random process. Some models use a lognormally distributed random process for the interest rates whereas others use a normally distributed random process. A few models use other processes such as a square root process or Bessel process but these are usually for specific analytic solutions in cases where a normal or lognormal distribution doesn't yield an analytic result. Other models have been based on a random description of the bond prices rather than the underlying rates, but these have to be radically adjusted in cases where the maturity of the bond is approached since at this time the price of the bond homes in on the face value and thus does not follow a square-root-of-time volatility process.

Other divisions within the family of interest rate models are based on the number of factors used in the construction of the term structure of interest rates. Many models have been developed which use a single factor, others are multi-factor. Single factor models are based on a single random variable in the interest rate process. This usually corresponds to a short-term interest rate. From this single source of uncertainty the whole term structure is derived, and the manner in which it evolves determined. This means that all other elements depend in some way on this single factor. An analogy can be

borrowed from hedging or risk management. Typically in the past, interest rate positions, particularly bond portfolios, had their risk assessed in terms of their “duration”. This was a single number that represented the average maturity of the portfolio and, using a basis point sensitivity, gave an indicator of the riskiness of the portfolio. This was based on a parallel shift of the yield curve and its impact on the portfolio. However, it has become clear that although this approach is a reasonable estimate, it is incomplete, because it takes no account of the different behaviour of different parts of the yield curve. This is because there isn’t perfect correlation in the movements of interest rates at different maturities on the yield curve—sometimes rates move in opposite directions. Thus, a parallel shift model for a yield curve gives an incomplete picture of the risk of a bond portfolio with bonds of different maturities. In the same way, single factor interest rate models relate all of the interest rates to a single short rate. Though different models build the term structure in different ways, the net result is that the rates are all perfectly correlated across maturities. To move beyond this, additional factors are required and these are based on additional random variables. For example, a two factor model may incorporate random variables (with some, but not perfect, correlation) to represent a short rate and a long-term interest rate (for example, Brennan and Schwarz).

The final family of models is the relatively new generation of yield curve models in which the whole term structure of interest rates is taken as the random variable, for example, Heath, Jarrow and Morton. In the single factor case, rather than a single interest rate varying, a vector of rates (the yield curve across the maturities) is varying, but based on a single random variable. Since the underlying “variable” is the whole yield curve, interest rate instruments can be valued directly at a future time by discounting back the cashflows down the forward curve. In order to preserve an arbitrage-free state in these cases there is generally a loss of the Markovian property. Alternatively, there can be a return to the Markovian nature of the process if specific assumptions are made about the form of the volatility. Generally, this reduces the model to one of the simpler cases, anyway.¹



1. For a more detailed family tree of option pricing models see Smithson (two articles in *Risk* magazine).

3. THE INTEREST RATE PROCESS

The fundamental level of differentiation of interest rate models is the underlying interest rate process. A Brownian motion or Wiener process or random walk all refer to a process by which a quantity such as an interest rate evolves through time in a manner which results in a normal distribution. There will be some drift or trend which is proportional to the length of time that passes and there will be a gradual spreading of the distribution in such a way that the standard deviation of the distribution is proportional to the square root of the length of time that passes. In simplest terms, all of this is represented in the following equation

$$dr = \mu dt + \sigma dz$$

where dr is the change in rates over a short time period dt . The drift coefficient is μ and the spread coefficient is σ , and dz represents a small change in a normal random variable with mean zero and standard deviation equal to the square root of dt .

One of the key elements to understanding different models is being able to interpret this type of stochastic (because it includes a random variable) differential (because it is expressed in terms of small changes or differences) equation (SDE). Since dz is normally distributed and dr is proportional to it, then we know that dr is also normally distributed. The spread of the probability distribution is determined by the constant σ which is usually referred to as the volatility of the rate. This value can be estimated from historic data as the standard deviation of the *absolute changes in rates* over a period containing a number of observations, for example, three months of daily rate movement data. As dz has a mean of zero, the average value for dr over a large number of observations is determined only by the μdt term. If μ is a constant, as in this simple case, then there is a steady linear trend in the average or expected rate changes. This is common for assets that return a steady long-term growth rate (for example, house prices, shares, et cetera) but it is not so appropriate for interest rates, as it means that the rates would continue along the linear upward or downward trend forever, thus reaching very large positive or negative values, given sufficient time. So, in short, the interpretation of a SDE as written above with constant drift and volatility coefficients would be that rates evolve following a normal distribution around a linear trend.

In the above example, if μ was negative then eventually rates would become negative and continue becoming more and more negative as time went on. Alternatively, a large value of σ would allow the random term to overwhelm the drift and again allow (with small but nevertheless finite probability) negative interest rates. Notwithstanding a couple of historic examples of their occurrence, it is generally accepted that negative rates are not a desirable feature in an interest rate process. The use of a lognormal distribution for the changes in rates puts a lower bound of zero on the interest rates. In processes of this type it is the logarithm of the rate that is normally distributed. In this case the SDE could be written as:

$$d \ln(r) = \mu dt + \sigma dz$$

where $\ln(r)$ is the natural (base e) logarithm. An alternative representation of this is to write:

$$dr = \mu r dt + \sigma r dz$$

This is equivalent after a slight change in the drift term (see below). Note that the volatility term contains the rate to the power one. In simple terms, when the volatility term doesn't contain r , the distribution is Normal and if it contains r to the power one it is Lognormal. The Cox, Ingersoll, Ross model has r to the power of a half, hence its description as a square root process. Strictly, the $d \ln(r)$ form of the dr expression above is:

$$d \ln(r) = \left(\mu - \frac{1}{2} \sigma^2\right) dt + \sigma dz$$

This can be shown using Ito's Lemma which is an important result in financial modelling. (A simpler, though not rigorous, derivation can be obtained by taking a Taylor's series expansion of the derivative of $\ln(r)$ and noting that dz^2 is equivalent to dt .)

Once again the spread of the distribution is determined by the value of σ , but in this case it can be estimated from historic data by taking the standard deviation of the logarithm of the ratio of rates from one observation to the next. This is true since the equivalent absolute change in the logarithm of the rate from one day to the next is the logarithm of one rate minus the logarithm of the previous rate, and the difference of two logarithms is equal to the logarithm of the ratio of the rates. This is the origin of the seemingly strange definition of historic volatility as the standard deviation of the log-ratio of rates.

If $\ln(r)$ trends down to be a negative number (that is, μ is negative), then this means the rate will still be positive as the exponential of a negative number is positive and lies between zero and one. It is still unbounded on the positive side so that very large rates of interest can be obtained. The BS model uses a lognormal distribution for stock prices since these can grow to large values with time and are always positive. Once again, this is less appropriate for interest rates, as unbridled growth upwards in rates is not likely, due to political and economic factors coming to bear so that they remain at "reasonable" levels.

This effect of rates returning to some long-term average value rather than growing to higher and higher levels is referred to as "mean reversion". This can be built into a SDE as follows:

$$dr = k(\theta - r)dt + \sigma dz$$

where θ is the long-term average value of r and k is a measure of how strong the mean reversion "force" is. As can be seen, if the level of r is above the long-term average, θ , then the drift term will be negative and the rate will trend down. If the level of r is below the long-term average then the drift term will be positive and the rate will trend up. The further the rate is from the long-term average the stronger this drift will be. This can also be applied to a lognormal process leading to a SDE of:

$$d \ln(r) = k[\theta - \ln(r)]dt + \sigma dz$$

The difficulty with the last two SDEs is that the inclusion of the rate or its logarithm in the drift term means that the change in rate is dependent on where the rate is currently, and thus there is a loss of the Markovian property (that is, the size of the next jump is dependent on where we are, rather than

being independent, which is the Markovian case). This leads to bushy lattices and difficult analytical solutions to the SDEs.

Thus far we have assumed that the coefficients in the SDEs have been constants but they can be chosen to be time dependent or, just as in the case of mean reversion, rate dependent. We have also only considered r to be a single short-term rate. If we wish to proceed to a multi-factor model so as to explain a greater range of possible term structures and term structure movements there are several different ways we can do so.

One alternative is to use a pair of rates—a short-term and a long-term rate. These can be represented by a pair of related SDEs as follows:

$$\begin{aligned} dr &= k(l - r)dt + \sigma_1 r dw \\ dl &= \sigma_2 ldz \end{aligned}$$

and where

$$dwdz = \rho dt$$

In this case (the model of Brennan and Schwartz), the short rate is a mean reverting lognormally distributed rate that has as long-term average a long-term rate that is also lognormally distributed. The two rates are correlated with correlation coefficient ρ .

There are many examples of models in which the coefficients are time dependent and this allows values to be selected to ensure that the evolving rates fit the market observed term structure of interest rates. What this means is that the average or expected value of the rates at forward dates are consistent with the forward rates observed today. Jamshidian, Hull and White and others have produced models of this type. The SDE in these cases is typically of the form:

$$dr = (\theta(t) - ar)dt + \sigma dz$$

The fitting of the rate process to the current term structure of interest rates and volatilities is a process known as calibration. This will be discussed in more detail in the next section.

The final type of interest process is that used in the yield curve models where the rate itself is a function of maturity rather than a single rate for a single maturity. Heath, Jarrow, Morton and Brace, Musiela, Gatarek are examples of this type of model. The forward rates are modelled in an arbitrage-free manner and calibrated to fit the existing term structures of interest rates and volatilities. The SDE looks like this:

$$df(t, T) = \alpha(t, T)dt + \sigma_f(t, T)dW(t)$$

where the volatility function, σ_f , is the forward volatility function. Models of this type are quite computationally intensive in their purest form as they are non-Markovian leading to bushy lattices. It is possible to make simplifying assumptions about the volatilities which reduce the model to simpler forms. In this way, many of the older and simpler models are special cases of the HJM approach. These are the state-of-the-art models and are likely to be the way of the future for pricing and risk management of interest rate derivative products.

How do we choose an appropriate interest rate process? This is very much a case of personal preference combined with an assessment of the complexity

of the uses for which it is to be implemented. The running of a complex book of derivative products by an investment bank is very different to a small corporate treasury that might purchase the occasional cap or floor. Clearly the former will need a high quality implementation of a yield curve model, whereas the latter can get by with one of the simpler BS style pricing models, provided a conceptual understanding of the risks inherent in option products is well understood.

What do we do with it? Once an interest process has been selected we can simulate the evolution of the rate going forward. This can be done in a variety of ways, including binomial or higher order lattices, Monte Carlo simulation or, in some cases, analytically, because the SDE is soluble algebraically as in the BS model.

To conclude the discussion of interest rate processes it is necessary to summarise what we have seen. An interest rate process is a description of the way in which rates evolve into the future. These processes can be described by a Stochastic Differential Equation (SDE). By examining the SDE of a model, or designing one of our own, we can determine the features we want to build into a model. These features include the type of probability distribution, whether or not we will try to fit the current market observables exactly, whether we will include mean reversion and so on. All of this is fundamental to the building of any model and in some senses from then on the rest is just algebra and, often, some numerical mathematics!

4. A BINOMIAL LATTICE ILLUSTRATION OF AN INTEREST RATE MODEL

Binomial (and higher order) lattices provide a means of modelling stochastic processes in a controlled manner and were first applied to option pricing by Cox, Ross and Rubinstein. An alternative approach, Monte Carlo simulation requires large numbers of trials to ensure that there has been a representative coverage of possible outcomes in a random process, whereas the lattice approach has a wide and representative coverage at each stage, and extending the number of steps in the lattice simply increases the resolution of the random process. It also allows for the pricing of American options since it provides a means of stepping back through time from expiry to determine whether it is better to exercise or to hold onto the option.

Let us work through an example, using a lattice model to price a caplet. Extension onto a cap is straightforward as it is merely the sum of a series of caplets. The caplet has an expiry two years hence and is based on a three month interest rate. For simplicity we will use just eight steps in the lattice, each representing one three month period. The interest rate process will be a simple lognormal one with the drift term being time-dependent to allow for the fitting of the current term structure in an arbitrage-free manner and with the volatility term being constant.

$$dr = \mu(t)rdt + \sigma rdz$$

As we stated, above, this is equivalent to:

$$d \ln(r) = (\mu(t) - \frac{1}{2}\sigma^2)dt + \sigma dz$$

Integrating both sides (ignoring for the moment the time-dependence of μ) gives:

$$\ln(r_T) - \ln(r_0) = (\mu(T) - \frac{1}{2}\sigma^2)T + \sigma\sqrt{T}U(0,1)$$

Or,

$$r_T = r_0 e^{(\mu(T) - \frac{1}{2}\sigma^2)T} e^{\sigma\sqrt{T}U(0,1)}$$

where $U(0,1)$ is a Standard Normal variable (that is, has mean zero and standard deviation 1).

This expression has three components, the initial interest rate, r_0 , the drift term (the first exponential) and the volatility term (the second exponential). We wish to adjust the drift term so that our rate lattice matches our current term structure so we will rewrite the equation above to be:

$$r_T = r_0 a_T e^{\sigma\sqrt{T}U(0,1)}$$

where we will determine the a_T terms to fit the term structure. This fitting process is the reason for not being rigorous with the integration of the time dependent μ term. We will replace the Standard Normal variable $U(0,1)$, with a Binomial variable “ $B(0,1)$ ”, as is the usual practice in building a binomial lattice.

A binomial variable, b_{ni} , is represented on a lattice at node i at time step n . It has value:

$$b_{ni} = \frac{(2i - n)}{\sqrt{n}}$$

and probability,

$$P_{ni} = {}^n C_i 2^{-n}$$

where the binomial co-efficient is

$${}^n C_i = \frac{n!}{i!(n-i)!}$$

and assuming a probability of one half for up and down movements. Consequently, this variable has mean zero and standard deviation one. (A further comment on our choice of probabilities is made below.) If we split our time to maturity T into n steps, then each step has length

$$\Delta t = \frac{T}{n}$$

Substituting into our expression for r_T ,

$$r_T(n, i) = r_0 a_T e^{\sigma\sqrt{n\Delta t} \frac{(2i-n)}{\sqrt{n}}}$$

Simplifying,

$$r_T(n, i) = r_0 a_T e^{\sigma\sqrt{\Delta t}(2i-n)}$$

So let us now move into the numerical part of the example. We will start with a continuously compounding zero coupon curve with rates for each quarter as below,

Exhibit 6.2

Years	0	0.25	0.5	0.75	1	1.25	1.5	1.75	2	2.25
Zero Rate	6.00%	6.15%	6.28%	6.40%	6.50%	6.58%	6.64%	6.69%	6.72%	6.75%

From this, we can generate a set of rates for each node using the formula above. A volatility of 12% per annum has been assumed.

Exhibit 6.3

Time Steps (quarters)

Node	0	1	2	3	4	5	6	7	8
0	6.1500%	6.7941%	7.4597%	8.0902%	8.6896%	9.2487%	9.8558%	10.3355%	11.0276%
1		6.0259%	6.6162%	7.1753%	7.7070%	8.2029%	8.7413%	9.1668%	9.7806%
2			5.8680%	6.3639%	6.8355%	7.2753%	7.7529%	8.1302%	8.6746%
3				5.6443%	6.0625%	6.4526%	6.8762%	7.2108%	7.6937%
4					5.3770%	5.7229%	6.0986%	6.3954%	6.8237%
5						5.0758%	5.4090%	5.6723%	6.0521%
6							4.7973%	5.0308%	5.3677%
7								4.4620%	4.7607%
8									4.2224%
a_T	1	1.040403	1.075797	1.098776	1.111456	1.114082	1.118078	1.104214	1.109543

This table represents a set of interest rates of period equal to a quarter of a year. The current spot rate, 6.15%, appears in the first cell corresponding to $r_0(0,0)$ and this is seen to progress either up to 6.7941% or down to 6.0259%. The probability of up or down movements is taken to be 0.5 and this reveals the reason for using a_T . We could use the formula value (from the integration above, done rigorously) for this and adjust the probability of up or down movements to fit the term structure. Alternatively, as has been done here, the probabilities can be chosen beforehand and the rates for the nodes adjusted to fit the term structure. The latter approach means that the formula for the probability at a node is the simple one shown above, rather than depending on a combination of the adjusted probabilities that would arise otherwise. Using probabilities of 0.5 makes a spreadsheet implementation much simpler.

How did we arrive at the values for a_T listed in the table? This is done iteratively. To illustrate using the first step we have rates $r_1(1,0)$ and $r_1(1,1)$ each occurring with probability 0.5. What is the expected two-quarter zero rate, R_2 , using these rates?

$$e^{R_2 \cdot 2\Delta t} = 0.5 * e^{r_0(0,0)\Delta t} * e^{r_1(1,0)\Delta t} + 0.5 * e^{r_0(0,0)\Delta t} * e^{r_1(1,1)\Delta t}$$

But we know that R_2 should equal our two quarter spot rate from our current zero coupon curve (6.28% in this example) so we can adjust the a_1 factor

which determines the r_I rates until the above condition is met. This is easy on a spreadsheet using a solve function (for example "Goal Seek" in Microsoft Excel) or can be done using any well-known numerical zero finding algorithm in a coded version. The two individual terms in the expression correspond to the node specific discount rates. These rates can be used to discount cashflows occurring on the nodes back to zero. This technique for building a lattice follows an approach which is known as forward induction. In this example the discount rate lattice has the following values:

Exhibit 6.4

Discount Rates	0	1	2	3	4	5	6	7	8
0	6.1500%	6.4721%	6.8013%	7.1235%	7.4367%	7.7387%	8.0412%	8.3279%	8.6279%
1		6.0879%	6.3921%	6.7413%	7.0873%	7.4188%	7.7448%	8.0522%	8.3668%
2			6.0146%	6.2435%	6.5610%	6.8994%	7.2439%	7.5739%	7.9088%
3				5.9220%	6.0787%	6.3420%	6.6572%	6.9831%	7.3246%
4					5.8130%	5.9087%	6.1215%	6.3901%	6.7018%
5						5.6902%	5.7437%	5.9001%	6.1347%
6							5.5626%	5.5753%	5.6966%
7								5.4250%	5.4180%
8									5.2914%
Expected Value	6.15%	6.28%	6.40%	6.50%	6.58%	6.64%	6.69%	6.72%	6.75%

From the first lattice of rates we can look at the value of our caplet and then using the second "discount" lattice we can get the net present value. If we take as our strike rate, 7.00%, then we need to compare each of our final node rates with this and construct the payoff. We then weight each payoff by the probability of occurrence and then discount back to today. The sum of these contributions is the caplet price.

Exhibit 6.5

Node	Rate	max ($r - K, 0$)	Value per \$1m	Discount Rate	Probability	Contrib- ution (\$)
0	11.0276%	4.0276%	10,068.96	8.6279%	0.00391	33.10
1	9.7806%	2.7806%	6,951.48	8.3668%	0.03125	183.76
2	8.6746%	1.6746%	4,186.51	7.9088%	0.10938	390.91
3	7.6973%	0.6937%	1,734.21	7.3246%	0.21875	327.66
4	6.8237%	0.0000%	—	6.7018%	0.27344	—
5	6.0521%	0.0000%	—	6.1347%	0.21875	—
6	5.3677%	0.0000%	—	5.6966%	0.10938	—
7	4.7607%	0.0000%	—	5.4180%	0.03125	—
8	4.2224%	0.0000%	—	5.2914%	0.00391	—
Price (\$)						935.43

The value per million is calculated as the percentage payoff applied to one million dollars for one quarter. The discounting raises an interesting point. Here we have discounted each caplet contribution by the discount factor that applies to its node—in effect assuming stochastic discount rates. You may recall that under BS assumptions there is a single constant discount factor for the period, that is discount factors are not stochastic. If we apply the two year spot rate (6.72%) as the discount rate for each contribution to the caplet we end up with a price of \$955.69 per \$m face value. The difference between \$955.69 and \$935.43 shows the relative size of the error made under a BS assumption of constant discount rates. A single constant volatility was used for the caplet valuation. Use of different volatilities and thus different lattices for the different caplets in a cap would allow the calibration of the cap price to the volatility term structure as well as the interest rate term structure.

We can make some concluding comments about this example. A simple interest rate process was assumed and from this a binomial process in which we made the assumption of 0.5 as our up and down probabilities. We built the lattice including a variable drift component which we then solved for iteratively to fit the current term structure of interest rates (in zero coupon curve form). Part of this iterative fitting yielded a discount lattice which is useful for discounting individual cashflows back to today from specific nodes. From the quarterly rate lattice we determined the payoff in percentage terms then converted this to a dollar amount. This was discounted back to today to provide us with a caplet price. We observed the different effect of discounting with the stochastic rates as opposed to using a single discount rate.

This approach can easily be extended to a full cap (or floor) and can also be easily implemented on a spreadsheet. Note that to increase the number of steps beyond one per quarter would require the final payoff rate of period one quarter being determined as an expected value from the rates beyond the expiry node.

5. CONCLUSION

In concluding this chapter I hope that I have provided the reader with an overview of the breadth of interest rate option pricing models and the various reasons for this diversity. The underpinning of all of these models is an assumption about the underlying interest rate process and interpreting the SDE allows us to identify what distribution of rates is being assumed, what factors are time-dependent and may thus be used to fit exact term structures, and whether such features as mean reversion have been included.

Through an example we looked at the choices required for building a lattice model and the way in which factors can be adjusted to calibrate the model to existing term structure information. Obviously in such a short chapter there are many things which have been simplified but I hope that the flavour of the topic has been made accessible.

There are a number of books which cover the field of interest rate option pricing but those of Hull and Jarrow are readily accessible to the non-mathematician and the review article by Ho and the chapters in the book by Rebonato provide discussion of individual models.

Chapter 7

Pricing Models for Complex/Exotic Options

by Dr Garry de Jager

1. INTRODUCTION

Until fairly recently, if reference was made to financial options, it almost certainly referred to either a simple call or put, or, in the field of interest rates, a simple cap or floor. Since that time, more complicated options and other derivative instruments have also become popular, notwithstanding the fact that the “vanilla” options and derivatives still constitute the bulk of the trades.

We can think about exotic derivatives, and exotic options in particular, as coming in two “waves”. The first wave consisted of products such as average rate options, barrier options and digital options, and these are often now not considered substantially different from the “vanilla” products. The second wave might be considered to be products such as knockout options with limited barrier periods, power options, cross product options, range binaries and spread options. In general we might say that the more difficult it is to understand the details of an option, the more difficult it is to assess a fair price and the more sources of uncertainty or risk, the more likely it is that the option will be considered “exotic”.

2. EMERGENCE OF EXOTIC OPTIONS

Exotic options tended to emerge in the face of three influences: customer needs, progress in pricing theory (often led by academics), and improved trader ability to manage risk.

2.1 Client needs

Standard European options have:

- (a) a payoff at one point in time, that is, the maturity of the option;
- (b) this “payoff” is a standard linear function of the price of the underlying asset; and
- (c) this “price” is the underlying’s value at maturity.

However, not all clients want to be rewarded only at maturity, nor might they necessarily want to be rewarded according to the level of the price of the underlying asset, and if they do, then it might not necessarily be a linear function. We can readily find examples for each of these three cases.

First, a client may wish to be rewarded when a particular event occurs, for example, at the exact moment when the price of the underlying moves

through a particular price level. A knock-in barrier option and knock-out barrier with rebate are examples of such a derivative. Secondly, a client may wish to be rewarded with a specific pre-determined cash sum if the underlying is above (or below) a certain price level at maturity. A digital option is an example of such a derivative. Thirdly, a client may prefer a payoff function that is different from the standard linear function of price at maturity. For instance, he/she may prefer a more generous payoff than the classical $MAX(0, spot - strike)$ for a standard call. Two such payoffs would be:

- (a) $MAX(0, spot * spot - strike)$; and
- (b) $MAX(0, spot - MIN(spot))$.

In the first case the payoff is usually much greater if the option is in the money at maturity, particularly when the asset price is unexpectedly high. A power option is an example of such an option. In the second case we are dealing with a lookback option, that is, at maturity we “lookback” over the past history of prices and pick the lowest one, and this replaces the standard strike in an ordinary call.

2.2 Progress in pricing theory

Following the publication of the seminal Black-Scholes¹ pricing formula for equity options, there rapidly followed a series of publications, primarily produced by academics, pricing options in other domains, for example, the Black² model for pricing options on commodities and futures, the Grabbe³ model for options on foreign exchange, the Roll/Geske/Whaley⁴ model for options on equities with dividends. There also followed a plethora of models to price more complicated styles of options, regardless of what the underlying might be—the Cox, Ross and Rubinstein⁵ binomial American option model; the Goldman et al⁶ lookback option model, which was further developed by Conze and Viswanathan;⁷ the arithmetic average rate and average strike option models; the Rubinstein and Reiner⁸ barrier option models; digital options; the de Jager/Winsen⁹ generalized power option models, and many more.

As option trading desks began hiring mathematical finance researchers (often referred to as “rocket scientists”), virtually any style of payoff

1. F Black and M Scholes, “The Pricing of Options and Corporate Liabilities” (1973) 81 *Journal of Political Economy* 637.
2. F Black, “The Pricing of Commodity Contracts” (1976) 3(1 & 2) *Journal of Financial Economics* 167.
3. J O Grabbe, “The Pricing of Call and Put Options on Foreign Exchange” (1983) 2 *Journal of International Money and Finance* 239.
4. R Roll, R Geske and R Whaley, “On the Valuation of American Call Options on Stocks with Known Dividends” (1981) 9 *Journal of Financial Economics* 207.
5. J C Cox, S Ross and M Rubinstein, “Option Pricing: A Simplified Approach” (1979) 7 *Journal of Financial Economics* 229.
6. M B Goldman, H B Sosin and M A Gatto, “Path Dependent Options: ‘Buy at the Low, Sell at the High’ ” (1979) 34(5) *Journal of Finance* 1111.
7. Conze and Viswanathan, “Path Dependent Options: The Case of Lookback Options” (1991) 46(5) *Journal of Finance* 1893.
8. M Rubinstein and E Reiner, “Barrier Options” (1991) *Exotic Options* (unpublished manuscript by M Rubinstein).
9. G de Jager and J Winsen (1992), “Power over Gamma: Curved Option Payoffs”.

became possible. Further, options depending on the paths followed by more than one asset, also became possible. "Rocket scientists", using various techniques usually derived from the original Black-Scholes PDE ("partial differential equation"), together with established statistical techniques involving the volatilities of, and the correlations between, two, three or even more assets, have been able to provide traders with quite robust pricing models for these "hybrid" style options.

Finally, in the interest rate domain, there has been a flurry of activity to improve on the standard market model which is a simple application of the Black¹⁰ model. These new improved models are usually referred to as "term structure" models, because they create a "tree" structure of possible future interest rates, and the option prices at each of the nodes are discounted depending on the interest rate at that node. These trees theoretically price exotic interest rate options more accurately than do standard formulae.

2.3 Traders and the management of risk

Unless traders feel confident in managing the not inconsiderable risks associated with exotic options, the financial community would be faced with the prospect of having a plethora of theoretical exotic option pricing formulae, but also having traders offering quotations with very wide bid/ask spreads, making it impracticable for clients to transact deals. Fortunately, both the power of computing hardware, and the availability of large-scale computing software systems, have enhanced the ability of traders to risk manage their portfolios more effectively. This increase in computer efficiency coincided with the entrance of exotics into the marketplace. Traders over the years have been able to develop sophisticated techniques to manage their books including vega and gamma immunization, and risks allocated to various "time buckets". The refinement of these techniques, prior to the arrival of exotics, resulted in the management of the latter being more smooth than might otherwise have been the case.

3. WHY EXOTICS ARE DIFFERENT

In the previous section, the implications were that exotics were different from the standard call and put options, and not only different, but far more complex to price and possibly to manage. Exotics are similar to the standard option in that they generally retain the concept of the "right, but not the obligation, to buy or sell a quantity of the underlying for what may prove to be a favourable price". The differences generally revolve around two issues (a) the activation of the option, and (b) the payout function of the option.

10. Black, op cit n 2.

3.1 Pricing

When pricing a standard European call or put option, it transpires that, in the words of Robert Merton (one of the doyens of option theory), “there is a trick” thanks to the technique of dynamic hedging:

- (a) we can actually assume that all investors are risk neutral—therefore they require only the risk-free rate on any investment—no matter how risky;
- (b) therefore in this (unreal) world all underlying assets will be priced on the basis that they will return, on average, no more (or less!) than the risk-free rate;
- (c) therefore we need only calculate the probabilities of the underlying’s prices at maturity (based on asset growth at the risk-free rate); and
- (d) hence in turn calculate the “expected” value of an option at maturity; and
- (e) finally, and most importantly, discount this expected value at the risk-free rate, back to today.

A problem with many exotics is that the technique above is not sufficient. We have seen, for instance, that some exotic options do not depend only on the price of the option at maturity. Thus the simple discounting from maturity may not apply—if we have a payoff that may occur at any time, then anticipating the discount factor will not be a trivial matter. Further, with the entire family of “term structure” models, even if we know the time of the payoff (maturity or some other time), we still don’t know the appropriate discounting factor since it depends on the path that interest rates have taken to that point. For these and other reasons, more complicated pricing techniques are required. We can see that these changes are basically examples of different “activation” and different “payoff functions”.

3.2 Managing risk

If pricing is not a simple matter of a fairly standard closed form mathematical expression, then the risk management is also often more difficult. We have to consider that the most elementary form of hedging is the basic “delta hedge” where a position in the underlying is used to counter the effects on option values of a move in the price of the underlying. “Smooth” option value functions usually lead to smooth hedging functions, that is, relatively small movements in the price of the underlying require relatively small trading strategies. However, with an exotic, the value function may not be “smooth”, and relatively small movements in the price of the underlying may lead to substantial changes in option revaluation—and hence to a very substantial trading strategy. These uncertainties make the risks associated with the option potentially more difficult to manage.

4. GOING ABOUT GETTING A FAIR PRICE

The pricing of exotic options needs to incorporate the path dependencies that may activate a payoff or the option itself, as well as the potentially difficult payoff functions. Faced with a difficult option valuation, a number of techniques may be employed by our “rocket scientist”. Seven techniques in his/her armory are:

- (a) integration of integrals using advanced mathematics;
- (b) solving single integrals, or multi-dimensional integrals using numerical techniques;
- (c) employing complex binomial or multinomial lattices (or trees);
- (d) employing finite difference techniques;
- (e) using monte carlo simulations;
- (f) creation of term structure trees; and
- (g) use of approximating techniques.

4.1 Integration/advanced mathematics

An example of complex integration is the limited period barrier option where the knockout period begins some time after option initiation, and ends some time prior to option maturity. To price such an option a triple integral is required and tedious integration is required to provide a solution. Complex probability distributions, affected by path dependencies, often contribute to the complexity.

4.2 Integrals/numerical techniques

Where integrals are difficult, or impossible, to solve, standard numeric integration techniques may be employed. This is often the case with complex options whose payoffs depend on more than one underlying asset. In the case of multiple integrals, it is often possible to reduce the problem to a single integral which requires a solution using a variant of a standard technique often employed in high school mathematics—calculating an area under a curve. In the high school approach, students plot the curve on graph paper, and physically count the area under the curve. Using computers, the “rocket scientist” will employ standard mathematical techniques such as *Simpson’s Rule* or *Romberg Integration* to calculate the area under the curve so as to calculate an option’s value.

4.3 Binomial/multinomial lattices

An appealing method for less technical people to cope with is the notion of a complicated *decision tree* with probabilities connecting each of the nodes on the tree. In the case of option pricing these are usually simplified into what we call *recombining trees* rather than *exploding trees*, that is, at each time step nodes are connected to multiple nodes on the next time step. As with decision trees, the calculations begin with option values at the maturity of the option, and we work backwards using the probabilities, until

we reach the node at the inception of the option. American options of all types are usually priced in this manner. Activation of options or determination of intermediate payoffs can be calculated more easily in this discrete time environment.

4.4 Finite difference schemes

A similar scheme to the above is the finite difference lattice. Although this system uses *coefficients* rather than *probabilities*, the effect is similar. This particular technique creates a lattice much like a chess board—in this case, rectangular with an identical number of nodes at each time step. The technique is widely used in mathematics, and option pricing fits neatly into a subset of the applications referred to as the *heat equation*. However, in the world of option pricing there is a case for maintaining that almost anything that can be achieved with a finite difference scheme, can be achieved more simply and more accurately with a multinomial tree.¹¹ The scheme is also used to value American options, and also options which exhibit a *volatility smile*.

4.5 Monte Carlo simulations

There has been an increasing interest in using a fairly “blunt instrument” known as *Monte Carlo* simulation. This method generates a path of prices for the underlying asset, and the option is priced at maturity. The method is repeated, with a second option price determined. Many thousands of trials follow, and the average of all the trials is deemed to be the option price. The field of monte carlo simulations is a science in itself, and many sophisticated techniques can be employed to speed the process or make it more accurate. In recent times the *pseudo random number* approach has been added to the *antithetic technique* to provide fast and relatively accurate option valuation. Typical options that are valued in this fashion include complicated discrete averaging payoffs where perhaps there are criteria for the number of business days that are required to be above or below a certain target price. This method generally avoids the complex mathematics associated with path dependencies.

4.6 Term structure trees

Standard binomial or multinomial trees where the underlying is an interest rate level suffer from a number of disadvantages:

- (a) Since interest rates themselves are not traded—products such as bonds and swaps which are dependent on the levels of interest rates are traded—the expected price of a bond some time in the future will be equal to the forward bond price, but the expected interest rate for a forward contract will generally not be equal to the forward interest rate. This gives rise to the well-known *convexity adjustment* between the

11. G. de Jager, “Option Pricing With Implicit Finite Difference Methods and Large-Scale Bounded Multinationals” (1994) Proceedings of Asia Pacific Finance Association Conference 1994.

forward yield and the expected yield, which complicates the construction of a tree of interest rates.

- (b) If future payoffs on such trees are to be discounted, the discount factors are not at all clear, and, theoretically should reflect the path of interest rates taken to the node in question, that is, path dependency is far more difficult than for other options.

The three most popular term structure methodologies have been the Heath, Jarrow and Morton method,¹² the Hull and White method,¹³ and the Black, Derman and Toy¹⁴ (or Black and Karasinski) model.¹⁵ The trees are *calibrated* against cap prices from the market model, and then used to price exotic options of almost every shade.

4.7 Approximating techniques

There will be cases where the problem is difficult to formulate, but there is an obvious connection with a similar problem that has a known answer. Examples include the desired arithmetic average versus the known geometric average; knockouts measured at discrete intervals, versus known probability distributions for continuous knockouts; control variate techniques in monte carlo simulations. The discrete time barrier is an interesting example. Broadie¹⁶ et al showed that a good approximation can be made by using the continuous barrier formula, but altering the barrier input into this formula by a specific multiple. Simulations indicate this approximating technique is effective.

5. BEHAVIOUR OF PRICES AND RISK MEASURES

The behaviour of option prices, deltas, gammas and vegas often distinguishes the exotic from the traditional European call and put and is therefore of particular interest. We illustrate this by graphing option prices and risk measures for three exotics—average rate option, down and out European barrier call, and binary range floater. The delta in the following figures is the traditional fraction of one unit of the underlying stock to buy/sell to hedge the option, but the vegas and gamma risk measures are *real* values rather than *theoretical* values and are calculated as follows:

- (a) *Dollar Gamma*: The price moves up 1%, the option delta changes—we sell more of the stock for a call—and then the price moves back to original level resulting in our buying back the extra hedge at a lower price. The gain made is called the *dollar gamma*.

12. D C Heath, R A Jarrow and A Morton, "Contingent Claim Valuation with a Random Evolution of Interest Rates" (1990) 9(1) *The Review of Futures Markets* 54.
13. J C Hull and A White, "Efficient Procedures for Valuing European and American Path-dependent Options" (1993) 1(1) *Journal of Derivatives* 21.
14. F Black, E Derman and W Toy, "A One-Factor Model of Interest Rates and its Application to Treasury Bond Options" (1990) (Jan-Feb) *Financial Analysts Journal* 33.
15. F Black and P Karasinski, "Bond and Option Pricing When Short Rates are Lognormal (1991) Jul/Aug) *Financial Analysts Journal* 52.
16. Broadie, Glasserman and Kov, "A Continuity Correction For Barrier Options" (1995) Proc of European Risk Conference Paris 1996.

- (b) *Dollar Vega*: The volatility moves up 1% and we note the change in the option price. The difference in price is termed the *dollar vega*.

5.1 Average rate call

The formula for a call on the geometric average stock price over the period to option maturity (rather than on the stock price at maturity), and without any past history of prices, is:

$$e^{-rT} \left[S e^{\text{mean} + 0.5 \text{var}} N(d1) - KN(d2) \right]$$

Where:

- T = time to maturity
 r = risk-free rate to maturity
 S = current spot price
 K = strike price
 mean = $0.5 * (r - d - 0.5 * \text{vol} * \text{vol}) * T$
 vol = expected volatility to maturity
 d = dividend yield/foreign exchange rate
 var = $\text{vol} * \text{vol} / 3$
 d2 = $[\log(S/K) + \text{mean}] / \text{sqrt}(\text{var})$
 d1 = $d2 + \text{sqrt}(\text{var})$



In practice we usually have to make two adjustments:

- (a) first, the traditional average option is calculated using arithmetic, rather than geometric averages, and
- (b) secondly, the averages are usually calculated using a discrete, say daily, set of future prices, rather than the “continuous” over time average implicit in the above formula.

We now price an average rate arithmetic call option which has a remaining 42 business days within 60 calendar days to run to maturity and the average is constructed using a 10 am fixing on each business day. Initially there were 142 business days to incorporate into the average, so there have been 100 previous business days, and the average spot on these was \$39.85. The option strike is \$40, market volatility is now 18.5%, and the risk-free rate to option maturity is 8% (expressed as an annually compounded in arrears rate). *Exhibits 7.1, 7.2, 7.3 and 7.4* plot the option price, delta, gamma and vega respectively, against those of a regular European call. The plottings vary the spot price of the underlying stock with values from \$32 to \$56 being employed.



It is interesting to see that the price of the average rate call is much more stable than that of its regular European call counterpart, and this is confirmed by the lower delta values in *Exhibit 7.2*. This is clearly intuitive, since the bulk of the averaging has already occurred, and thus changes in current price (which is the major predictor of the remaining 42 prices to be averaged) are somewhat overwhelmed by the history of past prices. This is not the case with the standard European call where the major pricing factor is the spot price of the underlying stock. We note that the standard European call has

deltas approaching 1 when the option is deep in the money, whereas for the same stock values, the average rate call deltas level off at just below 30%. The explanation for this is that for every dollar increase in the price of the stock, the average rate call does not increase by a dollar—since only 42 of a total of 142 prices are affected—and hence around a 30% hedge adequately covers changes in the option value.




Because the prices of the average rate call option are not as sensitive to price changes in the underlying stock as are standard calls, we would expect the gamma of the option to be less as well, and this proves to be the case in *Exhibit 7.3*. Finally, we would expect that changes in the volatility over the remaining life of the option will also have less effect than normal, since they

potentially only affect future prices of the underlying and cannot affect the 100 prices that have already occurred. This is confirmed in *Exhibit 7.4*.



We note that what we might term the “stabilising effects” outlined above are not so dramatic if the ratio of past prices to future prices in the average rate call option is reduced. *Exhibit 7.5* demonstrates that the prices and deltas are much closer when we consider a similar option to the above, but this time with 510 business days across 730 calendar days still to run to maturity, as against 100 past prices to be incorporated into the average.



5.2 Down and out European barrier call

The formula for a down and out call on stock varies depending on whether the barrier is above or below the strike. Where the barrier is below the strike, we have:

$$CALL = e^{-rT} \left[e^{(r-d)T} S \left(\frac{H}{S} \right)^{2\mu/(\sigma^2\sigma)+2} N(d1) - K \left(\frac{H}{S} \right)^{2\mu/(\sigma^2\sigma)} N(d2) \right]$$

Where:

- T = time to maturity
- r = risk free rate to maturity
- S = current spot price
- K = strike price
- H = barrier price
- σ = expected volatility to maturity
- $\mu = r - d - 0.5\sigma^2$
- d = dividend yield/foreign exchange rate

$$d2 = \frac{\ln\left(\frac{H^2}{SK}\right) + \mu T}{\sigma\sqrt{T}}$$


$$d1 = d2 + \sigma\sqrt{T}$$

In practice we usually have to make a price adjustment to this formula, since the barrier is often based on discrete time periods, say daily 10 am fixings, rather than the "continuous" barrier implicit in the above formula.

We now consider a down and out call option on a stock where there are 365 calendar days to maturity, with strike at \$40, market volatility at 18.5%,

risk-free rate to option maturity at 8%, and a barrier of \$36. *Exhibits 7.6, 7.7, 7.8 and 7.9* plot the option price, delta, gamma and vega respectively, against those of a corresponding regular European call. The plottings vary the spot price of the underlying stock from \$32 to \$44.


The first point of interest to note in *Exhibit 7.6* is that the barrier call price is zero if the spot is below the barrier, that is, the option has already been knocked out. At just above the barrier stock price of \$36, the barrier call has a very low value, as would be expected—there is a very good chance the option will be knocked out because the stock price is so close to the barrier. At higher stock prices, the barrier call is worth nearly the same as a standard call since, according to the statistical probabilities of a bell-shaped lognormal curve of future stock prices, there is minuscule probability that the stock price will drop back down to the barrier in the remaining year to option maturity.




The delta of the barrier option, however, exhibits what at first might seem “bizarre” behaviour. There are two determinants here. First, at high stock prices, the delta should be close to the same as that for a standard option, since there is a low probability of any knockout occurring. Secondly, at a low price, the delta will be either (a) zero for stock prices below the barrier, or (b) rising as the stock price moves away from the barrier. Just how rapidly the delta rises is the interesting thing in *Exhibit 7.7*. We note that it climbs to almost 1.2. This is explained as follows: At \$38, the stock price is still only \$2 above the barrier and the option is still in substantial danger of being quickly knocked out. To guard against this complete loss in value of the barrier call, the trader would sell 1.2 shares in the underlying stock—a type of “leveraged hedge”—so he/she is compensated for a falling stock price which knocks out, or threatens to knock out, the barrier call.


This rapid rise, and subsequent fall, in the delta as stock prices rise, means we have a large gamma which is initially negative, and then positive, and this is illustrated in *Exhibit 7.8*. (We should also note that there is a

discontinuity at the barrier price of \$36.) The negative gamma above a stock price of \$36 is obvious from *Exhibit 7.7*, which indicates that as stock prices rise, the knock out call delta reduces. Now we can see a clear contrast here as compared to a standard call. If a trader is long the standard call, then typically he will hold a negative position in the stock as a hedge. If the price of the stock rises, the delta rises, and the trader sells more of the stock. However, with a down and out call, although the trader, at stock prices just above \$36, will have sold more than one unit of the stock for every option held, if the stock price rises, the trader will actually buy back some of his/her hedge, rather than sell more of the stock. It is well known, that in these circumstances—that is, being negative gamma—a move up and then back in the price of the stock results in a “dollar gamma loss” for the portfolio of the long option and the short hedge.






Finally, we look at the dollar vega of the down and out call. The two determining factors which give rise to the patterns in *Exhibit 7.9* are (a) that the vegas of the European and down and out calls will converge as the stock price increases, and (b) the shape of the vega curve for the down and out call, whilst similar to that for the European call, is of course zero below the barrier, and the vega is always much smaller than that for the regular call, since the price of the option is always smaller.



5.3 Range binary

A standard range binary deal generally involves a client depositing, or notionally depositing, a principal with the trading house, but forgoing the interest thereon. In return, the client receives a payout that is potentially better than the standard libor interest rate, should three months libor remain within a defined band for most of the period. For instance, in one form of this deal, as represented in *Exhibit 7.10*, the client has a principal of \$10 million dollars on deposit with the trader, and the client forgoes the current market libor interest rate of 5.5% for the next three months. In return, the client will receive 6.15% for every day over the next three months that the daily libor fixing is within the band 5.15% to 6.15%, and zero for the days it is outside the range. If libor is within this band the entire time, then the client will receive \$152,500 in "interest" as opposed to the regular amount of \$138,700. This latter amount is represented by the horizontal line "fixed side". However, if halfway through the three months there is a half a per cent fall in rates and this situation persists for the remaining month and a half, then the client would only receive \$76,250 in "interest".

The *Exhibit 7.10* plots the expected payout to the client for a given range of volatility expectations. We note that the higher the expected volatility, the lower the expected payout to the client. On reflection, this should be obvious since (a) libor is currently within the range, and (b) a very low actual volatility would almost certainly ensure it stays within the range, and therefore (c) a high actual volatility will dramatically increase the chances of libor moving out of the range for at least part of the next three months.



6. CONCLUSION

The essential ingredients of exotic options seem to be path dependency and/or complex payoff functions. They are usually more difficult to price and more difficult to manage in a portfolio. There are many similarities between exotics and standard options, and in one sense we can say that we have merely been stretching the current concepts of optionality a little further. In the future, we might have newer concepts of optionality.

Much of the probability distribution work in option pricing is still based on the original Black-Scholes¹⁷ framework of log-normal brownian motion with constant volatility. As new stochastic processes are developed, we might logically see the emergence of newer exotic style options. Jump processes might logically be involved in such developments.

17. Black and Scholes, op cit n 1.

Appendix

A Glossary of Some Popular Exotic Options

Average Rate

The call or put payoff now depends on the average of the underlying between two specified dates, rather than on the underlying's price at maturity. The two selected dates are usually the option start and expiry. (Also called an Asian option).

Average Strike

The call or put strike is not set at option initiation, but is deferred till option expiry, whence the average of the underlying between two specified dates takes the place of the regular strike.

Barrier Option—In (Single)

An agreement whereby the purchaser obtains a vanilla European option under certain conditions, the central condition being that a particular price level of the underlying is breached. If the timing of the breaching of the barrier is restricted to option maturity, then sometimes it is called a "European barrier". Generally the barrier period is from option inception to option maturity. Some barrier options operate on a discrete basis (for example, the price level of the underlying is observed only once a day, whereas others operate on a continuous time basis). (Also known as "knock-in" options.)

Barrier Option—Out (Single)

The purchaser of the options stands to lose all the claims associated with the option should the price of the underlying breach a predetermined level; that is, the "barrier" level. (Also known as "knock-out" options.)

Barrier Options—Limited Period (Single)

These are similar to the above knock-in and knock-out options, except that the time period over which the barrier may be breached to effect a knock-in or knock-out is limited. If the time period is at the beginning, it is often referred to as a "front-end barrier", whereas if the time period is restricted to a period ending at option maturity, it is often referred to as a "back-end barrier". A more general form allows the barrier period to commence after option inception and end prior to option maturity.

Barrier Option (Double)

Similar to any of the above barrier options except that now we have two barriers—an upper and a lower barrier. For a "double knock-in" option the price of the underlying needs to breach either of the two barriers during the specified time period, whereas for a "double knock-out" option, the underlying needs to avoid breaching either of the two barriers else the option immediately lapses.

Barrier Option (Parisian Style)

This option can take the form of any of the above types of barriers, but now there is an additional requirement, namely that for the barrier to be considered "breached", the underlying must be beyond the barrier for a given number of days.

Ratchet Options

These options have their strikes reset (usually on a favourable basis) during the life of the option, or in the case of a cap, successive caplets may have their strikes reset favourably. Similar concepts underlie the *cliquet* or *resettable* options. The resetting will usually be the result of some level of the underlying being reached, and the possible reset times may be restricted to particular occasions (for example, the end of the previous caplet).

Bermuda Options

Options that may be considered half-way between European and American; that is, Bermudan (also known as mid-Atlantic options). The idea here is that on a number of particular intermediate dates, the option may be exercised.

Digital Options (European Style)

Basically any of the options in this section can also have a digital payoff. Unlike the standard call and put options, digital options have a set dollar payout at maturity if the underlying price is above the strike—for a digital call, or below the strike—for a digital put. Digital versions of barriers and double barriers are also possible.

Digital Options (American Style)

This form of the digital pays out the moment the strike is first reached during the life of the option. This option can be incorporated into the standard knockout option where a “rebate” is paid should the barrier be breached.

Optimal Rate LookBack Option

The call or put payoff depends on the most favourable price the underlying experienced between two specified dates (highest price for a call, lowest price for a put).

Optimal Strike LookBack Option

The setting of the call or put strike is deferred to option maturity whence the most favourable price the underlying experienced between two specified dates (lowest price for a call, highest price for a put) takes the place of the regular strike.

Power Options

A power call will payoff $\text{Max}(S^N - K, 0)$ instead of the usual payoff. If the spot is above 1, which it usually will be for these types of options, the payoff can be very high very quickly, especially if N is greater or equal to 2.

Quanto Options

An option where the underlying asset and the payoff calculation are in different currencies (for example, a cap on deutschmark interest rates where the notional principal payoff is in US dollars, and the “option compensation” is US dollars).

Shout Option

This option has some of the aspects of the *optimal rate lookback option* in that a favourable level achieved during the life of the option may be used to replace the underlying's price at option maturity. However, in this instance, the holder is not guaranteed the *most* favourable price of the underlying for the payoff settlement, but gets a single opportunity to nominate this price. Thus, if at any time during the life of the option the holder feels the *top of*

the market has been reached, he or she may nominate that price as a "fallback" level. Then at maturity, if this fallback level is more favourable to the outcome than the underlying's price at maturity, it will be used instead.

Spread Options

These payoff on the difference between the level of two underlying assets. A wide variety of spreads are possible (for example, the difference between two interest rates of different tenors, the difference between two metal prices, the difference between two futures prices on the same commodity).

Chapter 8

Estimating Volatility

by Satyajit Das

1. OVERVIEW

The concept of volatility of asset prices and returns is central to financial markets. Volatility provides essential data about the probability of achieving certain outcomes in terms of price levels which is intrinsic to key decisions in financial markets, such as asset allocation and construction of efficient asset or liability portfolios (in the context of a risk-return trade-off). In the context of option pricing, an estimate of volatility is essential to the valuation of the instrument. This chapter focuses on the problem of volatility estimation in the context of option pricing.

The structure of this chapter is as follows: a framework for volatility, covering causes of volatility in asset markets and the relationship between volatility and option pricing, is first examined. Approaches to volatility estimation are then considered, including historical volatility, implied volatility, as well as alternative approaches to volatility modelling. The behaviour of volatility, particularly the concept of the volatility smile and the term structure of volatility, are analysed. The Chapter concludes with an analysis of volatility estimation in practice.

2. VOLATILITY ESTIMATION—FRAMEWORK

2.1 Introduction

Volatility estimation, in the context of option pricing, must be considered in the broader context of asset price and return volatility generally. The framework for volatility estimation, in reality, must recognise the causes of volatility in asset prices and the inter-relationship between volatility and option pricing models.

2.2 Causes of volatility in asset prices

Price volatility in asset markets is caused by a variety of factors, the most important of which is information release. A second cause of volatility is the process of trading and market-making in financial instruments.

Information release generally fall into two categories: anticipated and unanticipated information. Anticipated information includes economic statistics as well as political or social information. This type of information release is anticipated. Market participants develop expectations regarding the

content of the informational release. The impact of the information is, often, driven by whether or not it corresponds to market expectations, reflecting the fact that asset prices will generally incorporate the content of expected information releases. The impact of information releases can be analysed, particularly with reference to past releases. It may be possible to develop probabilistic expectations of anticipated asset price volatility from the historical reaction of the market to prior actual data releases in combination with probabilities in relation to a variety of range outcomes for the relevant information release.

Unanticipated information releases typically relate to international events (wars, natural disasters, et cetera) and other unanticipated or unanticipatable events. This type of information can have substantial and unpredictable impact on asset price volatility. The difficulty in predicting this type of informational release (by definition) makes it extremely complex to incorporate these types of factors in forecasting future asset price volatility.¹

A newer area of financial economics research (the study of market micro-structure) seeks to isolate the impact of trading on volatility.² This research identifies the informational content of trading and its interaction with the institutional structure of markets as a possible source of volatility in asset markets.³

In seeking to isolate factors generating asset price volatility, the linkages between volatility in various market segments or across markets should be noted. Analysis indicates that there may be implied and historical volatility relationships between different markets. For example, bond market volatility across various currencies show significant correlation. Similarly, in certain currencies, volatilities in the foreign exchange market are, often, useful indicators of volatility in interest rate markets in the relevant currencies.

2.3 Relationship between asset volatility and option pricing

Mathematical option pricing models, such as the Black Scholes model, require estimation of the future volatility of the underlying asset price, as this item is a parameter which must be input into the model. Binomial models require similar inputs.

The volatility estimate used in option valuation is the annualised standard deviation of the logs of the asset returns (or the continuously compounded

1. For a recent analysis of the impact of release of information of volatility, see Louis Ederington and Ja Hae Lee, "The Impact Of Macroeconomic News On Financial Markets" (1996) 9 (Spring) *Journal of Applied Corporate Finance* 41.
2. For example, a number of studies have found that volatility as between close of trading on Friday and opening of trading the next Monday morning (when there is an interval of around three non-trading days) is only around 20% higher than that as between close of trading on a one day and open of trading the next day (when there are no intervening non-trading days) rather than the predicted three times higher, suggesting that volatility is higher when trading on exchanges is open than when it is closed: see E E Fama, "The Behaviour Of Stock Market Prices" (1965) 38 (Jan) *Journal of Business* 34; K R French, "Stock Returns And The Weekend Effect" (1979) *Journal of Financial Economics* 55.
3. See Kalman J, Cohen, Steven F Maier, Robert A Schwartz and David K Whitcomb (1986) *The Microstructure of Securities Markets* (Prentice-Hall Englewoods Cliffs, New Jersey, 1986); Robert A Schwartz, *Equity Markets* (Harper & Row, New York, 1988).

asset returns). The volatility estimate is a measure of the uncertainty about the returns on the asset. It is used to generate the distribution of asset prices as at the option expiry to calculate the fair value of the option.

There are several aspects of the volatility estimate which should be noted:

1. The volatility parameter required to derive option values is forward looking; that is, the relevant volatility is the asset return volatility in the period to option expiry.
2. Volatility is assumed to be constant as between the pricing date and option expiry.
3. Volatility is assumed to be time homogenous; that is, it is the same over the life of the option.
4. Uncertainty about the asset price at option maturity is assumed to be directly proportional to the asset price at commencement.

The estimation of the volatility of the underlying asset price is particularly problematic because it is the only parameter of most mathematical pricing models which is not observable directly. The sensitivity of the option value to this parameter places additional demands on the estimation of volatility.

3. VOLATILITY ESTIMATION

3.1 Approaches

Estimation of the *true* volatility of the underlying asset price is extremely difficult. In practice, a number of alternative approaches are utilised. The major approaches include:

- historical/empirical approach; or
- implied volatility approach.

In recent times, a number of other approaches to option volatilities have emerged. These include ARCH type volatility models.

3.2 Historical/empirical volatility

3.2.1 Calculation of historical volatility

Under the historical or empirical approach, volatility estimates are calculated as the standard deviation of logs of the price changes of a sample time series of historical data for the asset price.

This calculation procedure entails the following steps:

1. The time series of historical data is specified. This will usually be the series of daily, weekly or monthly price observations for the relevant asset. As discussed in more detail below, for debt securities either price or yield can be utilised.
2. The price changes are calculated to measure the periodic (daily, et cetera) return on the asset. In practice, the price relatives are utilised; that is, one plus the return or the observation at time t_1 divided by the observation at the previous point in the time series t_0 . While the

standard deviation can be calculated for either the returns or the price relatives the first leads to inaccuracies reflecting the nature of the log normal distribution based on the effect of compounding. This means that the calculation using the price relatives is preferred. The difference is not significant for calculations involving relatively short data series but becomes increasingly significant as the data series increases in size.

3. The standard deviation of the price relatives is then calculated.

The interpretation of the standard deviation is as follows:

- The standard deviation computed equates to the volatility over the relevant time interval (daily, et cetera).
- The periodic observation is then scaled to give the annualised volatility of the asset price returns. This is done using the following relationship:

$$\sigma_{\text{annual}} = \sigma_{\text{daily}} \times \sqrt{\text{number of days}}$$

where the number of days would be set at either 250 or 260 days.

Exhibit 8.1 shows a possible sequence of asset prices over a 20 day period. The data gives an estimate for the daily volatility of 0.00889 or 0.889%. Assuming that time is measured in trading days and that there are 250 trading days per year, the estimated annual volatility is 0.14051 or 14.051% pa.


Note that the daily volatility is scaled using the square root of the time interval. This reflects the fact that the assumed uncertainty about the asset price does not increase linearly.

Exhibit 8.1
Volatility Estimation—Historical/Empirical Approach

CALCULATING VOLATILITY

Period	Asset Price	Price Relative (St/(St-1))	Daily Return UI = ln (St/(St - 1))
0	8.2500		
1	8.1800	0.99152	(0.00852)
2	8.2000	1.00244	0.00244
3	8.2000	1.00000	0.00000
4	8.2100	1.00122	0.00122
5	8.1200	0.98904	(0.01102)
6	8.1500	1.00369	0.00369
7	8.0700	0.99018	(0.00986)
8	8.1300	1.00743	0.00741
9	8.1700	1.00492	0.00491
10	8.2000	1.00367	0.00367
11	8.1600	0.99512	(0.00489)
12	8.1900	1.00368	0.00367
13	8.1200	0.99145	(0.00858)
14	8.0400	0.99015	(0.00990)
15	8.1900	1.01866	0.01848
16	8.0900	0.98779	(0.01229)
17	8.1200	1.00371	0.00370
18	8.1800	1.00739	0.00736
19	8.2700	1.01100	0.01094
20	8.1500	0.98549	(0.01462)

STANDARD DEVIATION (PER PERIOD) 0.889%
 ANNUALISED VOLATILITY (DAYS) 250 14.051%



3.2.2 Calculating historical volatility—adjustments

In calculating the historical volatility of certain assets the asset price sequence must be adjusted to reflect the non-homogenous nature of the data series. A major cause of this non-homogeneity is the entitlement to income on the underlying assets; for example, coupons in the case of debt instruments, and dividends in the case of equity stocks.

The presence of these income flows has the impact of reducing the comparability of succeeding price observations because the transition from cum-interest to ex-interest or cum-dividend to ex-dividend will affect the price of the asset. In theory, the asset price should fall by the amount of the income entitlement (coupon or dividend). This necessitates intervention to adjust the price sequence when an asset goes ex entitlement as follows:

The price relative is restated as $\ln(S_t + D) / S_{t-1}$.

Some practitioners eliminate data around the ex entitlement data, particularly for equities, from the data series reflecting the impact of a variety of factors, such as tax, on asset prices during these transitions and the resultant potential distortion in the volatility estimate.

3.2.3 Calculating historical volatility—asset class extensions

The above analysis focuses on the volatility of asset prices in general. This would encompass equity stocks, equity market indexes, commodity prices, commodity price indexes and currency values. In the case of debt instruments, it is feasible to calculate the volatility of both asset prices (being volatility of the price of the underlying debt security) and yield volatility (being the volatility of the interest rate index itself).

Exhibit 8.2 sets out calculation of both the price and yield volatility for a debt instrument (in this case a three month or 90 day discount instrument). In practice, the two are related with the interest rate changes driving the changes in the price of the debt instrument. *Exhibit 8.3* sets out the normal method for converting yield volatility to price volatility. Both yield based and price based measures of volatility are utilised for debt instruments. As described above, in theoretical terms, price volatility is proportional to absolute yield volatility and modified duration.

Exhibit 8.2
Volatility Estimation—Historical/Empirical Approach for Debt Instruments
CALCULATING VOLATILITY—DEBT INSTRUMENTS

EXHIBIT 8.2 VOLATILITY ESTIMATION—HISTORICAL/EMPIRICAL APPROACH FOR DEBT INSTRUMENTS								
PERIOD	INTEREST RATE	ASSET PRICE	PRICE RELATIVE $(S_t / (S_{t-1}))$	DAILY RETURN $U_t - \ln(S_t / S_{t-1})$	PRICE RELATIVE $(S_t / (S_{t-1}))$	DAILY RETURN $U_t - \ln(S_t / S_{t-1})$	YIELD VOLATILITY CALCULATIONS PRICE RELATIVE $(S_t / (S_{t-1}))$	DAILY RETURN $U_t - \ln(S_t / S_{t-1})$
0	6.000%	98.5421						
1	6.125%	98.5122	0.98970	(0.00030)	1.02083	0.02062		0.03608
2	6.350%	98.4584	0.98945	(0.00055)	1.03673	0.03608		(0.02391)
3	6.200%	98.4943	1.00036	0.00036	0.97638	0.00036		0.00803
4	6.250%	98.4823	0.98988	(0.00012)	1.00806	0.00806		0.02372
5	6.400%	98.4464	0.98954	(0.00036)	1.02400	0.02400		0.01550
6	6.500%	98.4225	0.98976	(0.00024)	1.01563	0.01563		0.01550
7	6.650%	98.3867	0.98984	(0.00036)	1.02400	0.02400		0.01550
8	6.750%	98.3629	0.98976	(0.00024)	1.01563	0.01563		0.01489
9	6.850%	98.3867	1.00024	0.00024	0.98519	0.00024		(0.01058)
10	6.580%	98.4034	1.00017	0.00017	0.98947	0.00017		0.00303
11	6.800%	98.3987	0.98995	(0.00005)	1.00304	0.00304		0.01663
12	6.710%	98.3724	0.98973	(0.00027)	1.01667	0.01667		(0.00898)
13	6.850%	98.3867	1.00015	0.00015	0.98106	0.00015		0.00749
14	6.700%	98.3748	0.98988	(0.00012)	1.00752	0.00752		0.00298
15	6.720%	98.3700	0.98995	(0.00005)	1.00298	0.00298		0.00287
16	6.740%	98.3653	0.98985	(0.00005)	1.00593	0.00593		(0.00592)
17	6.780%	98.3557	0.98990	(0.00010)	0.98410	0.00010		(0.01044)
18	6.740%	98.3653	1.00010	0.00010	0.98961	0.00010		0.00449
19	6.870%	98.3820	1.00017	0.00017	1.00450	0.00450		
20	6.700%	98.3748	0.98983	(0.00007)				1.468%
STANDARD DEVIATION (PER PERIOD)				0.023%				0.000%
ANNUALISED VOLATILITY (DAYS)		250.0000		0.000%				0.000%

Exhibit 8.3**Relationship Between Price and Yield Volatility**

The following formula can be used to convert yield volatility into its price volatility equivalent:

$$\text{Price Volatility} = (\Delta \text{ Price} / \Delta \text{ Yield}) \times \text{Yield} \times \text{Yield Volatility}$$

For example, the formula can be used as follows to convert yield volatility of 20% for 91 day securities to its equivalent price volatility as follows:

$$(\Delta \text{ Price} / \Delta \text{ Yield}) = 24.08 \text{ for } 0.0001\% \text{ pa or } 1 \text{ bps change in yield (per } \$1,000,000 \text{ face value of the security at a yield of } 7.00\% \text{ pa)}$$

Therefore:

$$\text{Price Volatility} = 24.08 \times 7.00\% \times 20.00\% = 0.337\% \text{ pa}$$

An observable feature of the relationship between yield and price volatility is that yield volatility increases in a market with decreasing yields whereas price volatility decreases (for similar movements in outright yield).

In practice, both types of yields are utilised with the preferred volatility estimate parameter being driven, largely, by market convention. In utilising yield volatility estimates, the following points should be noted:

- Yield volatility is usually assumed to be constant for fixed interest instruments of the same yield implying a flat yield curve which does not change shape and trades at constant yield volatility across all maturities.
- The use of yield volatility has the potential to create confusion where the yield curve shape is, itself, volatile.
- Yield volatility is not affected by the changing duration associated with fixed coupon bonds.

Utilising price volatility estimates for debt instruments, the following points should be noted:

- The price volatility constantly changes with changing duration of the underlying fixed interest instrument.
- It is necessary to clarify where the price volatility being calculated are the basis of a "clean" (ex interest) or "dirty" (cum accrued interest) price.

In practice, price volatility is utilised in markets where the underlying security is traded in price (for example, the US treasury bond market). In addition, yield volatility is utilised in preference to price volatility in a variety of markets for options on short term interest rates because price volatility of these instruments is very low.

3.2.4 Considerations in using historical/empirical volatility

The historical/empirical techniques for volatility estimation seeks to quantify past market volatility and utilise this as a basis for forecasting *future* market volatility for asset prices.

The major difficulties with this approach include:

- assumed stationarity of volatility parameters;

- specifying the number of observations utilised;
- the availability of a variety of price observations;
- specifying the number of days utilised; and
- allocating relative importance to different components of the time series.

A major difficulty relates to the assumption implied by this technique that past volatility is a useful mechanism for deriving future asset price volatility. The assumed stationarity of the volatility estimate is neither logical nor supported by empirical evidence. Volatilities for a variety of assets demonstrate significant changes over time.

The period over which data is utilised to generate historical volatility is constantly debated. Proponents of utilising data over a very long period (say, five years) implicitly assume that volatility is constant over long periods of time or, alternatively, tends towards a quantifiable average level of volatility. This is consistent with a hypothesis regarding the mean reverting nature of volatility.

Proponents of utilising data over a shorter period (between one and three months) base their position on the implied view that volatility itself is not constant but varies significantly and prefer to use a shorter period to obtain a good estimate of the current level of volatility. Adherents to this theory would, as a consequence, adjust the volatility parameter input into option pricing formulas regularly.

It is clear that the larger number of observations utilised to estimate volatility, the higher the probability that changes in, for example, general economic conditions or other exogenous factors, will impact upon the calculated volatility causing a violation of the assumption that the standard deviation of the asset's return is constant over the life of the option. The trade-off between increasing the period over which data is utilised in order to achieve more efficient estimates and the probability that the volatility has altered is essentially not resolvable within the context of the mathematical option pricing models discussed.

A further complication arises from the fact that the price data utilised can take a variety of forms, including:

- close-to-close prices;
- open-to-close prices;
- close-to-open prices; and
- high or low prices.

The use of close-to-close prices is the most common measure utilised. This measure allows the full daily movement of the asset price to be captured and the impact of all information released over the relevant 24 hour period to be captured. The impact of non-trading days (such as holidays and weekends) tends to distort the data series as information may be released and impact upon prices but is not captured until a later date. The process of annualisation utilised (see above) does not adjust for this phenomenon.

Open-to-close prices provide a measure of intra-day volatility, which shows reaction to all information released during the trading day. Use of open-to-close prices creates problems of annualisation as it is necessary, to

be accurate, to adjust for the concentration of information released within an average 24 hour period.

Use of close-to-open prices allows a measure of overnight volatility. It shows reaction to information release outside trading hours and, in certain cases, measures the interrelationship between the domestic market and international market in other time zones. The close-to-open price series suffers from the same difficulties as the open-to-close price series because of difficulties of annualisation.

High to low prices again facilitate capturing of the range of a full day's price movements in the volatility estimate. It is typically useful to traders who have intra-day positions or are hedging positions intra-day.

The major value of the variety of price series and the different estimates of volatility that can be derived lies in the capacity to compare the relative price volatility of the various series.

As noted above, the daily volatility is scaled by the square root of the time to maturity to derive the annualised volatility. One commentator⁴ argues that there are at least three relevant choices as to the maturity estimate:

1. *Calendar days*—the number of actual calendar days between the time of valuation and option maturity.
2. *Trading days*—the number of days over the option life on which trading in the relevant asset is open.
3. *Economic days*—the number of days over the option life on which information *likely to impact the asset price* is released.

The last concept requires the information release to be *anticipated*. As noted previously, the asset price seems to react to both anticipated and unanticipated information. In addition, the use of economic days would necessarily require an understanding of and specific identification of information *likely* to affect asset prices.⁵

The distinction between calendar and trading days is more interesting. The volatility estimate utilised for short dated options is particularly problematic. For example, in the case of an option with a time to expiry of, say, seven days where there are intervening non-trading days (a weekend and perhaps a holiday), the proportion of non-trading days as a proportion of the life of the option is significant. Similarly, for a short dated option, there may be *no non-trading days*. In these cases, the use of the same volatility scaled over the *calendar time to maturity* appears inappropriate. In practice, some practitioners adjust the volatility of these option to reflect this factor. A common technique is to adjust the volatility used (see *Exhibit 8.4*).

4. See Kenneth Leong, "Exorcising the Demon" (1990) 3 (9) (Oct) *Risk* 29.

5. For discussion of research on economic days, see Galen Burghardt and Gerald A Hanweck Jr, "Calendar-Adjusted Volatilities" (1993) (Winter) *The Journal of Derivatives* 23.

Exhibit 8.4 Adjusting Volatility

Assume an option with seven days to expiry. Of these seven days, four are non-trading days. The market volatility on an annualised basis is 15.00% pa. The volatility of this option is rebased as based as follows to adjust for the short maturity:

1. Rebase the annual volatility to a daily basis:

$$.15 / \sqrt{250} = 0.009487$$

2. Re-annualise the daily volatility to an annualised basis adjusted for the lower number of trading days (three out of seven or 42.86%). The adjustment is based on the fact that normally the annualisation uses 250 days in a year (or 68.49% of the calendar year). Therefore, the daily volatility is scaled by 156.43 days (42.86% of 365 days).

$$.009487 \sqrt{156.43} = .118655$$

The adjusted volatility used would be 11.8655% pa.

If the option had five days to expiry, all of which were trading days, then the adjustment would be to rebase it using 365 days (100% trading days) as follows:

$$.009487 \sqrt{365} = .1812 \text{ or } 18.12\% \text{ pa}$$

An additional problem relates to the impact of high value price changes in the calculation of historical volatility estimates. This problem may also be stated as the problem of allocating relative importance to the sequence of data.

Consider an unweighted set of data which contains a single large price change. If that data point is removed, for example where a trailing average for a fixed number of days of price observations is utilised to calculate the volatility estimate, then the removal of the particular data point will have the impact of substantially altering the volatility either up or down. *Exhibit 8.5* sets out an example of this phenomenon.

Exhibit 8.5
Impact of Changes in Data Series on Volatility Estimate

CALCULATING VOLATILITY

Period	Asset Price	Price Relative (St/(St-1))	Daily Return $U_i =$ $\ln (St/(St - 1))$
0	9.6500		
1	8.2500	0.85492	(0.15674)
2	8.1800	0.99152	(0.00852)
3	8.2000	1.00244	0.00244
4	8.2000	1.00000	0.00000
5	8.2100	1.00122	0.00122
6	8.1200	0.98904	(0.01102)
7	8.1500	1.00369	0.00369
8	8.0700	0.99018	(0.00986)
9	8.1300	1.00743	0.00741
10	8.1700	1.00492	0.00491
11	8.2000	1.00367	0.00367
12	8.1600	0.99512	(0.00489)
13	8.1900	1.00368	0.00367
14	8.1200	0.99145	(0.00858)
15	8.0800	0.99507	(0.00494)
16	8.1900	1.01361	0.01352
17	8.0900	0.98779	(0.01229)
18	8.1200	1.00371	0.00370
19	8.1800	1.00739	0.00736
20	8.2700	1.01100	0.01094

STANDARD DEVIATION (PER PERIOD)

3.587%

ANNUALISED VOLATILITY (DAYS)

250

56.708%





Exhibit 8.5—continued
CALCULATING VOLATILITY

Period	Asset Price	Price Relative ($S_t/(S_{t-1})$)	Daily Return $U_t =$ $\ln(S_t/(S_{t-1}))$
0	8.2500		
1	8.1800	0.99152	(0.00852)
2	8.2000	1.00244	0.00244
3	8.2000	1.00000	0.00000
4	8.2100	1.00122	0.00122
5	8.1200	0.98904	(0.01102)
6	8.1500	1.00369	0.00369
7	8.0700	0.99018	(0.00986)
8	8.1300	1.00743	0.00741
9	8.1700	1.00492	0.00491
10	8.2000	1.00367	0.00367
11	8.1600	0.99512	(0.00489)
12	8.1900	1.00368	0.00367
13	8.1200	0.99145	(0.00858)
14	8.0800	0.99507	(0.00494)
15	8.1900	1.01361	0.01352
16	8.0900	0.98779	(0.01229)
17	8.1200	1.00371	0.00370
18	8.1800	1.00739	0.00736
19	8.2700	1.01100	0.01094
20	8.2400	0.99637	(0.00363)

STANDARD DEVIATION (PER PERIOD) 0.752%
ANNUALISED VOLATILITY (DAYS) 250 11.892%

Asset Prices



This problem, together with the desire to give greater weight to more recent data, has led to a number of weighting schemes being proposed. One example of this is the exponential weighting scheme proposed by J P Morgan in its RiskMetrics™ models. This approach uses a decay factor λ (equal to 0.94) to give greater weight to more recent data. Under this approach, the

forecast from the exponential estimator for the variance of returns over the next days for a given set of T returns is given as:

$$\sigma^2 = (1\lambda) \sum_{t=1}^T \lambda^{t-1} (R_t - \bar{R})^2$$

Where R_t is the daily return at time t

The decay factor is chosen to minimise the error between actual observed volatility and its forecast over the sample period utilised.⁶

3.3 Implied volatility

3.3.1 Calculation of implied volatility

The implied volatility approach calculates volatility implied by the current market value of options. This is undertaken by specifying the option price and calculating the volatility which would be needed in a mathematical option pricing formula, such as Black-Scholes to derive the specified market price as a fair value of the option. *Exhibit 8.6* sets out an example of the calculation of implied volatility.

6. For a detailed discussion of the exponential scheme, see J P Morgan Reuters, (1995) *RiskMetric™—Technical Document* (4th ed, J P Morgan Reuters Ltd, New York, 1996) at Chapter 5.

Exhibit 8.6
Volatility Estimation—Implied Volatility

MODEL OUTPUTS	CALL	PUT
OPTION PREMIUM	3.73	8.38
OPTION PREMIUM (% OF ASSET PRICE)	3.73%	8.38%

PRICING INPUTS		
UNDERLYING ASSET PRICE	100.00	
STRIKE PRICE	110.00	
TRADE DATE	1-Jan-X6	
EXPIRY DATE	1-Jul-X6	
RISK FREE RATE	10.00%	
CALL(0)/PUT (1)	0	1
	Call	Put
OPTION PREMIUM	4.50	9.00
OPTION PREMIUM (% OF ASSET PRICE)	4.50%	9.00%

IMPLIED OPTION VOLATILITY (% PA)	22.80%	22.26%
---	---------------	---------------

The calculation of implied volatility is usually done using an iterative procedure. *Exhibit 8.7* sets out two examples of iterative procedures commonly utilised.

Exhibit 8.7**Iterative Procedures for Calculation of Implied Volatility**

The following are two common methods used to calculate implied volatility:

Bisection method

The process utilised is as follows:

1. A low estimate for volatility is used to determine an option value.
2. A high estimate for volatility is used to calculate a second option value.
3. The next estimate is determined by an interpolation procedure:

$$\sigma_{low} + (P - P_{low}) \times [(\sigma_{high} - \sigma_{low}) / (P_{high} - P_{low})]$$

Where

σ_{high} ; σ_{low} is equal to the high or low volatility estimate

P ; P_{high} ; P_{low} is equal to the actual observed option premium, option premium for the high and low volatility estimate

4. If the interpolated option value is below (above) the actual option premium, then the procedure is repeated with the low (high) volatility estimate with the interpolated estimate.
5. The procedure entailing the above steps is repeated until the volatility estimate corresponds to the actual price of the option.

Newton Raphson method

This procedure is as follows:

1. A reasonable estimate for the implied volatility is used to calculate the option premium.
2. If the option premium does not correspond to the actual observed option price, then the first estimate is adjusted as follows:

$$(P - P_{\text{first estimate}}) / (\delta C / \delta \sigma)$$

Where

$P - P_{\text{first estimate}}$ is the actual observed premium and the premium calculated using the first volatility estimate

$\delta C / \delta \sigma$ is the derivative of the option formula with respect to volatility evaluated at the first estimate of volatility:

$$S \times \sqrt{T} (1 / \sqrt{2\pi}) e^{-d_1^2/2}$$

Where

C = Price of a call option

S = Asset Price

$D = d_1$ in Black-Scholes Option Pricing Formula (refer Chapter 5)

3. The volatility estimate generated is then used to recalculate the option premium.
4. The process is repeated until the implied volatility corresponding to the observed market premium is derived.

Source: Mark Kritzman, *The Portable Financial Analyst* (Probus Publishing, Chicago, Illinois, 1995), pp 113-122.

3.3.2 Considerations in using implied volatility

The use of implied volatility is generally deficient in that it intrinsically sanctions a circular process. The volatility implied in options currently trading, which measures the volatility level required to clear the market at a given point in time, is treated as being the true constant asset price volatility parameter.

There are additional technical difficulties:

- A major difficulty with implied volatilities is that options with different strikes with the same maturity often demonstrate different implied volatilities (the so-called “smile” effect which is discussed below).
- Where options of the relevant type or maturity are not traded, this technique is unavailable.

The major value of implied volatility techniques as a method of volatility estimation is that it provides an observable measure of the relevant option market expectations as to volatility.

3.4 Volatility modelling

Historical and implied volatility assume that volatility is stable in that changes in volatility are unpredictable in that volatility *changes* are uncorrelated with previous changes in volatility.

This condition would be satisfied if a regression of the squared value of the differences of the asset price changes or returns and the mean of the time series (the error squared) as at time $t(n)$ and the error squared as at time $t(n-1)$ (the previous observation) showed no significant relationship. This lack of relationship would be evident in:

1. the β or slope of the regression line should not be significantly different from 0;
2. the intercept of the regression line should approximate the average value of the errors squared; and
3. the residuals around the fitted values equal the differences between the actual values for the errors squared and the predicted values from the regression are randomly distributed around a zero expected value.

If these conditions were satisfied the residuals would be described as homoscedastic (that is, they were serially independent).

In practice, the residuals do not, in fact, satisfy the above conditions; that is, they are heteroscedastic. This is evident in the following:

1. the regression coefficients may or may not be significant; but
2. the errors squared are related to prior values (for example, there are clusters of positive as well as negative residuals (where the regression underestimates (overestimates) the actual errors squared));
3. the errors squared are related in usually a non-linear relationship.

The underlying logic of this approach is that of *volatility clustering* (the high value of the errors squared occur in clusters). Volatility clustering which is observable in financial markets suggests that volatility follows a predictable pattern. Large asset price changes seem to be succeeded by a

sequence of *further large changes*. This pattern causes volatility to be high after large asset price movements. Increasingly, a variety of statistical/econometric techniques are being applied to volatility estimation. One such technique is known as ARCH (an acronym for Auto Regressive Conditional Heteroskedasticity). A number of variations on ARCH techniques, representing extensions of the basic model, are also increasingly being utilised. A detailed analysis of estimating volatility with these types of models is set out in Chapter 9. In this Chapter a brief overview only is presented.

The basic insight underlying ARCH models is the concept that volatility follows clear patterns. Central to this approach is that the volatility of an asset today depends on the volatility of the asset yesterday and the "shock" in the price of the asset yesterday. A central tenet in this approach is that the intertemporal link in volatility changes over time is relatively constant or stationary. This implies that volatility changes are predictable on the basis of historical volatility. This approach is usually allied to an assumption that volatility regresses towards long-term long run means (that is, it shows a basic mean reversion tendency). ARCH models imply that the best estimate of volatility is not the volatility of the asset *today*. These models show how a change in volatility persists and decays gradually. For example, an increase in asset price leads to an underlying increase in asset volatility with a gradual decrease towards a mean level.

The ARCH models are predicated on correcting for the detected non-linearity through regressing the residuals at time $t(n)$ on the errors squared as at time $t2$. The coefficients of this second regression are then used to adjust (by adding) to the coefficients of the original regression. This assumes that the variance (the average value of errors squared) is *conditional on the heteroskedasticity*. *Exhibit 8.8* sets out the basic procedure for deriving ARCH estimates.

Exhibit 8.8 Arch Procedures

The ARCH methodology requires the following procedure:

1. The observed means are subtracted from their mean.
2. The difference calculated in the previous step (step 1) is squared to calculate the errors squared.
3. The errors squared as at time $t(n)$ are regressed against the errors squared as at time $t(n-1)$ (regression 1).
4. The fitted errors squared as at time $t(n)$ are subtracted from the observed errors squared as at time $t(n-1)$.
5. The residuals in the previous step as at time $t(n)$ are regressed on the errors squared as at time $t(n-1)$ (regression 2). Under the generalised least square approach, both sides of the regression equation should be divided by the fitted values from regression 1.
6. The regression coefficients (intercept (α) and slope (β)) from the regression 2 are added to the coefficients from regression 1 (regression 3).

The adjusted regression equation (regression 3) provides a more efficient predictor of variance than the original regression equation. The increased efficiency is only to the degree that the residuals from the original model are heteroskedastic.

Source: Mark Kritzman, *The Portable Financial Analyst* (Probus Publishing, Chicago, Illinois, 1995), pp 123-129.

An important element of the ARCH model is that it more readily explains “fat tailed” and leptokurtic distributions of asset price changes. The major applications of ARCH models, have, to date, been modelling correlations between assets and forecasting volatility.

3.5 Risk reversal volatility

Classical volatility estimates are to a degree non-directional. However, there is increasing interest in using observed market volatilities as a mechanism for inferring information about the future direction of volatility and also of the spot asset price. One such technique is that of risk reversal volatilities, which are used for these purposes.

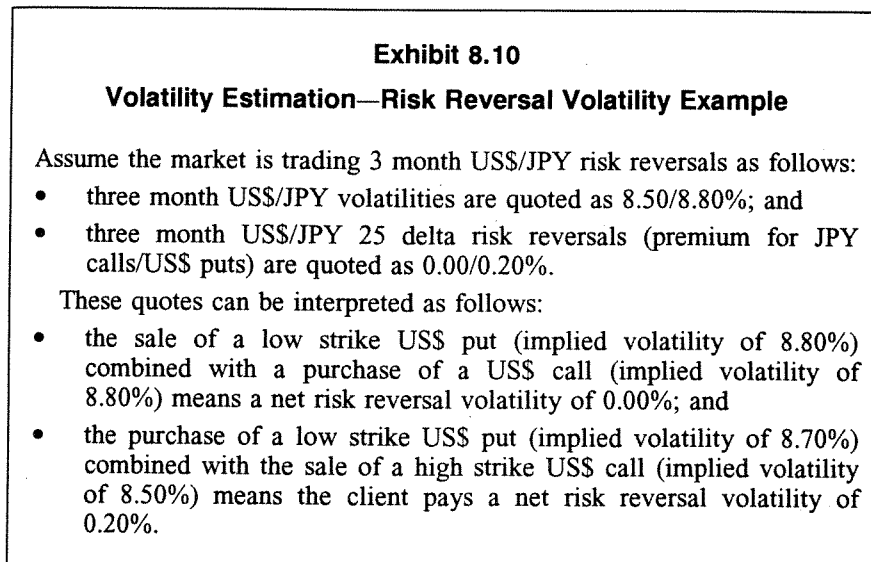

Risk reversal volatilities use implied volatilities as a mechanism for deriving the market’s view of the future path and volatility of asset prices. In itself, it is not a mechanism for deriving volatility estimates but is used in parallel with other techniques to provide additional information to calibrate volatility estimates. In particular, it provides valuable information regarding the pattern of volatility, in particular, the volatility smile (which is discussed below).

Risk reversals may be defined as the following transactions:

1. the purchase of an out-of-the-money call with the simultaneous sale of an out-of-the-money put; or
2. the purchase of an out-of-the-money put with the simultaneous sale of an out-of-the-money call.

Typically, both transactions are done, with the two options both usually having the same expiration date and being of equal size.

Transaction 1 above yields a view of the market which is biased to increases in the asset price, while transaction 2 evidences a view of the market which is biased to decreases in the asset price. This reflects the economic biases evident in the transactions. *Exhibit 8.9* sets out a diagrammatic view of the transaction using the payoff profiles. *Exhibit 8.10* sets out an example of risk reversal volatilities.



The demand and supply of risk reversals supplies valuable information on the following:

- the *expected* movement in implied volatility; and
- the *expected* movement in the spot.

The interpretation of the risk reversal volatilities is usually done within the following format which specifies the preferred trades for particular views of asset price and asset volatility:

Expectation As To	Decrease	Static	Increase
Asset Volatility			
Expectation As To			
Asset Price			
Decrease	Sell calls	Sell asset	Buy puts
Static	Sell calls & puts	Sell calls & puts	Buy calls & puts
Increase	Sell puts	Buy asset	Buy calls

3.6 Alternative approaches

One researcher has suggested an alternative approach to the estimation of volatility based on the concept of conservation of volatility.⁷

The basic premise of the theory is that actual observation reveals that there are days with high volatility and there are days with low volatility. Any estimate or forecast of volatility is essentially an average of these two elements based on historical data. Utilising this approach, volatility is defined as follows:

$$\sigma = w_1\sigma_n + w_2\sigma_h$$

Where

σ = volatility overall

w_1 = fraction of days of normal volatility

σ_n = normal volatility

w_2 = fraction of days of high volatility

σ_h = high volatility

This concept is similar to the concept of economic days identified above. This is relevant insofar as it would be expected that days when economic information is released would generally be volatile days.

This approach is used to estimate volatility as follows:

1. Existing historical price changes are segmented into two groups—high and normal volatility. This is done using a filter such as changes above a certain threshold level.
2. The volatility for each series is calculated normally.
3. The estimate for volatility is calculated using the traders estimate of the number of days of normal versus high volatility in the period to the expiry of the option. The weights are calculated using this scheme and applied to the historical volatility estimates derived for the respective series.

7. See Robert Tompkins, *Options Explained* (MacMillan Press, England, 1994), Ch 5.

4. BEHAVIOUR OF VOLATILITY


4.1 Key issues

Two aspects of the behaviour of volatility estimates require comment:

- the volatility smile—that is, the behaviour of volatility relative to the strike price or yield of the option; and
- the term structure of volatility—that is, the behaviour of volatility relative to time to expiry.

4.2 The volatility smile

In practice, at-the-money options, generally, are observed to trade with lower implied volatilities relative to out-of-the-money options and to a lesser extent in-the-money options. This phenomenon is described as the volatility smile. *Exhibit 8.11* sets out an example of a typical volatility which highlights both the higher volatilities for in and out-of-the-money options and the skew in the smile (usually, in favour of out-of-the money options).



The volatility smile reflects a variety of factors, including:

- adjustments for the distributional assumptions underlying standard option pricing models;
- directional assumptions regarding the movement in the underlying asset prices which are incorporated into the option volatility and price;

- clientele effects and the demand for out-of-the-money options;
- the management of option hedging risks by traders; and
- liquidity effects.

The volatility smile is capable of explanation in terms of the deviation of observed asset price movements from the assumed log normal distribution. In practice, asset price movements seem to be characterised by the following:

- The market distribution of asset prices changes appears to demonstrate fat tails (statistically described as the kurtosis of the distribution). This type of distribution is characterised by more larger value price changes than is consistent with a normal distribution.
- The fat tails are consistent with the presence of “jump” risk, that is, non-stochastic (or discontinuous) changes or movements in the price of the asset which cause deviation from the assumption of a normal distribution.

The *actual* observed pattern of price changes because of the above characteristics would systematically underestimate the value of deep in and out-of-the-money options. This reflects the fact that the log normal distribution *systematically* underestimates the expected values that the option may take at maturity in either tail of the distribution. The volatility smile is consistent with trader behaviour which seeks to adjust the option premium for these deficiencies in an option pricing model such as Black-Scholes.⁸ This adjustment is effected through an increase in the volatility for both deep in and out-of-the-money options to equate the premium received to the *expected* payouts under the option incorporating the *true* asset price change distribution.

The volatility smile, particularly the skew in the structure of the smile, may reflect expectations regarding the expected direction of price movements which are incorporated in the option price and by implication the implied volatility. For example, if the US\$/JPY is expected to decline from its current level of 110, then US\$ puts/JPY calls may be more valuable and US\$ calls/JPY puts with a strike price at or above the spot rate may be correspondingly less valuable. This directional view may be reflected in option price which will be higher than in the absence of this expectation and reflected in the implied volatilities.

The higher price and implied volatilities can be considered to be the higher expected economic cost of hedging or dynamically replicating the option. The smile and the skew are also consistent with the inherent nature of log normal distributions which have a natural skew to the right hand tail, implying a higher probability of a rise than a fall in the asset price. Implied volatilities, if it is sought to adjust for the skew, should be higher for options with strike prices below the implied forward rate (assumed to be the mean of the distribution) than for options with strike prices above the implied forward prices to correct for the natural bias in prices.

The market for options with different strike prices appears to exhibit significant biases in demand and supply (a clientele effect). Out-of-the-money

8. This approach is suggested by Fischer Black in “How to Use the Holes in Black-Scholes” (1989) 1 (4) *Continental Bank Journal of Applied Corporate Finance* 59.

options are attractive vehicles for speculative investment demand, reflecting the following factors:

- the gearing or leverage of the out-of-the-money options (expressed as the asset price divided by the option premium) is higher; and
- the low absolute cash investment entailed in the purchase of the option.

The presence of these factors dictates significant demand for these options. The supply for these types of options is constrained by the fact that option traders are reluctant to sell/write these out-of-the-money options because of the difficulty of hedging or replicating these options in the event of a jump in the asset price (high gamma risk).

In contrast, the position for in-the-money options is influenced by different factors. The dominating characteristic of these types of options is that they have a high delta and move closely with movements with the underlying asset prices. This allows in-the-money options to be used as a direct substitute for the asset itself. The primary demand for these options is from participants such as traders who, in replicating options through the process of delta hedging, need to trade in the underlying asset. The high delta of these options may enable these to be substituted for the asset in the replication process. This has the effect of lowering the financing costs in replicating the option synthetically. Similarly, other traders or participants seeking to synthesise positions in the asset at lower cost may find these in-the-money options attractive.

The supply of these options is limited. This reflects the reluctance to write a deep in-the-money option because in the absence of a large or extreme price movement the option will be exercised, requiring the seller to buy or sell the asset at a price which is disadvantageous to them. In addition, such options do not have significant time or volatility value, further reducing their attractiveness to the seller.

The interaction of supply and demand for these deep in-the-money options results in the option price and implied volatilities being bid up above comparable volatilities for at-the-money options of the same maturity.

The volatility smile also appears to incorporate the impact of traders seeking to manage the risk of option transactions. Traders seek to replicate options through a process of trading in the underlying asset. This approach to option portfolio management is consistent with standard option pricing models, such as Black-Scholes, which derive the fair value of the option as the cost of dynamically hedging a short position in the option through a position in the asset which continuously adjusted including the funding cost of the position. This process referred to as delta hedging or dynamic option replication is discussed in detail in Chapter 11.

The process of option replication, in the manner described, is, in practice, not free from risk. The primary risks are the possibility of a sudden gap or jump in asset prices (a large value price change) which requires a substantial adjustment to the position in the asset to be effected. This reflects the fact that the option delta has changed significantly. The option writer is also exposed to changes in volatility which will cause changes in both the value of the options and its delta requiring rehedging through trading in the asset. These risks usually referred to the gamma and vega risk cannot be hedged

through trading in the asset. This is because the asset itself has low or no gamma (that is, the convexity of the asset price is significantly lower than the convexity of the option) and no vega risk. Gamma and vega risk can only be offset by trading in an option on the underlying asset.

In practice, the problems in hedging require traders to trade in options to manage the risks in their option portfolios. Traders typically trade in short-dated at-the-money options to manage their gamma and to a lesser extent vega risks. This reflects the fact that gamma and vega are at their highest level for short-dated at-the-money options. This creates a significant demand for these options as traders rebalance their hedges frequently in response to market movements in the asset price and volatility. The supply for these options is also strong reflecting the high volatility or time value of these at-the-money options, which makes them attractive for sellers.


The combination of the above factors results in differential liquidity of options with different strike prices for a given maturity. The volatility of at-the-money options is lower reflecting the higher liquidity of these options from the greater balance between supply and demand for these options. In and out-of-the-money are less frequently traded and the imbalance of demand relative to supply is reflected in the higher implied volatility relative to the at-the-money options. The resultant volatility smile is, as noted above, often skewed with out-of-the-money options demonstrating higher volatilities than both at and in-the-money option. The volatility smile also appears to diminish with maturity reflecting the reduced impact of the factors identified.

4.3 Term structure of volatility

The term structure of volatility encompasses two separate issues:

1. the relationship between volatility and the time to expiry of the option;
and
2. the pattern of forward volatilities.

In general terms, volatility for shorter time to expiry option is higher than the volatility for options with longer times to expiry. This usual term structure of volatility is set out in *Exhibit 8.12*.




This pattern of decreasing volatility relative to option maturity essentially reflects:

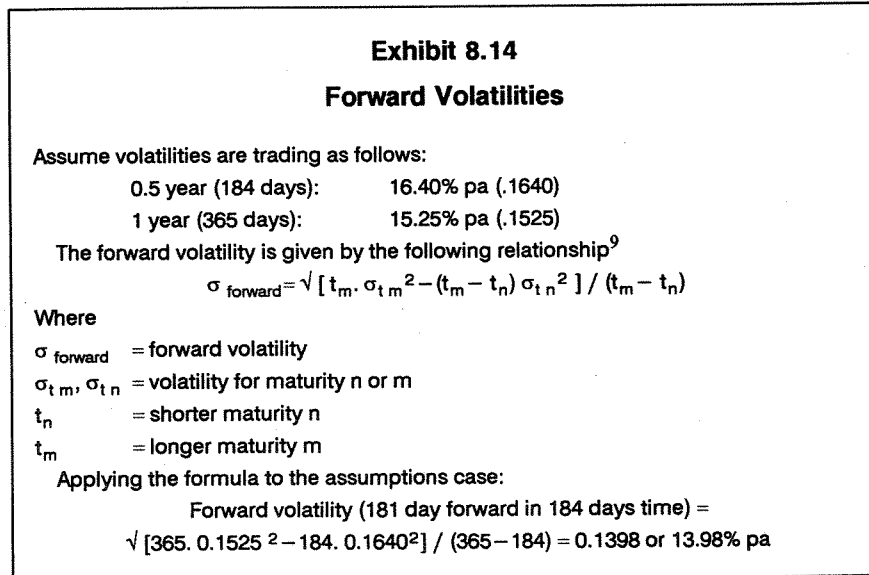
- the expectations of larger price movements in the very near future which drive the implied volatility levels to higher levels for short-dated options;
- the proportionately larger impact on option values of large asset price changes in the asset price on short-dated options and the relatively higher risk to the seller of these options, which must be compensated for through higher premiums and higher implied volatilities; and
- the mean reversion nature of volatility which seems to fall (rise) from high (low) absolute level towards a long run mean level.

There are two additional aspects of the term structure of volatilities: the volatility cone and the concept of forward volatilities.

The concept of the volatility cone is based on the fact that in projecting volatility over the life of the option the trader may seek to project, based on history, the highest and lowest volatility that are likely to occur to ensure that the estimated volatility is unlikely to be exceeded (assuming the option is sold). This entails the projection of maximum, minimum, average and (optional) historical implied volatilities for at-the-money options for different maturities. A sample volatility cone is set out in *Exhibit 8.13*. The trader would seek to use a volatility estimate that is favourable within the boundaries indicated by the cone.



The concept of forward volatilities is analogous to the concept of forward asset prices such as forward interest rates. It focuses on the implied forward volatility calculated as the forward variance. *Exhibit 8.14* sets out an example of the calculation of forward volatilities.




9. See Tompkins, op cit n 7, p 191.

4.6 Volatility surfaces

It is increasingly common to combine the patterns of volatility—the smile and term structure—into a volatility surface. *Exhibit 8.15* sets out an example of a volatility surface. The surface is typically generated for the purpose of valuation of portfolios of options which may contain a range of options with substantially different strike prices and maturities. The surface is derived from observed implied volatilities or specified volatilities obtained from other sources with interpolation procedures used to complete points on the surface which are either unavailable or not traded. The volatility estimates generated from the surface are then utilised to determine the value of the options in the portfolio.

Exhibit 8.15
Volatility Surface



4.7 Option pricing models incorporating stochastic volatility

A number of option pricing models have been developed to incorporate the impact of the volatility smile and term structure.¹⁰ These models effectively introduce a stochastic volatility term into the option pricing process. This replaces the traditional constant volatility term. The stochastic volatility term is usually a function of the strike price and the term of the option. The volatility function is derived from available market prices for options, and is calculated by calculating the implied option volatility using an option pricing model. This allows the construction of an option pricing model which uses and is consistent with historical volatilities structures.

5. VOLATILITY ESTIMATION IN PRACTICE

Volatility estimation in practice is a complex activity; it requires an understanding of the following:

- the various volatility estimates available and their significance; and
- behaviour of volatility parameters.

In practice, market participants take into consideration the following types of volatility:

- implied versus historical volatility;
- implied versus predicted volatility; and
- implied or historical volatility versus actual volatility.

The difference between implied and historical volatility is useful in predicting potential changes in volatility, while differences between implied versus predicted volatility determine a variety of trading strategies utilised by market participants. The relationship of the implied historical volatility versus actual realised volatility provides the basis for adjusting expectations of future asset price volatility.

In practice, the attitude to volatility estimation and the relative importance assigned to each of these methodologies will depend, to some extent, on the activities of the market participant and the risks sought to be managed. For example, an intra-day hedger or trader will take a different view to a participant seeking to hedge longer-term positions. Similarly, volatility estimation for hedging short versus long-dated option may as will generation of volatilities for the purposes of risk management as distinct from option valuation.

The volatility estimates available are, in practice, influenced by the underlying liquidity in assets. The lack of liquidity, signified by wider bid-offer spreads for particular securities, will generally be reflected in the pricing of options (that is, the implied volatility).

10. See Bruno Dupire, "Pricing with a Smile" (1994) 7 (1) *Risk* 18; Emanuel Derman and Iraj Kani, "Riding on a Smile" (1994) 7 (2) *Risk* 32.

Chapter 9

Estimating and Forecasting Volatility and Correlation Using ARCH and GARCH Models

by Carol Alexander

1. INTRODUCTION

Estimates and forecasts of volatility and correlation are one of the cornerstones of quantitative financial analysis. They are needed for risk management, investment analysis, capital allocation, trading, pricing and hedging—sometimes in the guise of covariance matrices.¹ In many cases these forecasts will be the only stochastic parameters in the model (as, for example, in certain Value-at-risk models) so their accuracy is crucial to the success of the analysis.

Two methods for generating volatility and correlation estimates and forecasts for financial returns are in common use: moving average methods and GARCH. There are also some more sophisticated mathematical techniques which can be applied (for example, neural networks) but these are beyond the scope of this chapter.² Moving average methods are standard statistical estimation techniques, and the estimates generated are taken as forecasts in financial applications. On the other hand GARCH models provide current estimates of volatility and correlation, that are then used to generate distinct forecasts of the whole term structure.

In this chapter the mathematical methods used to generate moving average and GARCH volatility and correlation are described, and the advantages and limitations of each method are explained. This is followed by a short survey of some of the most common applications of volatility and correlation to financial markets.

1. The covariance matrix is a square array of numbers with variances along the diagonal and covariances on the off diagonal. Volatility and correlation are just standardised forms of variance and covariance: In particular, if V is the variance of h -day returns, then h -day volatility is $100\sqrt{250V}/h\%$, where we have assumed there are 250 trading days per year. If V_1 and V_2 are the variances of two h -day returns and COV is their covariance, then their correlation is $COV/\sqrt{V_1V_2}$. Although covariance measures the same thing as correlation—degree of co-movement between two (jointly stationary) returns series—it depends on the units of measurement. Correlation is just a standardised form of covariance, so that it always lies between -1 and $+1$ and so that a correlation near zero is indeed insignificant.
2. C O Alexander and P M Williams, "Term Structure Forecasts of Foreign Exchange Volatility and Kurtosis: A Comparison of Neural Network and GARCH Methods" (1997).

2. MOVING AVERAGE METHODS

A moving average is an average taken over a rolling window of a fixed number of data points. Thus the average is first calculated on data points x_1, x_2, \dots, x_n , then on data points x_2, x_3, \dots, x_{n+1} and so on. Each time the window is rolled, one point is knocked off behind and another is added at the end, so that the sample size remains fixed. Recording the average in this way creates a new time series which begins at time period n of the original time series.

To generate variance estimates using moving averages, it is usual to apply the average to squared returns r_t^2 ($t = 1, 2, 3, \dots$) or squared mean deviations of returns $(r_t - \bar{r})^2$ where \bar{r} is the average return over the data window. Although standard statistical estimates of variance are based on mean deviations,³ empirical research on the accuracy of variance forecasts in financial markets has shown that it is often better not to use mean deviations of returns, but to base variances on squared returns and covariances on cross products of returns.⁴

2.1 "Historic" volatility and correlation


The n -period "historic" estimate of variance at time T is based on an equally weighted moving average of the n past one-period returns squared:

$$\hat{\sigma}_T^2 = \sum_{t=T-n}^{t=T-1} r_t^2 / n$$

The square root of this (that is, the standard deviation) is then converted into an annualised percentage in the usual way, to obtain the historic volatility.⁵ Traditionally, an n -day historic volatility estimate has been used as a volatility forecast over the next n days—for example, to price an option which matures in n days time. The rationale for this is that financial volatility tends to come in "clusters", where tranquil periods of small returns are interspersed with volatile periods of large returns.⁶ Long-term volatility predictions should be unaffected by this "clustering" behaviour, and so we need to take an average squared return over a long historic period. But short-term volatility predictions should reflect current market conditions, whether volatile or tranquil, that means that only the immediate past returns should be used.

3. The standard unbiased estimate of sample variance is $\sum(x_i - \bar{x})^2 / (n - 1)$ where n is the number of observations in the sample. Note that the use of $n - 1$ rather than n in the denominator is a bias correction which is not necessary when actual returns rather than mean deviations are used.
4. S Figlewski, "Forecasting Volatility Using Historical Data", Working Paper, Series No S94-13, Leonard N Stern School of Business, Salomon Center, New York University, 1994; C O Alexander and C Leigh, "On the Covariance Matrices Used in VAR Models" (1997) 4 (Spring) *Journal of Derivatives*.
5. See above n 1.
6. As long ago as 1963 Benoit Mandelbrot (of the "Fractal" fame) observed that financial returns time series exhibit periods of volatility interspersed with tranquillity, where "Large returns follow large returns, of either sign. . . ."

Although there is a rationale for using the current n -day volatility estimates as forecasts of average volatility over the next n -days, the equally weighted averaging induces some very misleading properties in historic volatility time series. *Exhibit 9.1* shows two such “historic” volatility series for the FTSE with $n=30$ and $n=60$ respectively (and both are calculated using daily returns). The figure illustrates a basic problem with the use of equally weighted averages that has motivated a general shift in methodology towards exponentially weighted moving averages for financial market analysis. First, when there is a jump in market price, an equally weighted average of squared returns will jump up the very next day. This is as an accurate reflection of the “clustering” behaviour of volatility in financial markets. But that one, large, squared return will continue to keep volatility estimates high for exactly 30 days in the 30-day moving average, and exactly 60 days in the 60-day moving average, whereas the underlying volatility will have long ago returned to normal levels. Secondly, exactly 30 days (or 60 days) after a major market event the equally weighted volatility estimate will jump down again as abruptly as it jumped up. What has been seen from the event until that day is just a ghost of what happened 30 days ago. But there was nothing special about that day—it was just the day on which a (long overdue) correction in the estimate occurred. Because the average is taken over fewer observations, this correction will be bigger in short-term volatility estimates. For example, the 30-day FTSE volatility jumped up to 68% for the 30 days after Black Monday before it jumped back to its normal levels around 13%. But the 60-day volatility only jumped up to around 50%—although it took twice as long to correct.



These “ghost features” are always going to be a problem when equally weighted averages are applied to financial market data.⁷ Not only do they keep volatility estimates artificially high for an arbitrary length of time, but they can induce an artificial sense of stability into covariance (and


7. Such as in the current practice of producing market betas or statistical hedge ratios (see section on Applications below).

correlation) estimates. An n-period “historic” covariance estimate between two one-period returns series $r_{1,t}$ and $r_{2,t}$ is given by

$$\hat{\sigma}_{12,T} = \frac{\sum_{t=T-n}^{t=T-1} r_{1,t} r_{2,t}}{n}$$

and this may be converted to n-period correlation by dividing by the product of the two n-period standard deviations. Again, it is common practice to use the current n-day correlation estimate as a forecast of correlation over the next n days.

Now, unlike volatility, stable correlations do not always exist.⁸ If a time series of instantaneous correlation estimates—such as GARCH estimates—were generated using a pair of returns series that are not “jointly stationary”, these correlations would be very unstable. However, if an equally weighted n-period moving average is used instead of GARCH, an artificial stability is induced in these correlation estimates by the “ghost features”. This feature is evident in *Exhibit 9.2*, which shows 30-day and 60-day historic correlations between MSCI indices for the UK and the US.



8. Volatility is applicable when financial returns are stationary—which they usually are—but correlations are applicable only when the two returns series are *jointly* stationary. Although joint stationarity can often be assumed, for example between two domestic bond returns, it is commonly found that two arbitrarily chosen returns are not jointly stationary. For example, returns to some South American Brady bonds are unlikely to be jointly stationary with an Indian equity—and in this case “correlation” estimates will be very unstable.

2.2 Exponentially weighted moving averages

The “ghost features” produced in equally weighted averages of financial market returns are clearly a problem, which is caused by the fact that all past returns that come into the average are equally important, however long ago. Exponentially weighted moving averages (EWMA) of squared returns do not exhibit these “ghost features” because past returns become less significant as time goes on. An EWMA variance estimate at time T , based on a time series of squared returns is

$$\hat{\sigma}_T^2 = (1-\lambda) \sum_{i=1}^{\infty} \lambda^{i-1} r_{T-i}^2$$

In an EWMA past observations are now weighted by the “smoothing constant” λ , which is between 0 and 1, so an observation n days ago is multiplied by λ^n , which is very small for large n . Thus extreme events have less of an impact on variances and covariances as they move further into the past. Since shocks are smoothed out by the exponential weighting, this technique is sometimes called “exponential smoothing”—the bigger the smoothing constant λ , the more weight is given to past observations and the smoother the resulting time series.

EWMA is a standard statistical estimation technique, but for volatility and correlation forecasting in financial markets it has some limitations. First, it is only really useful as a one-step-ahead forecast. That is, if daily squared returns are smoothed using EWMA, this provides a useful series of one-day variance forecasts (which are in fact similar to GARCH one-day forecasts).⁹ However there is only one way in which we can extend the forecast horizon and use EWMA methods for n -day forecasts, and that is to smooth n -day squared returns. Alternative methods—such as using the “square root of time” rule,¹⁰ or applying exponential smoothing to an equally weighted

9. C O Alexander, “Volatility and Correlation Forecasting” in C O Alexander (ed), *Handbook of Risk Management and Analysis* (Wiley, 1996).

10. The “square root of time” rule simply calculates h -day standard deviations as \sqrt{h} times the daily standard deviation. It is based on the assumption that daily log returns are normally, independently and identically distributed, so the variance of h -day returns is just h times the variance of daily returns. But since volatility is just an annualised form of the standard deviation, and since the annualising factor is—assuming 250 days per year— $\sqrt{250}$ for daily returns but $\sqrt{250}/h$ for h -day returns, this simply amounts to the Black-Scholes constant volatility assumption, i.e. that current levels of volatility remain the same.

variance series—are doomed to failure.¹¹ Another limitation of this technique is that—unlike GARCH—it gives no optimal method of estimating the parameter. Finally, there can be two problems when these methods are applied to generate large covariance matrices:

1. The same value of the smoothing constant needs to be applied to all series, otherwise the matrix will not be “positive semi-definite”, that means that it would give negative risk on certain portfolios;
2. EWMA covariance matrices effectively use only the last m data points, where m is such that $\lambda^{m-1} \approx 0$.¹²

Hence covariance matrices with k factors will have many zero eigenvalues when $k \gg m$, and special algorithms to cope with these nearly singular matrices are needed, otherwise many of these zero eigenvalues will be estimated as being negative and the “indefinite” matrices which result will give negative risk measures for certain portfolios anyway.¹³

3. GARCH METHODS

GARCH stands for Generalized Autoregressive Conditional Heteroscedasticity: “Generalized” because it is a general class of ARCH model, “Autoregressive” because the variances generated by ARCH models involve regression on their own past, and “Conditionally Heteroscedastic” literally means changing variance, or “volatility clustering” as it has become known. The first ARCH model was introduced by Rob Engle (1982). Tim Bollerslev (1986) developed the GARCH formulation of the model, which is more commonly used in financial markets.¹⁴

To understand GARCH we first need to recall the ideas of standard linear regression. Linear regression models are commonplace in financial market analysis; they are used to calculate market betas, statistical hedge ratios, and sensitivities in more general factor models (see the section below on “Applications”); they can be used for prediction, and for confidence limits and other types of statistical inference on model parameters. The innovation in GARCH is to augment the standard linear regression model with another

11. C O Alexander, “Evaluating the Use of RiskMetrics™ as a Risk Measurement Tool for your Operation: What are its Advantages and Limitations” (1996) 4 *Derivatives Use, Trading and Regulation*.

12. Let the original data be $X = \begin{bmatrix} Y \\ Z \end{bmatrix}$, $A = \text{diag}(1, \lambda, \lambda^2, \dots, \lambda^{t-1})$ $\tilde{X} = AX$
where X is $t \times k$, Y is $m \times k$ where m is the rank of $A'A$ and Z is $(t-m) \times k$.
Let $\Lambda = \text{diag}(1, \lambda, \lambda^2, \dots, \lambda^{m-1})$ so that $A'A$ can be partitioned $\begin{bmatrix} \Lambda' \Lambda & 0 \\ 0 & 0 \end{bmatrix}$

Then we can write $\tilde{X}'\tilde{X} = X' A' A X = (Y' Z') A' A \begin{bmatrix} Y \\ Z \end{bmatrix} = Y' \Lambda' \Lambda Y$

This proves that only the first m rows of the original data X are used in the calculation of the EWMA covariance matrix.

13. It also means that the Cholesky decomposition does not exist. The Cholesky decomposition of the covariance matrix is needed for evaluating Value-at-risk of options portfolios.
14. A survey of these models is given in Alexander, op cit n 9.

equation, the “conditional variance” equation—the equation in the original model now being termed the “conditional mean” equation. The parameters in both equations are estimated simultaneously, using maximum likelihood estimation, and the outputs of the GARCH model include two time series: the estimated conditional mean series (the “fitted” series in standard regression) and the estimated conditional variance series (the GARCH variance estimate). The standard input for a GARCH model for financial market volatility is a series of daily returns, r_t .¹⁵ This input is called the dependent variable, and its level is modelled by the conditional mean equation. The conditional mean equation can be anything reasonable,¹⁶ such as

$$r_t = \phi_0 + \phi_1 r_{t-1} + \varepsilon_t$$

If we assume that $V(\varepsilon_t) = \sigma^2$ the error process ε_t is called “homoscedastic” (that is, it has constant variance) and then this would just be an ordinary regression model. But in GARCH we allow the residual to have time-varying variance, in particular we have the conditionally heteroscedastic assumption that $V_t(\varepsilon_t) = \sigma_t^2$ and we model this time-varying variance with another equation—the conditional variance equation. The GARCH(1,1) model has the conditional variance equation

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2$$

where the restrictions $\omega > 0$, $\alpha, \beta \geq 0$ are placed on the parameters to ensure that variance is positive. Also $\alpha + \beta < 1$ ensures that GARCH volatility forecasts “mean-revert”, so that forecasts will get closer to the long-term average volatility as the maturity of the forecast increases.

An important difference between GARCH and the moving average methods described above is that the GARCH volatility estimates are different from the forecasts. GARCH forecasts of volatility of any maturity can be computed in a simple iterative manner: Put

$$\sigma_{t+1}^2 = \omega + \alpha \varepsilon_t^2 + \beta \sigma_t^2$$

and for $s > 1$

$$\sigma_{t+s}^2 = \omega + (\alpha + \beta) \sigma_{t+s-1}^2$$


Then, assuming returns have no autocorrelation the GARCH forecast of variance of h-day returns is

$$\sigma_{t,h}^2 = \sum_{i=1}^h \sigma_{t+i}^2$$


and the GARCH h-day volatility forecast is $100\sigma_{t,h}\sqrt{250}\%$, assuming 250 trading days per year. GARCH term structures behave in the intuitive “mean-reverting” fashion illustrated in *Exhibit 9.3*. In fact their behaviour can be very similar to implied volatility term structures.¹⁷


15. The length of historic data needed for GARCH needs to be long enough that the maximum likelihood estimation converges, but not so long that extreme market events that occurred a long time ago have an effect on current long-term GARCH forecasts.
16. However, the more parameters in the model the more flat the likelihood function and so the more difficult it becomes to get the estimation procedure to converge. Often a very parsimonious parameterization is used in the conditional mean equation, such as the AR(1) model in this example.
17. R F Engle and J Mezrich, “Grappling with GARCH” (1995) 8(9) *RISK*.

Exhibit 9.3
GARCH Volatility Term Structures of US\$ Rates (28 November 1994)



GARCH Volatility Term Structures of US\$ Foreign-exchange Rates (2 March 1995)






Moving average methods, on the other hand, assume constant volatility term structures.¹⁸ Thus current levels of volatility are supposed to persist forever, which is clearly a very unrealistic assumption. There is no doubt of the success of GARCH in volatility forecasting, but when we try to extend the univariate model so that GARCH covariances (and correlations) are obtained, we run into programming problems. The multivariate GARCH model takes as inputs a vector time series of returns r_t , for example the returns to all risk factors in a portfolio. The variance of a vector time series is a matrix time series H_t ; the covariance matrix. Since multivariate GARCH models need to output not a single scalar variance but a whole matrix at every point in time, there are a very large number of parameters in a multivariate GARCH model and this makes direct computation of GARCH covariance matrices very difficult. There is a multivariate GARCH method called “orthogonal GARCH” that gets around these programming difficulties by using a combination of principal components analysis and univariate GARCH.¹⁹ This orthogonal GARCH method has additional merits: the lack of dimensional restrictions, and the fact that these matrices are positive definite. Also, by adding or subtracting principal components, the model can be tailored to cut down the “noise” in the data if so desired, and this process makes correlation forecasts more stable. An example explaining the method for just two categories can easily be extrapolated to any number of risk factor categories: Suppose $P = (P_1, \dots, P_n)$ are the principal components of the first system (n risk factors) and let $Q = (Q_1, \dots, Q_m)$ be the principal components of the second system (m risk factors). Denote by A ($n \times n$) and B ($m \times m$) the factor weights matrices of the first and second systems. Within factor covariances are given by $AV(P)A'$ and $BV(Q)B'$ respectively and cross-factor covariances are ACB' where $V(P)$ and $V(Q)$ denote the diagonal matrices of GARCH variances of the principal components and C denotes the

18. See above n 10.

19. C O Alexander and A M Chibumba, “Orthogonal Factor GARCH” (1997).

mxn matrix of GARCH covariances of principal components. As an illustration of the method I have applied it to a system of equity indices (where the S&P has to be taken into a category of its own) and USD exchange rates. The system is small enough to compare the "splicing" method with full multivariate GARCH estimation (just!) and *Exhibit 9.4* shows just one of the resulting correlations obtained by each method—between the GBP/USD exchange rate and the Nikkei during the period 1 Jan 1993 to 17 Dec 1996.



However, the model can break down in stress scenarios, so although the stable variances and covariances it provides are perfect for measuring everyday risks they should be used with caution for volatility and correlation forecasting during extreme market conditions.

Although GARCH models are a little more complex than the moving average methods of volatility and correlation estimation, their sound mathematical structure makes them easy to adapt to different environments. For example, correlation estimates and forecasts will be upset if data are asynchronous, which they often are. Rather than invest huge amounts of time in the collection of synchronous data sets, a simple modification of the conditional mean equations in a multivariate GARCH model generates forecasts which are based on *two* period correlations which allow for feedback between markets which close at different times. This is the "asynchronous" GARCH model being developed by Rob Engle and Joe Mezrich at Saloman Bros, NY.

A final point which distinguishes GARCH from the moving average methods outlined above is that the estimated GARCH model, that is commonly used to generate well-behaved variance and covariance forecasts of any maturity, also provides a stochastic volatility/correlation model that is very effective for pricing and hedging options. The estimated parameters are

used in the conditional variance equations, and then Monte Carlo simulations are applied to a two factor model as outlined below (see “Option Pricing and Hedging”).

4. SOME APPLICATIONS OF VOLATILITY AND CORRELATION TO MARKET RISK

4.1 Measuring market betas in factor models

Simple factor models describe returns to an asset or portfolio in terms of returns to risk factors and idiosyncratic return. For example, in the simple Capital Asset Pricing Model (CAPM) the return to a portfolio, denoted r_t , is approximated by the linear regression model

$$r_t = \alpha + \beta R_t + \varepsilon_t$$

where R_t is the return to the risk factor, the error process ε_t denotes the idiosyncratic return and β denotes the sensitivity (or “beta”) of the portfolio to the risk factor. This beta is estimated by the ratio of the covariance (between the portfolio and the risk factor) to the variance of the risk factor.²⁰

More generally, in multivariate factor models, portfolio returns are attributed to several risk factors: the portfolio return is given by

$$\alpha + \beta_1 R_1 + \beta_2 R_2 + \dots + \beta_3 R_3 + \varepsilon$$

where R_1, \dots, R_n are the returns to different risk factors, β_1, \dots, β_n are the net portfolio sensitivities (that is, the weighted sum of the individual asset sensitivities) and ε is the idiosyncratic return of the portfolio (that is, that part of the return not attributed to the different risk factors). The net betas with respect to each risk factor can be calculated from the covariances (between risk factors and between the portfolio and risk factors) and the variances of the risk factors. EWMA or GARCH estimates of these quantities then give time-varying market betas which are a great improvement on the constant betas one usually obtains from standard data suppliers.

4.2 Measuring portfolio risk

It is standard procedure to measure risk in a linear portfolio in terms of variance and covariances.²¹ In simple portfolios which can be described by a weighted sum of the constituent assets, all we need to find is the covariance matrix of these asset returns. The return to the portfolio is $w'r$ where w is the vector of portfolio weights and r is the vector of asset returns so the portfolio variance is $w'Vw$ —that is, the variance of the return—where V is the covariance matrix of r . Provided V is a positive definite matrix, the portfolio risk will always be positive, whatever the weights in the portfolio.²²

20. So the beta is the correlation times the relative volatility.

21. Variance is not the only measure of portfolio risk—“Value-at-risk”, “regret” or “downside risk” are also used (see Beckers, 1996).

22. One of the problems with the RiskMetrics™ data is that they are not positive definite, they are “indefinite”, so certain portfolios will have negative VAR (see Alexander, 1996b).

In larger linear portfolios, that are best described not at the asset level but by the factor models described above, the portfolio variance is measured using the covariance matrix of risk factor returns and the vector of sensitivities to different risk factors. If we ignore the idiosyncratic risk, the portfolio variance from the factor model is simply $\beta'V\beta$ where β is the vector of (net) portfolio sensitivities and V is the covariance matrix of risk factor returns. Both V and β are obtained from variances and covariances, so the same remarks apply: it is better to use time-varying estimates such as those obtained from EWMA or GARCH models than the "constant" variances and covariances that one can get from ordinary least squares regression or equally weighted moving averages.

4.3 Constructing hedged portfolios to minimise portfolio risk

When hedging a portfolio with forwards or futures we regard their deltas as weights in the total portfolio. The problem is of choosing a vector of weights ω in a portfolio to minimise its variance. In mathematical notation:

$$\text{Min}_{\omega} \omega' V \omega \text{ such that } \sum \omega_i = 1$$

where V is the covariance matrix of asset returns. Assuming that weights can be negative or zero this is a straightforward linear programming problem, having the solution

$$\omega_i^* = \psi_i / \sum \psi_i \text{ where } \psi_i = \sum \text{ith column of } V^{-1}.$$

Suppose we take V to be a time-varying covariance matrix such as the GARCH matrix. This will give time-varying weights $\omega_{i,t}^*$ that can be used to re-balance the portfolio as volatilities and correlations between the assets change.

4.4 Option pricing and hedging

There are two types of application of volatility and correlation to pricing and hedging options. Either the relevant forecast of average volatility and/or correlation over the life of the option is "plugged" into the closed form solution (for example, into the Black-Scholes formula for vanilla options), or a two factor model with price diffusion and stochastic volatility model is used to simulate terminal price distributions from which the option price is calculated as the discounted expected option pay-off assuming risk-neutral evaluation.²³ Hedge ratios are also calculated using finite difference approximations on these simulations.²⁴

To fix ideas, first recall how simulation can be used to price and hedge a call option on an underlying price $S(t)$ which follows the Geometric Brownian Motion diffusion

$$dS(t)/S(t) = \mu dt + \sigma dZ$$

where Z is a Wiener process. Since volatility σ is constant this is a one factor model, and it is only necessary to use Monte Carlo on the independent increments dZ of the Wiener process to generate price paths $S(t)$ over the life

23. J C Hull, *Options, Futures and Other Derivative Securities* (Prentice Hall, 1993).

24. M Broadie and P Glasserman, "Estimating Security Price Derivatives Using Simulation" (1996) *2 Management Science* 42.

of the option. This is done on the discrete form of Geometric Brownian Motion, namely

$$S_t = S_{t-1} \exp(\mu - 0.5\sigma^2 + \sigma z_t)$$

where $z_t \sim \text{NID}(0,1)$.²⁵ So, starting from the current price S_0 , Monte Carlo simulation of an independent series on z_t for $t=1,2,\dots,T$ will generate a terminal price S_T . Thousands of these terminal prices should be generated starting from the one current price S_0 , and the discounted expectation of the option pay-off function gives the price of the call. For example for a plain vanilla option

$$C(S_0) = \exp(-rT) E(\max\{S-K,0\})$$

where r is the risk-free rate, T the option maturity and K is the strike. So from the simulated distribution $S_{T,i}$ ($i = 1, \dots, N$) of terminal prices we get the estimated call price

$$\hat{C}(S_0) = \exp(-rT) (\sum_i \max\{S_{T,i} - K, 0\} / N)$$

Simulation deltas and gammas are calculated using finite difference approximations, such as the central differences

$$\delta = [C(S_0 + \eta) - C(S_0 - \eta)] / 2\eta$$

$$\gamma = [C(S_0 + \eta) - 2C(S_0) + C(S_0 - \eta)] / \eta^2$$

In this example, with constant volatility GBM for the simulations, the delta and gamma should be the same as those obtained using the Black-Scholes "plug-in" formulae. But in practice, simulation errors can be very large unless the time is taken to run large numbers of simulations for each option price.²⁶

Now consider how to extend the standard GBM model to a two factor model, where the second factor is GARCH(1,1) stochastic volatility. The two diffusion processes are

$$S_t = S_{t-1} \exp(\mu - 0.5\sigma_t^2 + \varepsilon_t)$$

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2$$

where $\varepsilon_t = \sigma_t z_t$. Starting with current price S_0 and unconditional standard deviation σ_0 , an independent set of Monte Carlo simulations on z_t ($t = 0, 1, \dots, T$) will now generate σ_t at the same time as S_t for $t = 1, \dots, T$. Note that the simulated price paths are already based on expected volatility levels, so the GARCH delta and gamma hedge ratios do not require additional vega hedging.²⁷

GARCH is not the only way of putting stochastic volatility into option prices and hedge ratios: there are alternatives such as the "autoregressive in variance" model.²⁸ However, moving average methods do not yield stochastic volatility models and so these methods can only be used to "plug" into closed form solutions; they cannot be used to generate option prices and

25. To derive this from the continuous form use Ito's lemma on $\log S$ and then make time discrete.

26. Simulation errors are reduced by using correlated random numbers—the variance of delta estimates is reduced when $C(S_0 + \eta)$ and $C(S_0 - \eta)$ are positively correlated.

27. R F Engle and J Rosenberg, "GARCH Gamma" (1995) *Journal of Derivatives*.

28. S J Taylor, "Modelling Stochastic Volatility: A Review and Comparative Study" (1994) 4(2) *Mathematical Finance* 183.

hedge ratios which already account for stochastic volatility, so moving average deltas and gammas do need additional vega hedges.

4.5 Fitting the smile with GARCH

GARCH models can be estimated using cross-section data on the market implied smile surface, and then the dynamics of the GARCH model can be used to predict the smile. Initial values for the GARCH model parameters are fixed, and then GARCH option prices obtained, as explained above, for options of different strikes and maturity. These prices are then put into the Black-Scholes formula, and the GARCH “implied” volatility is backed out of the formula (just as one would do with ordinary market implied volatilities, only this time the GARCH price is used instead of the price observed in the market). Comparison of the GARCH smile surface with the observed market smile surface leads to a refinement of the GARCH model parameters (that is, iteration on the root mean square error between the two smiles) and so the GARCH smile is fitted. It turns out that the GARCH parameters estimated this way are very similar to those obtained from time series data, so using the GARCH smile to predict future smiles leads to sensible results.²⁹

4.6 Volatility and correlation trading

We have just seen that options prices are based on assumptions about volatility and correlation during the lifetime of the option. Options sellers will account for volatility and correlation in their prices, and each market price of an option has forecasts of volatility and/or correlation implicit in its price. This “implied” volatility and correlation needs to be distinguished from the statistical volatility and correlation forecasts which are calculated using moving averages or GARCH. In fact, implied and statistical methods use different data and different models³⁰ to forecast the same thing: the volatility of the underlying assets over the life of the option.

Implied volatility is like an option price—it contains all the forward expectations of investors about the likely evolution of the underlying. In fact implied volatility is an inverse option price: if we know the market price of the option and all the other parameters in a particular pricing functional used for the option, we can invert this price functional to solve for volatility: this is the implied volatility. If the price functional were an accurate representation of reality, and if the statistical forecasts were known to be accurate, then any observed differences between implied and statistical volatility would reflect inappropriate expectations from investors. So volatility traders use the relationship between statistical and implied measures to form expectations of implied volatility and so also the likely movements in options prices.

29. J Duan, “Cracking the Smile” (1996) *RISK*.

30. Implied methods use current data on market prices of options—statistical methods use historic data on the underlying asset price. Implied methods use an assumed diffusion process for the underlying asset price—statistical methods assume only (conditional) normality of underlying asset returns.

For example, 95% confidence limits of GARCH n -period volatility forecasts can be generated using the covariance matrix of estimated coefficients in the GARCH model. These are then tracked over a period of time and implied volatility is examined in relation to these limits. A volatility purchase (such as buying an at-the-money straddle) would be appropriate when implied volatility falls below the lower 95% limit. On the other hand, when implied volatility seems too high, above the 95% GARCH upper limit, the volatility trader should rather sell volatility.

The P&L of certain volatility trades may be reduced by trading pairs of options on assets which have been highly correlated historically. For example, a trade on the relative volatility V_1/V_2 of two highly correlated equities may consist of two straddles where one is bought and one is sold in vega neutral proportions.³¹ Tracking relative implied volatilities and their statistical historical levels provides a means of determining whether current relative volatility is out of line: If their relative volatility is thought to be too high, the straddle on asset 1 is sold and the straddle on asset 2 is bought, with the opposite trade being relevant for traders who believe relative volatility is undervalued.

4.7 Value-at-risk modelling

Value at risk (VAR) has become central to financial decision making, not just for risk managers and regulators, but for anyone concerned with the actual numbers produced: traders, fund managers, corporate treasuries, accountants etc. Not only is capital set aside for regulatory compliance on the basis of these measures—such things as trading limits or capital allocation decisions may be set. Providing accurate VAR measures therefore becomes a concern for many.

A VAR measure depends on two parameters, the holding period h and the significance level, α . A good VAR model should provide a convergent and consistent sequence of VAR measures for every holding period from one day (or even less than a day) to several years.³² However the significance level is just a matter of personal choice: Do you want normal circumstances VAR measures to reflect potential losses one day in twenty ($\alpha=0.05$) or one day in a hundred ($\alpha=0.01$)?

The generally accepted methods for linear portfolios (“Covariance” methods) and options portfolios (“Structured Monte Carlo”) both require an accurate, positive definite covariance matrix. In the “covariance” method, denote by ΔP_t the forecast P&L over the next h days, and by μ_t and σ_t^2 its mean and variance. Then assuming P&Ls are conditionally normally distributed we have³³

$$\text{VAR} = Z_\alpha \sigma_t - \mu_t$$

31. M D Fitzgerald, “Trading Volatility” in C O Alexander (ed), *Handbook of Risk Management and Analysis* (Wiley, 1996).

32. Currently only GARCH models can really do this effectively—moving average methods have problems because they are too simple, and other more advanced models such as neural networks still have some way to go.

33. Apply the standard normal transformation to ΔP_t .


where Z_α denotes the critical value from the standard normal distribution corresponding to your choice of significance level. It is often assumed that μ_t is zero, and Z_α is just looked up in tables, so the sole focus of this method is on σ_t , the standard deviation of forecast P&L over the holding period. But

$$\sigma_t^2 = \mathbf{P}'\mathbf{V}\mathbf{P}$$

where \mathbf{V} is the covariance matrix of asset returns (or risk factor returns) over the holding period and \mathbf{P} is a vector of the current mark-to-market values of the assets (or of the current sensitivities times the current mark-to-market values of the risk factors). Since \mathbf{V} is the only stochastic part of the model it is crucial to obtain good estimates of the parameters in \mathbf{V} , that is, the variance and covariance forecasts of asset (or factor) returns over the holding period. For example, the model may give zero or negative VAR measures if \mathbf{V} is not positive definite.

VAR measures for options portfolios are often calculated using Monte Carlo methods to simulate a distribution of P&Ls over the forthcoming holding period. To calculate this distribution, correlated vectors of underlying returns are simulated over the holding period. These are generated using the Cholesky decomposition of the covariance matrix of returns to the underlyings³⁴ and it is (again) important to obtain positive definite covariance matrices, because otherwise the Cholesky decomposition does not exist.³⁵ Each set of correlated simulations gives one value of the portfolio at the end of the holding period. Lots of such values are simulated and used to calculate a simulated P&L distribution from which the VAR measure can be read off directly as $\text{Prob}(\Delta P_t < -\text{VAR}) = \alpha$ (see *Exhibit 9.5*).

34. If \mathbf{r} is a vector of standard independent returns, and $\mathbf{C}'\mathbf{C} = \mathbf{V}$ where \mathbf{C} is the (triangular) Cholesky matrix, then $\mathbf{C}'\mathbf{r}$ is a vector of returns with variance covariance matrix \mathbf{V} .
35. So what can be done with the non-positive semi-definite matrices, produced by RiskMetrics™? Ad hoc methods—such as shooting the negative or zero eigenvalues to something small and positive—must be resorted to. Unfortunately this changes the original covariance matrix in an arbitrary way, without control over which volatilities and correlations are being effected.



The current emphasis on integrated risk systems, to measure VAR across all the risk positions of a large bank, requires very large covariance matrices indeed. Regulators require the use of at least one year of historic data, so exponentially weighted moving average methods are ruled out. Currently there are only two realistic alternatives: equally weighted moving averages or GARCH. Equally weighted moving average covariance matrices are easy to construct, and should be positive definite (assuming no linear interpolation of data along a yield curve, et cetera) but will be contaminated by any stress events which may have occurred during the last year. These should be filtered out of the data before the moving average is applied, and saved for later use when investigating the effect of real stress scenarios on portfolio P&Ls. Multivariate GARCH models of enterprise-wide dimensions would be computationally impossible. The only way that one can construct large dimension covariance matrices using GARCH is to employ the orthogonal GARCH model and “splice” together the full covariance matrix.³⁶

5. SUMMARY

This chapter has described current state-of-the-art techniques for estimating and forecasting covariance matrices, and surveyed some of their applications to financial markets analysis. In the earlier part of the chapter it was shown that one has to be very careful about using moving average techniques, however simple they may appear. Their drawbacks include: “indefiniteness” of covariance matrices estimates, so that certain portfolios will have negative risk measures; “ghost features” in equally weighted average methods; the absence of any model for stochastic volatility; and problems with extending estimates to forecasts—in fact only their one-step ahead forecasts are meaningful. Thus moving average techniques have very limited applicability. On the other hand, although the computational side of GARCH models is

36. C O Alexander, “‘Splicing’ Methods for Value-at-Risk” (1997) *Derivatives Week*.

relatively complex, utilisation of GARCH techniques is becoming more and more common. GARCH is just another name for “volatility clustering”, so it is particularly relevant to modelling financial returns. Also, since GARCH models are based on sound statistical techniques such as maximum likelihood, it is easy to apply these models in a number of ways without having to resort to bits of string and elastoplast. In the survey of applications it has been shown how elegantly GARCH models can be adapted and applied to enhance a number of areas in the management of risk, investment, pricing, hedging and trading.

Chapter 10

Measuring Option Price Sensitivity—The “Greek Alphabet” of Risk

by Satyajit Das

1. INTRODUCTION

The fair value of an option is a function of five parameters: the price of the asset (S); the option strike price (K); the time to option expiry (T); the risk free rate (R_f); and the volatility of the asset returns (σ). Changes in each of these variables directly impact upon the price of the relevant option (see discussion in Chapter 5). In practice, precise and quantitative estimates can be obtained of the *directional* as well as *quantum* effect of changes in the option pricing input parameters. These sensitivities of the option premium are represented by a number of Greek letters (following the notation conventions of mathematics), which are used to quantify and estimate the risk of options.

In this chapter, the techniques of quantifying the sensitivity of option prices using the Greek alphabet of risk is examined. An overview of option price sensitivities and risk, including potential uses, is first considered. Each of the individual risk elements—delta (Δ), gamma (γ), vega (κ),¹ theta (τ) and rho (ρ)—are separately considered, including their function and significance as well as their behaviour. Other risk factors—such as lambda (λ), charm, speed, colour and fugit—are also identified as well as a generalisation of the concepts to cover option risks where the underlying option involves different assets. The chapter also examines the use of extending the overall conceptual framework to cover risk generally (as distinct from risk in relation to option transactions).

2. OPTION PRICE SENSITIVITIES AND RISKS—AN OVERVIEW

As noted above, the sensitivities and risks of option transactions with references to changes in option pricing parameters are obtained through the mathematical option pricing models which not only provide a closed form means of valuing option contracts, but, in addition, provide a wealth of additional data with respect to the various formula variables. For example, the delta, that is, the derivative of the option premium with respect to asset price, provides investors, portfolio managers or market makers with the exact

1. Vega is not a Greek letter and, in practice, the Greek letter kappa (κ) is used to denote the sensitivity of the option premium to changes in volatility. However, the term vega is used in the text to denote this sensitivity. Other Greek letters used to denote this sensitivity include epsilon (ϵ) and lambda (λ).

hedge ratio required to hedge their portfolio position in options or in the underlying assets.

These derivatives allow market participants to identify the short-term sensitivity of option premiums to changes in the underlying security price, volatility, time to expiration, et cetera. In mathematical terms, these sensitivities are partial derivatives of the premium with respect to these parameters. *Exhibit 10.1* summarises the key partial derivatives in relation to a European option on a non income paying asset. *Exhibit 10.2* sets out the actual mathematical partial derivatives for a European option on a non-income paying asset.

Exhibit 10.1
Option Risk Measures—The Greek Alphabet

Option Derivative	Concept
Delta (Δ)	Delta is the derivative of the option pricing formula with reference to the asset price (S). It measures the estimated change in the option premium for a change in S.
Gamma (γ)	Gamma is the second derivative of the option pricing formula with reference to the asset price (S). It measures the estimated change in the delta of the option for a change in S.
Vega (κ)	Vega is the derivative of the option pricing formula with reference to the volatility of the asset returns (σ). It measures the estimated change in the option premium for a change in σ .
Theta (τ)	Theta is the derivative of the option pricing formula with reference to the time to option expiry (T). It measures the estimated change in the option premium for a change in T.
Rho (ρ)	Rho is the derivative of the option pricing formula with reference to the risk free rate (Rf). It measures the estimated change in the option premium for a change in Rf.

Exhibit 10.2
Option Derivatives—European Options on a
Non-income-producing Asset

Option Derivative	Partial Derivative
Delta—Call	$\Delta = N(d1)$
Delta—Put	$\Delta = N(-d1)$
Gamma—Call & Put	$\gamma = N(d1)/S \cdot \sigma \cdot \sqrt{T}$
Vega—Call & Put	$\kappa = S \cdot \sqrt{T} \cdot N(d1)$
Theta—Call	$\tau = (S \cdot \sigma \cdot N(d1)/2 \sqrt{T}) - R_f \cdot K \cdot e^{-R_f \cdot T} N(d2)$
Theta—Put	$\tau = (S \cdot \sigma \cdot N(d1)/2 \sqrt{T}) - R_f \cdot K \cdot e^{-R_f \cdot T} N(-d2)$
Rho—Call	$\rho = K \cdot T \cdot e^{-R_f \cdot T} N(d2)$
Rho—Put	$\rho = K \cdot T \cdot e^{-R_f \cdot T} N(-d2)$

The major applications of these measures of risk include:

1. The measurement of the risk of options in terms of its behaviour in response to changes in an individual market parameter, such as the asset price.
2. Facilitating the replication of asset by synthetically creating the option's economic payoffs by trading in the underlying asset.
3. Using the sensitivity and behaviour of the option as a means for precisely defining the hedging or other objectives of the option trader and enabling the creation of more efficient strategies.

In analysing the risk measures, individual components are illustrated using a series of examples based on a consistent set of inputs. The example utilised is set out in *Exhibit 10.3*. The example used is that of a European option on a non-income-paying asset.

**Exhibit 10.3
Option Risk Measures—Example**

**OPTION PRICING MODEL—BLACK-SCHOLES
INPUTS/OPTION PRICING/SENSITIVITIES**

PRICING INPUTS	
UNDERLYING ASSET PRICE	100.00
STRIKE PRICE	100.00
TRADE DATE	1-Jan-95
EXPIRY DATE	1-Jul-95
VOLATILITY	20.00%
RISK FREE RATE	10.00%
INCOME ON ASSET	0%
MODEL OUTPUTS	
OPTION PREMIUM	8.23
OPTION PREMIUM (% OF ASSET PRICE)	8.23%
	CALL
	PUT
	3.39
	3.39%
PREMIUM SENSITIVITIES/"GREEKS"	
DELTA	0.6637
GAMMA	0.0259
VEGA	0.26
THETA (pa)	10.99
THETA (per day)	0.03012
RHO	0.2883
	CALL
	PUT
DELTA	-0.3363
GAMMA	0.0259
VEGA	0.26
THETA (pa)	1.48
THETA (per day)	0.00405
RHO	-0.1836

Description of outputs
 A 1.00 increase in S will change premium by = 0.6637
 A 1.00 increase in S will change delta by = 0.0259
 A 1% increase in vol will change the premium by 0.2569
 Divide by 365 to get daily time decay 10.9950
 A 1% incr. in int rate will change by premium by 0.2883

Description of outputs
 A 1.00 increase in S will change premium by = -0.3363
 A 1.00 increase in S will change delta by = 0.0259
 A 1% increase in vol will change the premium by 0.2569
 Divide by 365 to get daily time decay 1.4787
 A 1% incr. in int rate will change by premium by -0.1836

3. DELTA

3.1 Concept

Delta is the first derivative of the option pricing formula with respect to the asset price. The delta measures the expected change in the value of the option for a given change in the asset price.

The expected price change is given by:

Change in the asset price times Delta = Expected Change in Option premium

In *Exhibit 10.3* (above) the deltas are as follows:

Call option—0.6637

Put option—(0.3363)

This implies that for a small change of say .10 in the asset price the option value will change by:

Call option—0.6637 times .10 = 0.07

Put option—(0.3363) times .10 = 0.03

Exhibit 10.4 shows the result of the change in terms of its impact on the premium, which is consistent with the predicted changes.

**Exhibit 10.4
Option Delta**

**OPTION PRICING MODEL—BLACK-SCHOLES
INPUTS/OPTION PRICING/SENSITIVITIES**


PRICING INPUTS		CALL	PUT
UNDERLYING ASSET PRICE	100.10		
STRIKE PRICE	100.00		
TRADE DATE	1-Jan-95		
EXPIRY DATE	1-Jul-95		
VOLATILITY	20.00%		
RISK FREE RATE	10.00%		
INCOME ON ASSET	0%		
MODEL OUTPUTS		CALL	PUT
OPTION PREMIUM	8.30		3.36
OPTION PREMIUM (% OF ASSET PRICE)	8.29%		3.36%
PREMIUM SENSITIVITIES/"GREEKS"			
		CALL	PUT
DELTA		0.6663	-0.3337
GAMMA		0.0258	0.0258
VEGA		0.26	0.26
THETA (pa)		11.01	1.49
THETA (per day)		0.03017	0.00409
RHO		0.2896	-0.1823
Description of outputs			
		A 1.00 increase in S will change premium by =	0.6663
		A 1.00 increase in S will change delta by =	0.0258
		A 1% increase in vol will change the premium by	0.2564
		Divide by 365 to get daily time decay	11.0103
		A 1% incr. in int rate will change by premium by	0.2896
Description of outputs			
		A 1.00 increase in S will change premium by =	-0.3337
		A 1.00 increase in S will change delta by =	0.0258
		A 1% increase in vol will change the premium by	0.2564
		Divide by 365 to get daily time decay	1.4941
		A 1% incr. in int rate will change by premium by	-0.1823

3.2 Behaviour

The delta of an option is characterised by the following pattern of behaviour:

1. The delta of a call (put) option is positive (negative), reflecting the direction of change of the option value for a given increase in the asset price.
2. The option delta is between 0 and 1 for a call option and 0 and -1 for a put option. The delta of the asset is always 1; a long position has a delta of $+1$ while a short position has a delta of -1 .
3. Deep out-of-the-money options have deltas close to zero because they are not very responsive to changes in the underlying asset price. Deep in-the-money options have deltas close to $+1$ or -1 because they move in step with the underlying price. At-the-money options tend to have deltas close to 0.5.
4. The higher the delta the closer the option price changes are to the changes in the asset price and consequently are to the gains and losses that would derive from a position in the underlying asset.

Exhibit 10.5 sets out a delta surface which illustrates the behaviour of the delta for the option depicted in *Exhibit 10.3*.



In utilising delta, it is important to recognise that it holds for only small movements in the asset price. This can be illustrated by comparing the expected change in the option premium and the *actual* change in the option premium. This comparison is set out in *Exhibit 10.6*.

Exhibit 10.6
Delta Performance

Change in Asset Price	Estimated Change in Call Option Premium Using Commencing Delta	Actual Change in Call Option Premium	Difference
+20	+13.27	+16.92	3.65
+10	+6.64	+7.74	1.10
+5	+3.32	+3.62	0.30
+1	+0.66	+0.68	0.02
+0.10	+0.07	+0.07	0
-0.10	-0.07	-0.06	0.01
-1	-0.66	-0.65	0.01
-5	-3.32	-2.97	0.35
-10	-6.64	-5.21	1.43
-20	-13.27	-7.59	5.68

Change in Asset Price	Estimated Change in Put Option Premium Using Commencing Delta	Actual Change in Put Option Premium	Difference
+20	-6.73	-3.08	3.65
+10	-3.36	-2.26	1.10
+5	-1.68	-1.38	0.30
+1	-0.34	-0.32	0.02
+0.10	-0.03	-0.03	0
-0.10	+0.03	+0.03	0.01
-1	+0.34	+0.35	0.01
-5	+1.68	+2.03	0.35
-10	+3.36	+4.79	1.43
-20	+6.73	+12.42	5.69

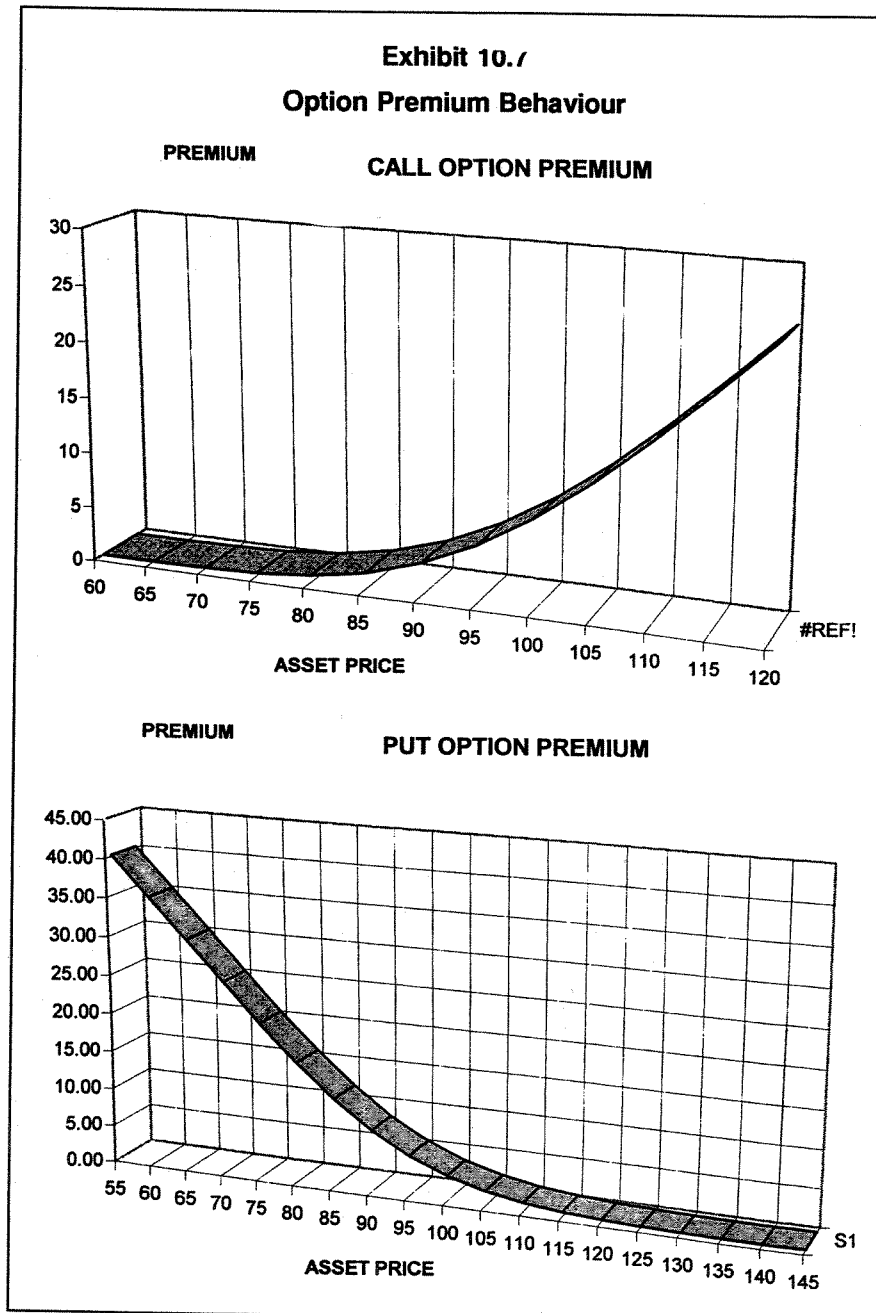
Notes:

The calculations are based on the data in *Exhibit 10.3*.

As is evident, for other than relatively small changes in the asset price, the delta predicted changes in option premium over or under estimate the actual change in the option premium.

This reflects the fact that the option delta is, in effect, the slope of the curve which plots the change in the option premium for given changes in asset price. *Exhibit 10.7* shows the curve relating changes in option premium to changes in asset price. The graph illustrates the curvilinear relationship. The slope of the curve changes with respect to the underlying asset price. As the asset price moves by large increments and larger portions of the curve are covered by the jump, the delta estimate becomes increasingly inaccurate as a predictor of option price movements. This is effectively because delta itself

changes (see *Exhibit 10.4*). This change is measured by gamma, which is described below.



A feature of option deltas is the fact that they are additive; that is, the delta of a portfolio of options is the sum of the face value weighted deltas of the options contained within the portfolio. This relationship can be stated as:

$$\Delta_p = \sum_{n=1}^n w_n \cdot \Delta_n$$

Where

Δ_p = the portfolio delta

Δ_n = the delta of the nth option in the portfolio

w_n = the weight for the nth option calculated as the face value of the option divided by the total face value of options in the portfolio

The additive nature of delta greatly facilitates the management of a portfolio of options.

3.3 Application

Delta is the most important measure of option price sensitivity. It conveys a wide range of information, including:

- the asset content of the option; and
- the probability of the option being exercised.

The delta of an option is often referred to as its *equivalent asset position*. This reflects the fact that a holding of delta amount of the asset (say, a long (short) position of .6637 (−0.3363) in the asset) will provide an economic result in terms of gains or losses for small movements in the asset price which are identical to purchases of the option itself.

The delta of an option also gives the probability of the option being exercised. For example, in the example above, the deltas indicate that the call (put) has an approximately 66.37% (33.63%) probability of exercise. The delta in effect gives the probability that the asset price will be above or below the strike price of the option at maturity. This provides important information in terms of enabling the risk of the option to be assessed.

These two properties of the option delta are intrinsic to the application and pricing and trading of options. The asset content implicit in delta allow the assessment and comparison of option strategies. An important aspect of this process is the ability to use the delta equivalence to compare positions as between the *asset* and the *option*.

This equivalence has an important implication for replication of options, which is central to the trading of options. The fact that delta provides the asset content as well as the probability of the option being exercised allows the option to be replicated by maintaining and dynamically adjusting a portfolio consisting of the asset and cash. This asset portfolio will, under certain conditions, give the same economic payoffs as the option being

replicated. This relationship is central to the derivation of the option value.² This process, which is referred to as delta hedging and forms the basis of all trading and risk management of options, is discussed in detail in Chapter 11.

4. GAMMA

4.1 Concept

Gamma is the second derivative of the option premium (delta is the first derivative) with respect to the underlying security price. It indicates how quickly delta will change as the underlying price changes, that is, the change in the price delta given the unit change in the underlying price.

The expected price change is given by:

Change in the asset price times Gamma = Expected Change in Option Delta

In *Exhibit 10.3* (above) the gammas are as follows:

Call option—0.0259

Put option—0.0259

This implies that for a small change of say 1.00 in the asset price the option value will change by:

Call option—0.0259 times 1.00 = 0.0259

Put option—0.0259 times 1.00 = 0.0259

Exhibit 10.8 shows the result of the change in terms of its impact on the premium, which is consistent with the predicted changes. The slight mis-estimation reflects the curvature of the option premium with respect to asset price, which gives gamma a certain curvature.

2. Black-Scholes and other option pricing approaches, such as the binomial models, are predicated on the concept of this replicating portfolio which is free of risk, which allows derivation of the option price as the combined portfolios should only yield the risk free rate of interest.

Exhibit 10.8
Option Gamma

OPTION PRICING MODEL—BLACK-SCHOLES
INPUTS/OPTION PRICING/SENSITIVITIES

PRICING INPUTS	
UNDERLYING ASSET PRICE	101.00
STRIKE PRICE	100.00
TRADE DATE	1-Jan-95
EXPIRY DATE	1-Jul-95
VOLATILITY	20.00%
RISK FREE RATE	10.00%
CALL (0)/PUT (1)	0
EUROPEAN (0)/AMERICAN (1) EXERCISE	0
INCOME ON ASSET	0%
MODEL OUTPUTS	
OPTION PREMIUM	CALL 8.91
OPTION PREMIUM (% OF ASSET PRICE)	CALL 8.82% PUT 3.07 3.04%

PREMIUM SENSITIVITIES/"GREEKS"	
DELTA	CALL 0.6891
GAMMA	0.0248
VEGA	0.25
THETA (pa)	11.14
THETA (per day)	0.03051
RHO	0.3009

Description of outputs	
A 1.00 increase in S will change premium by =	0.6891
A 1.00 increase in S will change delta by =	0.0248
A 1% increase in vol will change the premium by	0.2513
Divide by 365 to get daily time decay	11.1352
A 1% incr. in int rate will change by premium by	0.3009

Exhibit 10.8—continued


PREMIUM SENSITIVITIES/"GREEKS"	P/LI	Description of outputs
DELTA	-0.3109	A 1.00 increase in S will change premium by =
GAMMA	0.0248	A 1.00 increase in S will change delta by =
VEGA	0.25	A 1% increase in vol will change the premium by
THETA (pa)	1.62	Divide by 365 to get daily time decay
THETA (per day)	0.00444	
RHO	-0.1710	A 1% incr. in int rate will change by premium by

4.2 Behaviour

The gamma of an option is characterised by the following pattern of behaviour:

1. Deep out-of-the-money or in-the-money options have gammas close to zero because they are not very responsive to changes in the underlying asset price.
2. At-the-money options particularly where the time to maturity is relatively short have the highest gammas reflecting the fact the option is likely to be exercised or expire unexercised and the resulting option delta will go to either 1 (in the case of exercise) or 0 (in the case of non-exercise).
3. A low gamma indicates that the option delta changes slowly for a given change in the underlying asset price. A high gamma indicates that the option delta is very sensitive to changes in the price of the underlying asset.
4. Assets have minimal gamma. For most assets, the gamma is zero. The only exception to this is fixed interest securities which have some gamma because of the convex nature of price changes for given changes in interest rates.

Exhibit 10.9 sets out a gamma surface which illustrates the behaviour of the gamma for the option depicted in *Exhibit 10.3*.



4.3 Application

Gamma can be best regarded as a measure of the convexity of the option and the resultant hedging risk of replicating the option with a position in the asset through delta hedging.

In this regard it is a measure of hedging risk. It provides information about the sensitivity of the asset equivalent position of the option for given changes in asset prices. Similarly, it provides a measure of the convexity of the option position.

5. VEGA

5.1 Concept

Vega is the first derivative of the option premium with respect to the volatility. The vega measures the expected change in the value of the option for a given change in the volatility parameter.

The expected price change is given by:

Change in the volatility times Vega = Expected Change in Option premium

In *Exhibit 10.3* (above) the vegas are as follows:

Call option—0.26

Put option—0.26

This implies that for a small change of say 1.00% in the volatility the option value will change by:

Call option—0.26 times 1.00 = 0.26

Put option—0.26 times .10 = 0.26

Exhibit 10.10 shows the result of the change in terms of its impact on the premium, which is consistent with the predicted changes.

Exhibit 10.10
Option Vega

OPTION PRICING MODEL—BLACK-SCHOLES
INPUTS/OPTION PRICING/SENSITIVITIES

PRICING INPUTS	
UNDERLYING ASSET PRICE	100.00
STRIKE PRICE	100.00
TRADE DATE	1-Jan-95
EXPIRY DATE	1-Jul-95
VOLATILITY	20.00%
RISK FREE RATE	10.00%
CALL (0)/PUT (1)	0
EUROPEAN (0)/AMERICAN (1) EXERCISE	0
INCOME ON ASSET	0%
MODEL OUTPUTS	
OPTION PREMIUM	CALL 8.49 PUT 3.65
OPTION PREMIUM (% OF ASSET PRICE)	CALL 8.49% PUT 3.65%

PREMIUM SENSITIVITIES/"GREEKS"	
DELTA	CALL 0.6588
GAMMA	0.0248
VEGA	0.26
THETA (pa)	11.21
THETA (per day)	0.03071
RHO	0.2846

Description of outputs	
A 1.00 increase in S will change premium by =	0.6588
A 1.00 increase in S will change delta by =	0.0248
A 1% increase in vol will change the premium by	0.2584
Divide by 365 to get daily time decay	11.2098
A 1% incr. in int rate will change by premium by	0.2846

Exhibit 10.10—continued


PREMIUM SENSITIVITIES/"GREEKS"	PUT	Description of outputs
DELTA	-0.3412	A 1.00 increase in S will change premium by =
GAMMA	0.0248	A 1.00 increase in S will change delta by =
VEGA	0.26	A 1% increase in vol will change the premium by
THETA (pa)	1.69	Divide by 365 to get daily time decay
THETA (per day)	0.00464	A 1% incr. in Int rate will change by premium by
RHO	-0.1873	-0.1873

5.2 Behaviour

The vega of an option is characterised by the following pattern of behaviour:

1. Deep out-of-the-money or in-the-money options have lower vegas because they are not very responsive to changes in volatility, reflecting the fact that these options either have values dominated by the intrinsic value of the option or have low values.
2. At-the-money options, particularly where the time to maturity is long, have the highest vega.
3. A low vega indicates that the option delta changes slowly for a given change in the volatility. A high vega indicates that the option delta is very sensitive to changes in the price of the volatility.
4. Assets have zero vega. This reflects the fact that volatility is not a parameter relevant to pricing of *assets*.

Exhibit 10.11 sets out a vega surface which illustrates the behaviour of the vega for the option depicted in *Exhibit 10.3*.



5.3 Application

Vega measures the sensitivity to changes in volatility of a single option or a portfolio. In this regard it measures the exposure or risk to volatility changes present in the position under consideration.

6. THETA

6.1 Concept

Theta is the first derivative of the option premium with respect to the time to expiry. The theta measures the expected change in the value of the option for a given change in the option's time to expiry. The option theta is an annualised estimate and is usually divided by 365 days to get the daily theta estimate.

The expected price change is given by:

$$\text{Change in the time to expiry times Theta} = \text{Expected Change in Option premium}$$

In *Exhibit 10.3* (above) the vegas are as follows:

$$\begin{aligned} \text{Call option} & -0.03 \\ \text{Put option} & -0.004 \end{aligned}$$

This implies that for a small change of say 1 day (7 days) in the time to expiry the option value will change by:

$$\begin{aligned} \text{Call option} & -0.03 \text{ times } 1.00 = 0.03 \text{ (0.21)} \\ \text{Put option} & -0.004 \text{ times } 1.00 = 0.004 \text{ (0.03)} \end{aligned}$$

Exhibit 10.12 shows the result of the change in terms of its impact on the premium, which is consistent with the predicted changes.

Exhibit 10.12
Option Theta

OPTION PRICING MODEL—BLACK-SCHOLES
INPUTS/OPTION PRICING/SENSITIVITIES

PRICING INPUTS	
UNDERLYING ASSET PRICE	100.00
STRIKE PRICE	100.00
TRADE DATE	2-Jan-95
EXPIRY DATE	1-Jul-95
VOLATILITY	20.00%
RISK FREE RATE	10.00%
CALL (0)/PUT (1)	0
EUROPEAN (0)/AMERICAN (1) EXERCISE	1
INCOME ON ASSET	0%
MODEL OUTPUTS	
OPTION PREMIUM	CALL 8.20
OPTION PREMIUM (% OF ASSET PRICE)	PUT 3.39
	8.20%

PREMIUM SENSITIVITIES/"GREEKS"	
DELTA	CALL 0.6632
GAMMA	0.0260
VEGA	0.26
THETA (pa)	11.01
THETA (per day)	0.03017
RHO	0.2866

Description of outputs

A 1.00 increase in S will change premium by = 0.6632
 A 1.00 increase in S will change delta by = 0.0260
 A 1% increase in vol will change the premium by 0.2564
 Divide by 365 to get daily time decay 11.0106
 A 1% incr. in int rate will change by premium by 0.2866

Exhibit 10.12—continued

PREMIUM SENSITIVITIES/"GREEKS"	PUT	Description of outputs
DELTA	-0.3368	A 1.00 increase in S will change premium by =
GAMMA	0.0260	A 1.00 increase in S will change delta by =
VEGA	0.26	A 1% increase in vol will change the premium by
THETA (pa)	1.49	Divide by 365 to get daily time decay
THETA (per day)	0.00409	
RHO	-0.1828	A 1% incr. in int. rate will change by premium by

Exhibit 10.12—continued

OPTION PRICING MODEL—BLACK-SCHOLES
INPUTS/OPTION PRICING/SENSITIVITIES

PRICING INPUTS	
UNDERLYING ASSET PRICE	100.00
STRIKE PRICE	100.00
TRADE DATE	8-Jan-95
EXPIRY DATE	1-Jul-95
VOLATILITY	20.00%
RISK FREE RATE	10.00%
CALL (0)/PUT (1)	0
EUROPEAN (0)/AMERICAN (1) EXERCISE	0
INCOME ON ASSET	0%
	1
MODEL OUTPUTS	
OPTION PREMIUM	CALL 8.02
OPTION PREMIUM (% OF ASSET PRICE)	CALL 8.02% PUT 3.37
	PUT 3.37%
PREMIUM SENSITIVITIES/"GREEKS"	
DELTA	CALL 0.6607
GAMMA	CALL 0.0265
VEGA	CALL 0.25
THETA (pa)	CALL 11.11
THETA (per day)	CALL 0.03043
RHO	CALL 0.2767

Description of outputs
 A 1.00 increase in S will change premium by = 0.6607
 A 1.00 increase in S will change delta by = 0.0265
 A 1% increase in vol will change the premium by 0.2528
 Divide by 365 to get daily time decay 11.1075
 A 1% incr. in int rate will change by premium by 0.2767

Exhibit 10.12—continued

PREMIUM SENSITIVITIES/"GREEKS"	PUT	Description of outputs
DELTA	-0.3393	A 1.00 increase in S will change premium by =
GAMMA	0.0265	A 1.00 increase in S will change delta by =
VEGA	0.25	A 1% increase in vol will change the premium by
THETA (pa)	1.57	Divide by 365 to get daily time decay
THETA (per day)	0.00431	
RHO	-0.1778	A 1% incr. in int rate will change by premium by


6.2 Behaviour

The theta of an option is characterised by the following pattern of behaviour:

1. The theta is negative reflecting the loss of value as the time to expiry diminishes. This fall in value accrues to the seller as a gain which is offset by the loss suffered by the purchaser.
2. The theta of an option typically for a deep in or out-of-the-money option is lower than for a corresponding at-the-money option.
3. The highest theta is for an at-the-money option with a short time to expiry.

Assets, interestingly, will generally have thetas. The source of theta will vary. For example, options on debt instruments will exhibit a theta consistent with the accrual of interest income or expense on the instrument. Similarly, the impact of dividends or commodity interest rates will be reflected in the asset thetas for equities/equity indexes or commodities/commodity indexes. The changes in futures price reflecting changes in the carry cost as a result of changes in the time to delivery/settlement will provide futures/forward contracts with their theta.

Exhibit 10.13 sets out a theta surface which illustrates the behaviour of the theta for the option depicted in *Exhibit 10.3*.



6.3 Application

Theta measures the sensitivity to changes in the time to expiry of a single option or a portfolio. In this regard it measures the exposure or risk to the passage of time present in the position under consideration.

7. RHO

6.1 Concept

Rho is the first derivative of the option premium with respect to the risk free interest rate. The rho measures the expected change in the value of the option for a given change in the risk free interest rate.

The expected price change is given by:

Change in the risk free rate times Rho = Expected Change in Option premium

In *Exhibit 10.3* (above) the vegas are as follows:

Call option 0.29
Put option -0.18

This implies that for a small change of say 1% pa in the risk free rate the option value will change by:

Call option 0.29 times 1.00 = 0.29
Put option -0.18 times 1.00 = 0.18

Exhibit 10.14 shows the result of the change in terms of its impact on the premium, which is consistent with the predicted changes.

Exhibit 10.14
Option Rho

OPTION PRICING MODEL—BLACK-SCHOLES
INPUTS/OPTION PRICING/SENSITIVITIES

PRICING INPUTS	100.00	
UNDERLYING ASSET PRICE	100.00	
STRIKE PRICE	100.00	
TRADE DATE	1-Jan-95	
EXPIRY DATE	1-Jul-95	
VOLATILITY	20.00%	
RISK FREE RATE	11.00%	
CALL (0)/PUT (1)	0	1
EUROPEAN (0)/AMERICAN (1) EXERCISE	0	
INCOME ON ASSET	0%	
MODEL OUTPUTS	CALL	PUT
OPTION PREMIUM	8.52	3.21
OPTION PREMIUM (% OF ASSET PRICE)	8.52%	3.21%
PREMIUM SENSITIVITIES/"GREEKS"	CALL	
DELTA	0.6764	
GAMMA	0.0255	
VEGA	0.25	
THETA (pa)	11.60	
THETA (per day)	0.03179	
RHO	0.2932	

Description of outputs
 A 1.00 increase in S will change premium by = 0.6764
 A 1.00 increase in S will change delta by = 0.0255
 A 1% increase in vol will change the premium by 0.2530
 Divide by 365 to get daily time decay 11.6049
 A 1% incr. in int rate will change by premium by 0.2932

Exhibit 10.14—continued


PREMIUM SENSITIVITIES/"GREEKS"	PUT	Description of outputs
DELTA	-0.3236	A 1.00 increase in S will change premium by =
GAMMA	0.0255	A 1.00 increase in S will change delta by =
VEGA	0.25	A 1% increase in vol will change the premium by
THETA (pa)	1.19	Divide by 365 to get daily time decay
THETA (per day)	0.00326	
RHO	-0.1764	A 1% incr. in int rate will change by premium by

6.2 Behaviour

The rho of an option is characterised by the following pattern of behaviour:

1. The rho for a call (put) is positive (negative) reflecting the directional impact of an increase in interest rates on the value of the option.
2. The rho decreases with time to maturity, reflecting the diminished impact of the discounting of the exercise price.

Exhibit 10.15 sets out a rho surface which illustrates the behaviour of the rho for the option depicted in *Exhibit 10.3*.



6.3 Application

Rho measures the sensitivity to changes in the risk free interest rate of a single option or a portfolio. In this regard it measures the exposure or risk to changes in the discount rate present in the position under consideration.

8. INTERACTION OF RISK MEASURES

Exhibit 10.16 summarises the option sensitivities for underlying asset and option positions.



Exhibit 10.16				
Option Sensitivities—Summary				
Asset or Option Position	Delta	Gamma	Theta	Vega
Purchased Asset	Positive	Not applicable	Not applicable	Not applicable
Sold Asset	Negative	Not applicable	Not applicable	Not applicable
Purchased Call	Positive	Positive	Negative	Positive
Sold Call	Negative	Negative	Positive	Negative
Purchased Put	Negative	Positive	Negative	Positive
Sold Put	Positive	Negative	Positive	Negative

The analysis of the individual risk measures, to date, has been undertaken on a separate basis. In practice, there is substantial interaction between the risk measures. The most significant interactions are between delta, gamma, vega and theta.

Delta and gamma are both affected by changes in volatility. The pattern of interaction is as follows:

- Increases in volatility result in higher deltas for out and at-the-money options. Increases in volatility reduce the delta for in-the-money options. This pattern reflects the fact that delta is, in one sense, a measure of the probability of exercise of the option. The increase in volatility increases or decreases the probability of exercise for out and at-the-money options and in-the-money options respectively.
- Gamma generally decreases with an increase in volatility and vice versa. The change in gamma for a change in volatility is greatest for an at-the-money option, while the change in gamma is much lower for changes in volatility for in and out-of-the-money options.

The relationship between delta and volatility is set out in the delta-volatility surface set out in *Exhibit 10.17*. The relationship between gamma and volatility is set out in the gamma-volatility surface set out in *Exhibit 10.18*. In both cases, an option with a term to expiry of 30 days is utilised.



The implication of this relationship is that where the delta of an option or portfolio is utilised to hedge or replicate the underlying option, changes in volatility can significantly erode the efficiency of the hedge.

The relationship between gamma and theta is more complex. Where theta is large and positive, the corresponding gamma of the option also tends to be large and negative.³ This means that the gamma of a portfolio, under certain conditions, offsets the theta of that portfolio.

9. OTHER RISK PARAMETERS

A number of other risk parameters are utilised to obtain additional information about the sensitivity or risk profile of an option. These include: lambda (λ); speed; charm; colour; and fugit.

Lambda (λ) measures the percentage change in the value of the option for a percentage change in the asset price. It is usually calculated as:

$$\text{Lambda} = \text{Delta} \times (\text{Asset Price} / \text{Option Value})$$

It is effectively a measure of the leverage inherent in the option position. A large lambda indicates a greater sensitivity *in proportional terms* to movements in the asset price.

Speed, charm and colour were all derivatives suggested by Mark Garman.⁴ Each of these terms measures additional risk aspects of option transactions:

- *speed* is a measure of the change in gamma for a given change in the asset price;
- *charm* is a measure of the change in delta for a given change in time to expiry; and
- *colour* is a measure of the change in gamma for a given change in time to expiry.

Each of these measures, in particular, charm and colour, have significance for traders seeking to manage a portfolio of options.

Fugit was also suggested by Mark Garman.⁵ It is the expected value of the time to exercise of an American option. It is usually calculated using an iterative procedure within a binomial option pricing framework.

10. EXTENSIONS FOR OPTIONS ON DIFFERENT ASSETS

In the above analyses, the focus has been on a European option with non-income-paying asset. The same approach can be extended to cover options on different underlying assets. *Exhibit 10.19* sets out the partial derivatives for these types of options.

3. For a formal mathematical proof of this see John C Hull, *Options Futures and Other Derivatives* (3rd ed, Prentice Hall, Englewood Cliffs, NJ, 1996), pp 327-328.
4. See Mark Garman, "Charm School" (1992) 5(7) *Risk* 53.
5. See Mark Garman, "Semper Tempus Fugit" (1989) 2(5) *Risk* 34; "Semper Tempus Fugit" in "From Black Scholes To Black Holes" (1992) *Risk* 89.

Exhibit 10.19
Option Derivatives—Extensions

Option Derivative	European Option on a Non-asset Paying Asset (Black-Scholes)	European Option on Asset Paying Continuous Income (Y) (Amended Black-Scholes)	European Option on Forward/Futures Contract (Black)	European Option on Currency (Garman-Kohlhagen)
Delta—Call	$\Delta = N(d1)$	$\Delta = e^{-Y \cdot T} N(d1)$	$\Delta = e^{-Y \cdot T} N(d1)$	$\Delta = e^{-R_{f,T} \cdot T} N(d1)$
Delta—Put	$\Delta = N(-d1)$	$\Delta = e^{-Y \cdot T} N(-d1)$	$\Delta = e^{-Y \cdot T} N(-d1)$	$\Delta = e^{-R_{f,T} \cdot T} N(-d1)$
Gamma—Call & Put	$\gamma = N(d1)/S \cdot \sigma \cdot \sqrt{T}$	$\gamma = N(d1) \cdot e^{-Y \cdot T} / S \cdot \sigma \cdot \sqrt{T}$	$\gamma = N(d1) \cdot e^{R_{f,T} \cdot T} / F \cdot \sigma \cdot \sqrt{T}$	$\gamma = N(d1) \cdot e^{-R_{f,T} \cdot T} / S \cdot \sigma \cdot \sqrt{T}$
Vega—Call & Put	$\kappa = S \cdot \sqrt{T} \cdot N(d1)$	$\kappa = e^{-Y \cdot T} \cdot S \cdot \sqrt{T} \cdot N(d1)$	$\kappa = e^{R_{f,T} \cdot T} \cdot F \cdot \sqrt{T} \cdot N(d1)$	$\kappa = e^{-R_{f,T} \cdot T} \cdot S \cdot \sqrt{T} \cdot N(d1)$
Theta—Call	$\tau = (S \cdot \sigma \cdot N(d1) / 2 \sqrt{T}) - R_f \cdot K \cdot e^{-R_{f,T} \cdot T} N(d2)$	$\tau = (S \cdot \sigma \cdot N(d1) \cdot e^{-Y \cdot T} / 2 \sqrt{T}) - Y \cdot S \cdot e^{-Y \cdot T} N(d1) - R_f \cdot K \cdot e^{-R_{f,T} \cdot T} N(d2)$	$\tau = (F \cdot \sigma \cdot N(d1) \cdot e^{R_{f,T} \cdot T} / 2 \sqrt{T}) - R_f \cdot F \cdot e^{R_{f,T} \cdot T} N(d1) - R_f \cdot K \cdot e^{R_{f,T} \cdot T} N(d2)$	$\tau = (S \cdot \sigma \cdot N(d1) \cdot e^{-R_{f,T} \cdot T} / 2 \sqrt{T}) - R_{f,T} \cdot S \cdot e^{-R_{f,T} \cdot T} N(d1) - R_{f,T} \cdot K \cdot e^{-R_{f,T} \cdot T} N(d2)$
Theta—Put	$\tau = (S \cdot \sigma \cdot N(d1) / 2 \sqrt{T}) - R_f \cdot K \cdot e^{-R_{f,T} \cdot T} N(-d2)$	$\tau = (S \cdot \sigma \cdot N(d1) \cdot e^{-Y \cdot T} / 2 \sqrt{T}) - Y \cdot S \cdot e^{-Y \cdot T} N(-d1) - R_f \cdot K \cdot e^{-R_{f,T} \cdot T} N(-d2)$	$\tau = (F \cdot \sigma \cdot N(d1) \cdot e^{R_{f,T} \cdot T} / 2 \sqrt{T}) - R_f \cdot F \cdot e^{R_{f,T} \cdot T} N(-d1) - R_f \cdot K \cdot e^{R_{f,T} \cdot T} N(-d2)$	$\tau = (S \cdot \sigma \cdot N(d1) \cdot e^{-R_{f,T} \cdot T} / 2 \sqrt{T}) - R_{f,T} \cdot S \cdot e^{-R_{f,T} \cdot T} N(-d1) - R_{f,T} \cdot K \cdot e^{-R_{f,T} \cdot T} N(-d2)$
Rho—Call	$\rho = K \cdot T \cdot e^{-R_{f,T} \cdot T} N(d2)$	$\rho = K \cdot T \cdot e^{-R_{f,T} \cdot T} N(d2)$	$\rho = K \cdot T \cdot e^{-R_{f,T} \cdot T} N(d2)$	$\rho = K \cdot T \cdot e^{-R_{f,T} \cdot T} N(d2)$
Rho—Put	$\rho = K \cdot T \cdot e^{-R_{f,T} \cdot T} N(-d2)$	$\rho = K \cdot T \cdot e^{-R_{f,T} \cdot T} N(-d2)$	$\rho = K \cdot T \cdot e^{-R_{f,T} \cdot T} N(-d2)$	$\rho = K \cdot T \cdot e^{-R_{f,T} \cdot T} N(-d2)$
Rho—Call (Foreign Interest Rate)				
Rho—Put (Foreign Interest Rate)				

The options covered include:

1. European option where the underlying asset pays continuous income;
2. European option on a forward or futures contract; and
3. European option on a currency.

The risk measures for a European option on a commodity or equity would be identical to that for either an option where the underlying asset pays continuous income or an option on a currency with the commodity payout or convenience yield or equity dividend yield substituted for either the income of foreign interest rate.

11. SUMMARY

The derivatives of the valuation formula for options provide a quantifiable measure of the sensitivity and risk of the option to changes in one or more of the underlying factors which determine the price of the option. These risk measures are central to the management of the risk of options, hedging options, and using options to manage risk in asset portfolios.

In recent times, the derivative measures of risk have been used more generically to capture and present the risk of *all instruments*. This encompasses *assets* and *forwards*, as well as options. This reflects the fact that the risk or value sensitivities of all instruments are capable of being measured, expressed, and communicated easily using this general vocabulary of risk. *Exhibit 10.20* sets out this generalised language of risk which is increasingly utilised.

Exhibit 10.20
Generalised Risk Measures

Risk Measure	Concept for Options	Concept for General Risk
Delta (Δ)	Delta is the derivative of the option pricing formula with reference to the asset price (S). It measures the estimated change in the option premium for a change in S.	Measures exposure to price change of underlying asset; equivalent to Present Value of Basis Point (PVBP) or Dollar Value of 1 Basis Point (DVO1).
Gamma (γ)	Gamma is the second derivative of the option pricing formula with reference to the asset price (S). It measures the estimated change in the delta of the option for a change in S.	Measures exposure to change in delta; equivalent to measure of convexity.
Vega (κ)	Vega is the derivative of the option pricing formula with reference to the volatility of the asset returns (σ). It measures the estimated change in the option premium for a change in σ .	Measures exposure to changes in volatility (only applicable to options).
Theta (τ)	Theta is the derivative of the option pricing formula with reference to the time to option expiry (T). It measures the estimated change in the option premium for a change in T.	Measures exposure to or change in value arising from the effluxion of time; analogous to carry income or expense.
Rho (ρ)	Rho is the derivative of the option pricing formula with reference to the risk free rate (R_f). It measures the estimated change in the option premium for a change in R_f .	Measures exposure to changes in the discount rate(s) applicable.

Chapter 11

Option Replication Utilising Delta Hedging

by Satyajit Das

1. INTRODUCTION

It is feasible to replicate or synthesise the economic return profile of an option transaction by a process of trading in the underlying asset. The capacity to replicate the option using the underlying asset is implicit in the inherent nature of delta which, amongst other information, provides a guide as to the asset content of the option. This concept which is referred to variously as delta hedging, option replication, synthetic options, or dynamic hedging as well as a number of proprietary terms, is central to both the pricing and trading/hedging of options.

The requirement to synthetically replicate the economic profile of an option rather than trading in the underlying option may arise for a number of reasons. These include:

- Financial institutions active in trading may enter into transactions with clients where they either sell or buy options which they must offset the market risk. The dealer can offset the client position with an offsetting option transaction (either on an exchange or over-the-counter). However, where it is not possible to offset the exposure in this way, it may be necessary for the dealer to hold and hedge the position either till maturity or, more often, until an offsetting transaction can be entered into. In the period between the initial transaction and the offsetting transaction, the option dealer would delta hedge its exposure to reduce any exposure to movements in the price of the underlying asset.
- Participants may need to synthesise options where, first, such options are not readily available or traded, or, secondly, it is more cost effective to create such options rather than purchasing the options.
- The creation of structured investment products such as capital protected funds or portfolio insurance products which require either the creation or purchase of options. In these situations, the fund may prefer to synthesise the option for cost or other reasons, such as customisation et cetera. The use of synthetic option technology in portfolio insurance applications is discussed in Chapter 14.

In this chapter, the process of option replication is examined. The process of delta hedging is first explained with a series of examples of increasing complexity. This analysis focuses on the cost of hedging and the risks inherent in synthetically creating the option. The risk and risk management of this methodology for trading and hedging options is then considered.

2. DELTA HEDGING¹

2.1 Concept

The concept of utilising a position in the underlying asset to replicate the economic profile of an option is central to both the *pricing or valuation* of an option and the trading and hedging of these instruments.

As noted previously, delta, that is, the partial derivative of the option valuation formula with respect to asset price, provides a measure of the change in the option premium for a small change in the asset price. This allows a portfolio, consisting of the asset and the option, to be constructed which is immune in relation to value to small changes in the asset price.

For example, a short call position may be hedged with a long position in the asset. The amount of asset held would be equivalent to delta times the face value of the option.² For small changes in the asset price, the change in the value of the asset portfolio would offset equal and opposite changes in the value of the option portfolio. This portfolio would economically be a riskless portfolio which to avoid arbitrage would only earn the risk free rate of interest on the investment in the portfolio. This basic insight is utilised in all option pricing models, such as Black-Scholes and the binomial pricing approaches, to derive the fair value of the option. *Exhibit 11.1* sets out the riskless fully-hedged portfolio positions possible.

Exhibit 11.1
Riskless Hedge Positions

Position	Hedge
long position in calls	short Δ assets for each call held
short position in calls	long Δ assets for each call sold
long position in puts	short Δ assets for each put held
short position in puts	long Δ assets for each put sold

Delta (Δ)³ refers to the sensitivity of the option premium to changes in the asset price.

Implicit in this approach is the capacity for the final payoff from the option to be replicated by trading in the asset underlying the option to create an instantaneous hedge against small movements in asset price. This approach can then be extended by dynamically adjusting the hedge by altering the holding of the asset and the amount of the borrowing and lending to replicate the option until maturity. This approach, which has been

1. For an excellent analysis of delta hedging see Nassim Taleb, *Dynamic Hedging: Managing Vanilla and Exotic Options* (John Wiley & Sons, New York, 1997).
2. Alternatively, the position could be constructed as for each unit of asset held, it would be necessary to sell $1/\text{delta}$ of the call option on the asset.
3. See detailed discussion in Chapter 10.

recognised by a number of writers, allows the synthetic creation of an option by holding delta weight of the asset underlying the option and adjusting the holding of the asset in accordance with changes in the delta of the option.⁴ The synthetic option can therefore be created from a portfolio of existing traded instruments, which, with proper management over time, can replicate the return characteristics of an option. This is what is referred to as delta hedging.

2.2 Example of delta hedging

Delta hedging entails the creation of a synthetic option by using a portfolio consisting of two instruments:

1. the asset into which the option can be exercised, whether it be a cash market instrument or futures and forwards on the underlying asset; and
2. a risk-free asset, usually cash or high quality securities.

The key to creating a synthetic option is to determine the proportion of cash and asset to maintain in the portfolio. This proportion is adjusted through time in a very specific way to replicate the price behaviour of a call option. In practice, such a portfolio can be created and properly managed to give approximately the same premium and outcome as a traded call option on the underlying asset.

The intuition for synthetic options derives from the fact that, for example, the price behaviour of a call option is similar to that of a portfolio with combined positions involving the underlying asset and cash. Although the price of a call option and the price of the underlying asset change in the same direction, the effect on the price of the call option of a given change in the asset price depends on the current price level of the asset. This is because the number of units of the asset held in the replicating portfolio must be sufficient to equate to the slope of the call option price curve at that particular price level or the particular option's delta.

Exhibit 11.2 sets out a simple example of a delta hedge to replicate a call option.

4. See, for example, Mark Rubinstein and Hayne E Leland (July-August 1991).

Exhibit 11.2
Delta Hedge—Example 1

Assume a call option on a non-income-producing asset where the asset is trading at 100. The call option is on the following terms:

- Strike prices = 100
- Time to expiry = 0.50
- Risk free rate = 10% pa
- Volatility = 20% pa

The call option premium (using Black Scholes) is 8.26. The option delta is .6641.

Using this information two portfolios can be constructed:

1. a short position in the options, say 100 options (each on 1 unit of the asset); and
2. a long position in 66.41 units of the asset (calculated as delta times the face value of the options).

The changes in value of the two portfolios for small changes in the asset price are set out below

Increase In Asset Price (from 100 to 100.10)

Initial Portfolio			Final Portfolio			
Portfolio 1 = Option			Asset Price = 100.10			
Asset Price = 100			Asset Price = 100.10			
Amount	Value (per Option)	Portfolio Value	Amount	Value (per Option)	Portfolio Value	Change in Portfolio Value
100	8.26	826	100	8.33	833	-7
Portfolio 2 = Assets			Asset Price = 100.10			
Asset Price = 100			Asset Price = 100.10			
Amount	Value (per Option)	Portfolio Value	Amount	Value (per Option)	Portfolio Value	Change in Portfolio Value
66.41	100	6641	66.41	100.10	6648	7
Change In Overall Portfolio Value			0			

Decrease In Asset Price (from 100 to 99.90)

Initial Portfolio			Final Portfolio			
Portfolio 1 = Option			Asset Price = 99.90			
Asset Price = 100			Asset Price = 99.90			
Amount	Value (per Option)	Portfolio Value	Amount	Value (per Option)	Portfolio Value	Change in Portfolio Value
100	8.26	826	100	8.20	820	6
Portfolio 2 = Assets			Asset Price = 100.10			
Asset Price = 100			Asset Price = 100.10			
Amount	Value (per Option)	Portfolio Value	Amount	Value (per Option)	Portfolio Value	Change in Portfolio Value
66.41	100	6641	66.41	99.90	6635	-6
Change in Overall Portfolio Value			0			

As is evident, the portfolio of the short call options and the offsetting portfolio of the delta amount of the asset is insulated from small changes in asset price. The change in the value of the options for a given movement in the asset price is offset by an equal but opposite change in the value of the asset. While the example focuses on short call position, an identical logic can be utilised in relation to all option positions. For example, the technique for replicating put options is similar. It entails maintaining a portfolio consisting of cash and a short position in the underlying asset, which is adjusted as the price of the asset changes.

The replicating portfolio must be adjusted as the asset price changes. This will usually entail, in the case of a call option, selling assets as the asset price falls and buying assets as the asset price rises. If the asset price declines, then the holding of the assets is sold, with the proceeds being used to partially repay the borrowings used to finance the position. If the asset price at maturity is below the strike price, then the replication portfolio should end up with no holding of the asset at maturity and no corresponding borrowing. This portfolio would have a value of zero. This corresponds to the value of the option. If the asset price increases, then the holding of the assets is increased, with borrowings being undertaken (for the value of the asset less the premium received) to finance the position. If the asset price at maturity is above the strike price, then the option will be exercised. The replication portfolio should equal the units of asset underlying the option financed by borrowings, which is equal to the strike price. The difference corresponds to the intrinsic value of the option.

The portfolio is never fully invested when the asset price increases, nor fully disinvested when the asset price falls. The process of portfolio adjustment will reduce the initial investment. Theoretically, by the expiration of the call option, the cumulative depreciation should approximately equal the initial theoretical value of the call option.

From a theoretical standpoint, the process of delta hedging entails the final option payoff being replicated by the position in the stock financed by borrowings and the option premium received. This can be seen by examining a model such as Black-Scholes, which can be reduced as follows:

$$\text{Option Premium} = \text{Delta times Asset} - \text{Amount of Borrowing}$$

This more formally can be given as:

$$P_{ce} = S \cdot N(d1) - K e^{-Rt} \cdot N(d2)$$

The asset holding is dictated by the delta or sensitivity to asset price changes of the option. The amount of borrowings is equivalent to the present value of the exercise price of the option adjusted for the probability that the call option expires at maturity in-the-money. The rationale, in a risk-neutral and arbitrage-free world, being that the expected amount due for repayment to the lender is equivalent to this amount.

Two aspects of the replication process should be noted:

1. the replication portfolio is inherently self financing; and
2. the replication portfolio must be dynamically managed reflecting changes in the factors affecting the option price leading to changes in the relative holding of the asset and the amount of borrowings.

The concept of creating synthetic options permits option granters to replicate not only call options, but many other option positions. Using replicating portfolios the granter of options can, where it has created a risk position by writing options, cover these open positions, creating synthetic options which hedge the existing exposure.

2.2 Costs of replication

The process of replication is not costless. The major costs include:

- financing of the asset position or investment of the proceeds of a short sale;
- trading costs; and
- loss of hedge efficiencies.

The financing cost element is self-explanatory. The holding of assets required to hedge a short call option or a long put option will require financing, with the resulting interest expense being incurred. In contrast, a long call or short put option will require short selling the asset, releasing cash which can be invested and earn interest.

Trading costs cover a variety of items:

- The transaction costs of trading—the bid-offer spread—in actual markets means that each transaction results in a cost to the trader seeking to replicate the option.
- The gains or losses on trading—the process of replication requires the trader to, for a call option, buy high, sell low in the process of dynamic hedge management whereby the holding of assets must be increased as the price goes up and decreased as the price of the asset falls. For a put option, a similar process is applicable, with the short position in the asset having to be increased as the price decreases and decreased as the price increases, resulting in the same pattern of selling low and buying high. This pattern of trading inherently creates losses which will represent the cost of replicating the option.

Loss of hedge efficiency covers a variety of items:

- Delta slippage—as discussed in greater detail below in the context of risk management, it may not be possible to maintain a delta neutral position under all circumstances. This may be the result of either the use of *periodic rebalancing* rather than *continuous rebalancing*, reflecting the incorporation of transaction and trading costs as well as price movements which are discontinuous and introduce lags in adjustment of the delta hedge. The failure to maintain perfect delta neutrality creates exposures to the movement in the *price of the underlying asset* with resulting gains and losses.
- Rebalancing costs—these reflect the impact of the decision to periodically rebalance the portfolio, which creates the delta slippage noted above and the resultant gains and losses from exposure to asset price movements. In this regard, all delta hedging is a compromise between increased frequency of rebalancing the hedge (which incurs additional trading costs) and less frequent rebalancing (which reduces efficiency of the hedge and exposes the portfolio to gains and losses from changes in the asset price).

In perfect, frictionless capital markets, the costs of the delta hedge should exactly equal the theoretical premium of the option. In practice, the factors identified make this unlikely.

2.3 Issues in delta hedging

As illustrated above, in most situations, the value of an option behaves in a manner which is similar to a portfolio consisting of short or long positions in the asset and cash borrowings or investments. This allows the process of replication described.

However, when put to practice, the theory of delta hedging suffers from a number of difficulties:

1. market structure issues, such as the capacity to borrow and finance or lend the asset;
2. the impact of transaction costs;
3. the pattern of changes in the asset price, such as the absence of large jumps or non-continuous price changes;
4. the constancy of volatility and interest rates over the life of the option; and
5. issues arising from the use of forwards or futures contracts on the asset to replicate the option.

The process of option replication inherently assumes that, where the delta hedge requires holding a position in the asset, it is possible to finance. Similarly, where the delta hedge requires a short sale of the asset, it is assumed that, first, it is feasible to effect the short sale, and, secondly, it is possible to borrow the asset to implement the short sale. These conditions may not be able to be met in every market structure.

As noted above, the requirement to trade in the underlying asset results in the trader incurring trading costs. The trading costs are related to the frequency of re-balancing of the hedge. The more frequent the re-balancing the more accurately the replicating portfolio tracks the underlying option being hedged. However, the frequency of trading obviously affects the cost of replication by way of increased transaction costs. A related problem is the inability to accurately predict the *level* of transaction costs. The position in the asset must be adjusted periodically as the option's delta changes. As the delta changes are unknown (being a function of the path of asset prices), the exact number or quantum of hedge adjustments is not known in advance. This means that the transaction costs are not known in advance.

The ability to replicate the value changes in the option through trading in the asset will only be applicable for small changes in the asset price and where the asset price changes are continuous. A large discrete change in the asset price—in effect, a discontinuous movement or gap—will significantly impair the effectiveness of the hedge. This reflects the fact that in the event of a discontinuity or gap the change in the asset position cannot be undertaken sufficiently quickly. The price of the option will adjust immediately but the hedge cost will lag behind the change in the asset price, creating hedging errors and a divergence in the value of the hedge relative to the option.

The theoretical model of delta hedging assumes that both the volatility of asset price changes is negligible and the risk-free interest rate is constant over the life of the option. In reality, both of these variables are uncertain. Changes in each of these terms will, in fact, impact on the efficiency of the hedge. For example, a change in volatility levels may not be accompanied by a change in the price of the asset. In these circumstances, the value of the option will change but will not be accompanied by a corresponding change in the value of the asset position. Changes in volatility will also alter the option delta, resulting in changes in the required hedge which will have an impact on the cost of replicating the option.

Changes in interest rates will affect the process of delta hedging in a number of ways. Changes in interest rates will alter the forward price of the asset and hence the value of the option. In addition, it will affect the cost of the hedge, as it will increase the cost of holding the asset and reduce the cost of shorting.

The final issue relates to the use of forwards or futures on assets to replicate the option. The major rationale for this is the off-balance sheet nature of these instruments and the accompanying efficiencies in the use of capital. However, the use of forwards or futures introduces additional risk factors into the hedging process. The use of forward/futures contracts assumes the fair value pricing of these contracts; that is, their prices are free of arbitrage. In reality, the forward contract may be priced away from fair value, thereby introducing additional errors in the hedging process. The changes in the cash-futures basis (effectively, the cost of carry) may result in higher hedging costs.

Similarly, where the forward or futures contract used is of a different maturity to that of the option, the maturity mismatch introduced creates an exposure to changes in the shape of the yield curve. This reflects the fact that changes in shorter term rates will affect the forward/futures contract value differently from the impact of longer-term rates on the value of the underlying option.

2.4 More complex examples of delta hedging

The concept of utilising option replication techniques is further illustrated in *Exhibit 11.3* with a series of more complex examples. These examples are designed to highlight some of the risks and difficulties in seeking to replicate options through trading in the underlying assets.

Exhibit 11.3

Delta Hedge—Example 2

DELTA HEDGING

14-Aug-96

CALL OPTION EXAMPLE (OPTION EXPIRES IN-THE-MONEY)

INPUTS	
ASSET PRICE	\$100.00
STRIKE PRICE	\$100.00
TRADE DATE	01-Jan-96
EXPIRY DATE	29-Jan-96
TIME TO EXPIRY (DAYS)	28.00
VOLATILITY (INITIAL)	20.00%
INTEREST RATE	10.00%
HOLDING COST	0.00%

HEDGE PROFIT AND LOSS	
PREMIUM RECEIPT	\$21,928
INTEREST ON PREMIUM	\$168
HEDGE COSTS	(\$19,554)
NET	\$2,541

NO OF ASSETS:	10,000	10,000	10,000	10,000	10,000
DATE	01-Jan-96	08-Jan-96	15-Jan-96	22-Jan-96	29-Jan-96
TIME TO EXPIRY (DAYS)	28.00	21.00	14.00	7.00	0.00
ASSET PRICE	\$100.00	\$101.50	\$99.00	\$100.50	\$102.50
VOLATILITY (CURRENT)	20.00%	20.00%	20.00%	20.00%	20.00%
INTEREST RATE (CURRENT)	10.00%	10.00%	10.00%	10.00%	10.00%
OPTION PREMIUM	\$2,1928	\$2,7542	\$1,1014	\$1,3729	\$2,5000
DELTA	0.5071	0.6273	0.4048	0.5758	1.0000
GAMMA	0.0714	0.0770	0.0996	0.1404	0.0000
THETA (per day)	0.0385	0.0427	0.0532	0.0773	0.0000
VEGA	0.1096	0.0913	0.0749	0.0544	0.0000
RHO	-0.0017	-0.0016	-0.0004	-0.0003	0.0000

Exhibit 11.3—continued

DELTA HEDGE HEDGE REQUIREMENT	5,071	6,273	4,048	5,758	10,000
HEDGE TRANSACTIONS					
PURCHASE ASSETS	\$507,143	\$507,143	\$405,935	\$405,935	\$405,935
PURCHASE ASSETS	1,202	\$121,980			
PURCHASE ASSETS	(2,225)				
PURCHASE ASSETS	1,710			\$171,860	\$171,860
PURCHASE ASSETS	4,242			\$434,829	\$434,829
VALUE OF HEDGE PORTFOLIO					
NO OF ASSETS	5,071	6,273	4,048	5,758	10,000
AVERAGE VALUE OF ASSETS	\$100.00	\$100.29	\$100.29	\$100.35	\$101.26
TOTAL	\$507,143	\$629,123	\$405,935	\$577,795	\$1,012,624
GAIN ON HEDGE ADJUSTMENT		(\$972.60)	(\$1,206.54)	(\$778.51)	(\$1,108.10)
INTEREST COST		(\$972.60)	(\$5,044.13)	(\$5,822.63)	(\$19,554.48)
CUMULATIVE COST					