BUTTERWORTH-HEINEMANN FINANCE

CD ROM INCLUDED

QUANTITATIVE FINANCE SERIES

derivative instruments

a guide to theory and practice

Brian A. Eales & Moorad Choudhry

Foreword by Kenneth Garbade, Federal Reserve Bank of New York

Derivative Instruments: A Guide to Theory and Practice

Butterworth-Heinemann Finance

QUANTITATIVE FINANCE SERIES

Aims and objectives

- books based on the work of financial market practitioners, and academics
- presenting cutting edge research to the professional/practitioner market
- combining intellectual rigour and practical application
- covering the interaction between mathematical theory and financial practice
- to improve portfolio performance, risk management and trading book performance
- covering quantitative techniques

Market

Brokers/Traders; Actuaries; Consultants; Asset Managers; Fund Managers; Regulators; Central Bankers; Treasury Officials; Technical Analysts; and Academics for Masters in Finance and MBA market.

Series titles

Return Distributions in Finance
Derivative Instruments: a guide to theory and practice
Managing Downside Risk in Financial Markets: theory, practice, & implementation
Economics for Financial Markets
Performance Measurement in Finance: firms, funds and managers
Real R&D Options
Forecasting Volatility in the Financial Markets
Advanced Trading Rules
Advances in Portfolio Construction and Implementation

Series editor

Dr Stephen Satchell
Dr Satchell is a Reader in Financial Econometrics at Trinity College, Cambridge; Visiting Professor at Birkbeck College, City University Business School and University of Technology, Sydney. He also works in a consultative capacity to many firms, and edits the journal *Derivatives: use, trading and regulations* and the *Journal of Asset Management.*

Derivative Instruments: A Guide to Theory and Practice

Brian A. Eales
Moorad Choudhry

BUTTERWORTH HEINEMANN

OXFORD AMSTERDAM BOSTON LONDON NEW YORK PARIS
SAN DIEGO SAN FRANCISCO SINGAPORE SYDNEY TOKYO

Butterworth-Heinemann
An imprint of Elsevier Science
Linacre House, Jordan Hill, Oxford OX2 8DP
200 Wheeler Road, Burlington MA 01803

First published 2003

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication Data
A catalogue record for this book is available from the Library of Congress

ISBN 0 7506 5419 8

For information on all Butterworth-Heinemann finance publications
visit our website at www.bh.com/finance

Typeset by Integra Software Services Pvt. Ltd, Pondicherry, India. www.integra-india.com
Printed and bound in Great Britain

To my wife and family for their continuing support.
Brian A. Eales

To anyone who's never let me down.
Moorad Choudhry

This Page Intentionally Left Blank

Contents

Foreword

Derivative instruments have a remarkable history, a ubiquitous present and a promising future.

First and foremost, derivatives have materially enhanced the ability of market participants to isolate, price and trade a wide variety of risks, including interest rate risk, foreign exchange rate risk, equity risk and credit risk. The idea that someone might buy a convertible bond and sell off the embedded cash flows, credit exposure and equity participation in separate pieces was hardly dreamed of 30 years ago, yet has become a driving force in the equity-linked markets during the past several years. Connections between credit risk and equity risk that were once matters of theory have become issues of profit and loss as market participants spread one against the other – without ever trading the underlying instruments. And by lowering transaction costs, derivatives have opened global capital markets to a far wider array of participants, enhancing liquidity, facilitating diversification and sharpening prices. Derivatives have even forced us to re-examine our understanding of the roles played by different market participants. Labelling an institution as a financial intermediary used to say something about its standing between capital surplus units and capital deficient units. Today the label is as likely to say something about its standing between those that want to shed risk and those that want to bear risk.

The breadth and complexity of the topics addressed in this book testify to the importance of derivatives in contemporary financial markets. From relatively simple forward rate agreements and interest rate swaps to swaptions and collateralised debt obligations, derivatives are the work horses of finance; neither investment management nor corporate finance can be appreciated without understanding their use and pricing. When a transaction can be structured in any of several different ways, it is imperative that analysts be able to see through the form to the substance in sorting out the alternatives.

Looking to the future, it appears that derivatives can continue to expand in at least two dimensions. First, they can be used to price and transfer more different kinds of risk. The relatively young markets in weather derivatives and economic derivatives come to mind as examples. More speculatively, one might imagine derivatives for isolating, pricing and transferring risks associated with uncertainty over discretionary corporate actions like recapitalisations, dividends and spin-offs. Second, derivatives can be used to intermediate more intensively similar risks in different packages, leading to more homogeneous and transparent pricing of the common risks borne by different classes of securities issued by a single corporation.

Whether one is interested in current or prospective methods in investments and finance, acquiring a solid grounding in derivatives is crucial. Messrs Choudhry and Eales here offer a wide-ranging, yet detailed, guide for practitioners and practitioners-to-be. I enjoyed reading it, and I hope you do too.

Kenneth D. Garbade
Research and Market Analysis Group
Federal Reserve Bank of New York

The views expressed are those of the author and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

Preface

This book is a concise practitioners' guide to the main derivative instruments used in financial markets today. It should also be suitable for graduate students undertaking courses in business finance, economics and financial engineering. The focus is on analysis and valuation techniques, presented for the purposes of practical application. Hence institutional and market-specific data is largely omitted for reasons of space and clarity, as this is abundantly available in existing literature. Students and practitioners alike should be able to understand and apply the methods discussed here. The book attempts to set out a practical approach in presenting the main issues and the reader should benefit from the practical examples presented in the chapters.

The contents are presented as one unified volume, with coverage of interest rate derivatives such as bond futures and interest rate swaps, and equity derivatives. There is also a detailed introduction to credit derivatives. The final chapter looks at structured products such as convertible bond structures and guaranteed equity products. The authors hope that readers find the book thought provoking and that it enables them to become familiar with practical concepts of importance in the derivatives markets.

Comments on the text are welcome and should be sent to the authors care of Butterworth-Heinemann Finance.

Layout of the book

This book is organised into 12 chapters. The first chapter presents a broad-based introduction to derivatives. Chapters 2–6 look at interest rate derivatives, including futures and swaps. Chapter 7 is an introduction to credit derivatives, and contains coverage of synthetic collateralised debt obligation (CDO) structures, which use credit derivatives to transfer credit risk, or for credit speculation purposes. Equity derivatives are considered in Chapters 8 and 9, while option instruments are considered in Chapters 10–11. The final chapter looks at convertible bonds and other structured products.

Students are encouraged to use the spreadsheet examples, contained on the CD with this book, as part of their study into financial engineering as used in the markets.

Further market research can be viewed at
www.YieldCurve.com

Acknowledgements

Thanks to everyone that helped me out while I worked on this book. No thanks to those who doubted, bureaucrats and the rest. Help.

Moorad Choudhry
Surrey, England
October 2002

About the Authors

Brian A. Eales is a principal lecturer in the Department of Economics, Finance and International Banking at London Metropolitan University. He lectures on risk hedging and financial engineering using derivative securities on Masters courses at London Metropolitan, City University Business School and Stockholm University. His publications to date include texts on Financial Risk Management and Financial Engineering as well as articles on the use of derivatives in several practitioner journals.

Moorad Choudhry is a vice-president in structured finance services with JPMorgan Chase Bank in London. He joined the City in 1989 from the University of Reading, and worked as a gilt-edged market maker and treasury trader at ABN Amro Hoare Govett Sterling Bonds Limited, and as a sterling proprietary trader at Hambros Bank Limited, before joining JPMorgan Chase in February 2000. He is a Fellow of the Securities Institute and a member of the Chartered Institute of Bankers. He is a Senior Fellow at the Centre for Mathematical Trading and Finance, CASS Business School, London.

Moorad lives in Surrey, England.

Do not worry about your difficulties in mathematics. I can assure you mine are still greater.

Albert Einstein (1879–1955)

1 Introduction to Derivatives

Derivative instruments have been a feature of modern financial markets for several decades. They play a vital role in managing the risk of underlying securities such as bonds, equity, equity indexes, currency, short-term interest rate asset or liability positions. In the commodity markets they have, in general, been around for a great deal longer. In modern times the Chicago Board of Trade, for example, was set up in 1848 for the exchange trading of agricultural products such as wheat and corn. The Exchange put in place a mechanism that would play an important role in helping the agricultural community to plan for the future by enabling users of derivatives to lock-in the prices they will receive for their goods before they were even ready for harvesting. In 1865 the Chicago Board formally established its General Rules. This opened the floodgates for spot and forward trading of commodities – which ultimately would be delivered to an end user – and in 1870 the New York Cotton Exchange was established.

Financial futures contracts, where in many cases no delivery of a physical security is involved (rather settlement is in the form of a cash payment), had to wait another century before taking off. However, once present they succeeded in generating trading volumes in stock indexes, stocks, foreign exchange, bonds and short-term interest rate instruments to unprecedentedly high levels.

Just what are these derivatives, though, and what role or roles do they play? A working definition of a derivative, which will help lay the foundation of this text, is: "an instrument whose existence and value is contingent upon the existence of another instrument or security". The security on which the derivative is created may itself be a derivative and this can give rise to potentially dangerous inverse financial pyramids.

The major derivative instruments, which in some respects may be regarded as building blocks, can be categorised as follows:

futures
forwards
swaps, and
options

Each instrument has its own characteristics, which offer advantages in using them but bring with them disadvantages. The disadvantages may not always be apparent to the end user and these days it is crucial that end users are made aware of the risks associated with the derivative contracts they enter into and are made aware of the instrument's appropriateness for the purpose it is to perform.

Why should market participants seek to use derivatives? There are several answers to this question: derivatives may be used to:

- speculate;
- hedge a portfolio of shares, bonds, foreign currency, etc.;
- undertake arbitrage – i.e. benefit from mispricing;
- engineer or structure desired positions.

Each of these functions will be described in detail in later chapters.

One justification for their use in hedging can be found in the domain of portfolio diversification. The classical portfolio allocation methods of Markowitz and Sharpe, and extensions of those methodologies provide accepted, well-documented cases for creating a collection of risky assets and combining them in a portfolio so that the risk of loss is spread, or diversified, over the constituents of that portfolio.[1] Adopting such a strategy will result in risk reduction as illustrated in Figure 1.1.

From Figure 1.1 it can be seen that, up to a point, the inclusion of additional risky assets reduces the overall risk of the portfolio. However, beyond a certain point risk reduction hardly occurs but an identifiable risk still remains. This component of risk is known by

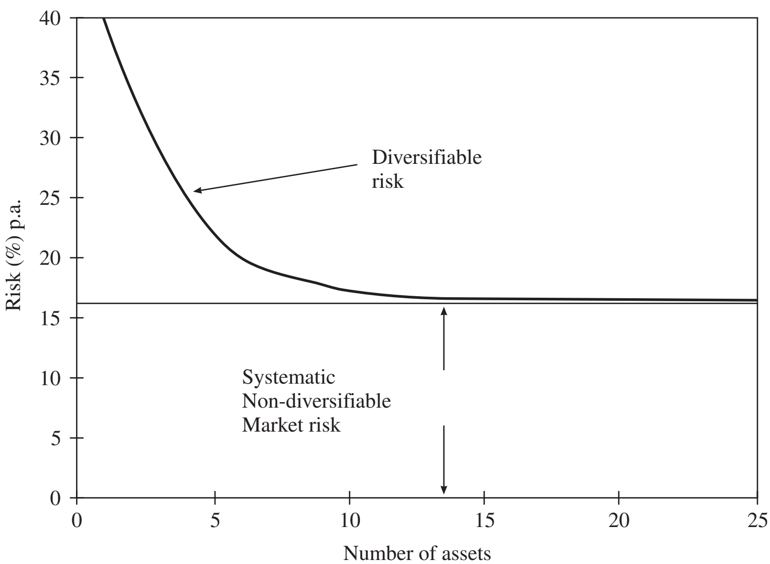

Figure 1.1: Effects of diversification on portfolio risk

[1] The interested reader is referred to classic texts such as Levy, H. & Sarnat, M., *Portfolio and Investment Selection: Theory and Practice*, Prentice Hall (1984) or Elton, E. J. & Gruber, M. J., *Modern Portfolio Theory and Investment Analysis*, 4th edition, Wiley (1991), for background on approaches to portfolio asset allocation.

several names: market risk, systematic risk and non-diversifiable risk, as indicated on the figure. It is the hedging of this form of risk to which some derivatives are directed.

These days a great deal of effort is directed towards constructing portfolios that mimic the behaviour of a specially constructed index with minimum tracking error. Such portfolios are described as index-tracking portfolios and – in the case of zero-tracking error[2] – enjoy a portfolio beta of one (i.e. $\beta_p = 1$). They replicate market movements up and down exactly. Even if $\beta_p \neq 1$ but is very close derivatives on the appropriate index should be able to serve as very good hedge vehicles. This idea will be examined in Chapter 8.

To demonstrate this idea consider a portfolio's exposure to market risk. For many short-term interest rate, equity, foreign exchange positions, etc., this can be expressed in the form of a simple profit/loss diagram as depicted in Figure 1.2.

If the current value of the portfolio is based on an index value of 5000 any positive movement in the index will result in an increase in the value of the portfolio while any fall from the value of 5000 will see the value of the portfolio decline.

An obvious way of countering the effect of a downward movement in the index is by adding to the portfolio a security whose value will fall when the index rises. The payout of such an instrument is exhibited in Figure 1.3.

By combining the long position depicted in Figure 1.2 and the short position in Figure 1.3 a flat position is obtained, the downside risk has been removed but the upside profit potential has also been eliminated. In other words it does not matter whether the index rises or falls, the value of the portfolio today (5000) is, to a great extent, assured – at least in theory and an ideal world.

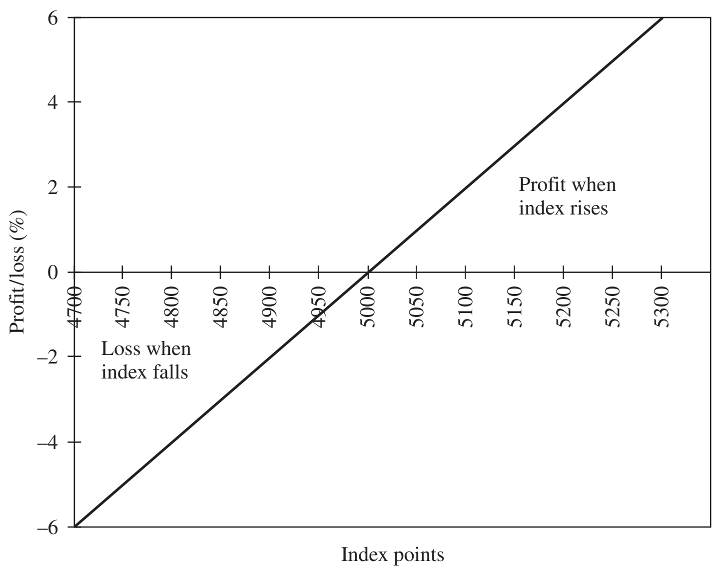

Figure 1.2: Long equity index position

2 It is not feasible to construct and maintain a perfect index-tracking portfolio. A defined minimum tracking error is the goal frequently sought by fund managers.

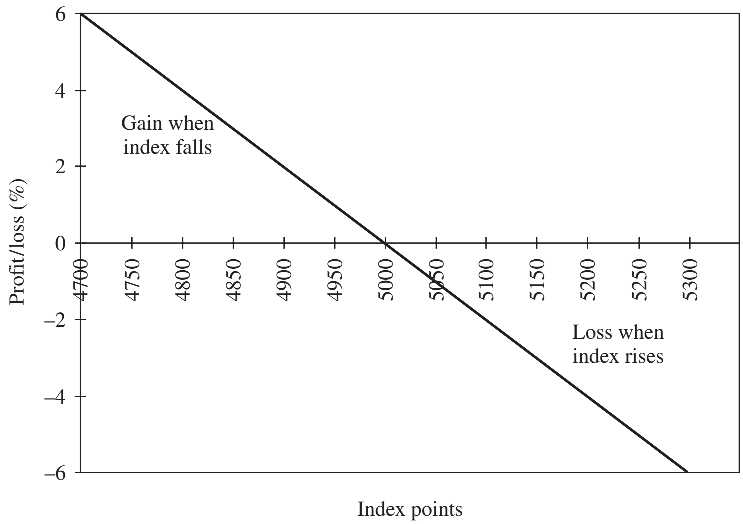

Figure 1.3: Short equity index position

1.1 Exchange-based futures contracts

Exchange-based futures contracts can be used to manipulate a portfolio's risk exposure. By adjusting the number of contracts shorted or bought, the portfolio can be made neutral to market moves by creating a zero portfolio beta. The portfolio can be made bullish by increasing the portfolio's beta to a value greater than one, or can be made bearish by reducing the portfolio's beta to a value of less than zero. Being exchange-based derivatives, futures contracts are very tightly defined and regulated to ensure that all parties to a transaction are aware of the instrument's operational characteristics. The use of futures contracts will be explained in Chapters 4, 5 and 8.

1.2 Forward contracts

Forward contracts could also be used to hedge the position in the same way as futures contracts.

These contracts, however, are off-exchange products (also known as OTCs). They offer far more in the way of flexibility from the end user's point of view. However, despite the fact that British Bankers' Association (BBA) and International Swaps and Derivatives Association (ISDA) standardised documentation exists for many OTC instruments, caution in their use on the part of the end user is still of paramount importance. For example, there may be penalty clauses if a contract is wound up before it reaches its maturity date. There may also be large bid–offer spreads when buying and later trying to close the position. And unless some form of monitoring takes place on a periodic basis there is a risk of counterparty default.

Notwithstanding these disadvantages, transactions in forwards in the world's interest rate, foreign exchange, equity, equity index markets run into billions of dollars daily.

In their favour, forward contracts do not suffer from basis risk, they are not restricted by the imposition of standardised contracts, boasts, in many cases 24-hour/day trading, and are not open to potential squeezes.

Pricing of these contracts is similar across all markets and considers two strategies. In the interest rate domain, in setting up a standardised forward rate agreement (FRA) a forward rate will need to be calculated and the results lead to forward/forward interest rates.

1.2.1 *Forward/forward interest rates*

The idea of rolling over an investment with its associated reinvestment risk or going for one long investment is one way of examining the problem of "fair" pricing.

The diagram below shows the rollover investment problem:

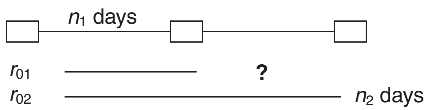

The question mark indicates that we do not know with certainty at what rate we will be able to reinvest when the first n_1 days have elapsed. Given that the spot rates r_{01} and r_{02} are known, an *implied forward rate* (r_{12}) can be calculated using equation (1.1).

$$r_{12} = \left[\left(\frac{1 + r_{02} \dfrac{n_2 \text{ days}}{\text{Days in year}}}{1 + r_{01} \dfrac{n_1 \text{ days}}{\text{Days in year}}} \right) - 1 \right] \left(\frac{\text{Days in year}}{n_2 - n_1} \right). \tag{1.1}$$

These implied rates play an important role in the valuation of *plain vanilla* as well as more complex swap structures.

Generally, the pricing of a forward contract can be regarded as the interplay between two strategies. In the case of a forward interest rate derivative, a fair forward price could be obtained by using the strategies in Table 1.1.

For there to be no profit accruing at some specified forward date to either of the counterparties, the two strategies should have equal values at the maturity of the contract. Thus when the contract is held for $n_1/365$ days, the value of Strategy 1 at maturity will be:

$$S_T - S_0 - S_0(r - q)(n_1/365) \tag{1.2}$$

where

- S_T represents the value of S_0 at maturity, T;
- r represents all the annualised costs of holding the asset (including the interest on the borrowed funds, storage, deterioration, insurance, etc.);
- q represents the annualised return on holding the asset.

Note that both r and q are expressed as proportions.

Strategy 1	**Strategy 2**
Borrow funds and buy one unit of the underlying security at price S_0.	Buy a forward contract at price F.

Table 1.1

The value of Strategy 2 at maturity will be the difference between the purchase price of the forward contract at initiation and the realised spot price at maturity as shown in equation (1.3):

$$S_T - F. \tag{1.3}$$

Since there is to be no profit, equation (1.2) must be equal to equation (1.3) at maturity.

$$\begin{aligned} S_T - F &= S_T - S_0 - S_0(r-q)(n_1/365) \\ F &= S_0 + S_0(r-q)(n_1/365). \end{aligned} \tag{1.4}$$

This same approach can be adopted for the pricing of a futures contract. Notice that if $r > q$ then F today will be higher than today's spot or cash market quote, and if $r < q$ the reverse is true. In either case this difference between the spot and forward, futures quote is called *basis* (sometimes raw basis). In the case of futures contracts the basis at some future point in time is, more often than not, unknown – when the contract is closed out before maturity, for example. This leads to the problem of basis risk mentioned earlier.

Interest rate forwards will be addressed in detail in Chapters 4 and 5 while fair pricing in an equity arena will be looked at in Chapter 8.

1.3 Options

Options too come in both exchange-based and OTC varieties. Basically there are two types of option classes available: calls and puts. Buying an option bestows on the purchaser (the holder) a privilege. That privilege is the right but not the obligation to buy (call) or sell (put) a specified underlying security at a specified price (strike or exercise price) on a specified future date (European-style options) or on or before a specified future date (American-style options).

Options bring with them the possibility of creating some interesting structures. The kinked – sometimes described as the "hockey stick" – shape of an option's payoff profile at expiration comes in four shapes: long or short calls, long or short puts as illustrated in Figures 1.4, 1.5, 1.6 and 1.7. How these profiles are obtained and to what uses they can be put is reserved for Chapter 6.

Figures 1.4–1.7 represent the kinked payoff profiles of long call, long put, short call and short put options, respectively.

Armed with the profiles and operational characteristics of the derivative instruments discussed in this text, it becomes possible to embark on the creation of derivative-based structured products in virtually any domain. Good examples of this are the introduction of energy derivatives, catastrophe derivatives, property index forwards (PIFs), and weather derivatives.

1.4 Swaps

These instruments are based on an agreement between two counterparties to exchange a series of cash flows. The cash flows are almost always calculated by reference to the behaviour of an index and are scaled by an agreed nominal principal. Swaps are now available on interest rates, currencies, equity, credit, property, weather indices and, of course, many types of commodities.

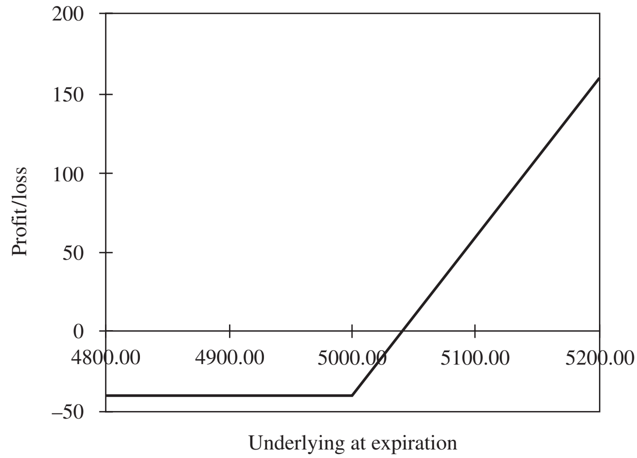

Figure 1.4: Options only position (long call)

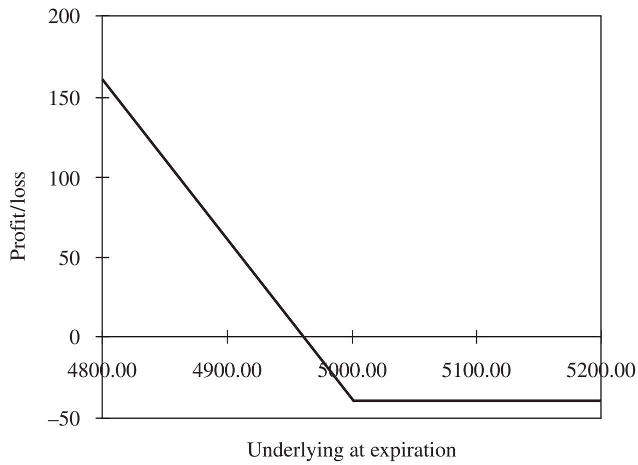

Figure 1.5: Options only position (long put)

Figure 1.8 gives an idea of the structure of a standard (often called plain vanilla) interest rate swap. The counterparties to the transaction have each raised funds in different markets. Company A has raised funds in the money market and will have an obligation to pay LIBOR on specified dates in the future. Company B has raised funds through the issue of a bond and will need to pay coupons on specified dates in the future. The two parties then agree to swap their periodic commitments. Company A pays Company B a fixed rate when due and receives in return LIBOR. In this way the original commitments will be altered from floating rate to fixed rate and from fixed rate to floating rate.

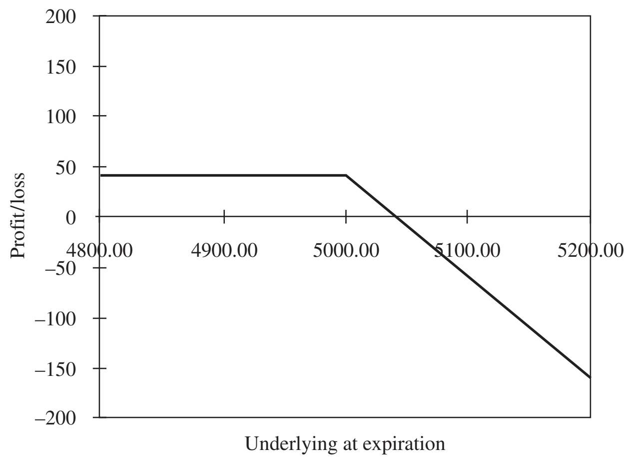

Figure 1.6: Options only position (short call)

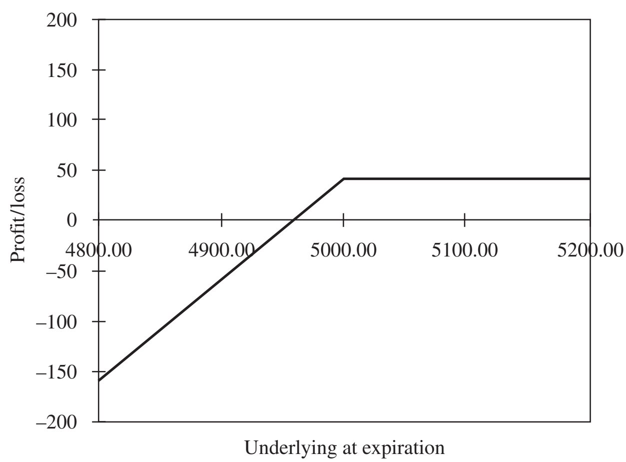

Figure 1.7: Options only position (short put)

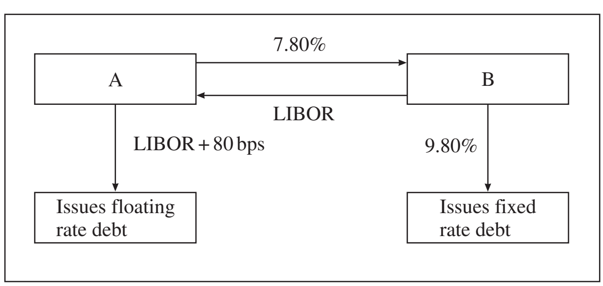

Figure 1.8: Structure of a standard interest rate swap

1.5 The future?

A few years ago a senior colleague of one of the authors at a highly respected City institution was heard to remark: “Derivative securities have reached their peak. There’s nowhere left to go. In the next few years we’ll see a decline in the use of derivatives”. At the time he was speaking, LIFFE was still operating an open-outcry market, the introduction of electronic trading was still several years in the future and the Euro was still at the design stage. The market though had already seen the introduction of principal guaranteed equity-linked bonds, reverse floaters, a variety of interest rate, currency, and equity swaps, and exotic options.

What happened? The markets are always on the lookout for new ideas, new instruments and structures, new ways of satisfying corporate funding needs, new ways of delivering investors enticing risk/reward investments. From one point of view the colleague was right. Some of the new instruments that came along were merely replays of existing instruments but with innovative twists. For example, embedding exotic options into other products rather than relying on expanding the exotic option market in its own right; or tweaking instruments to side step regulatory issues. However, two important aspects of financial market instruments were still waiting to receive attention, namely credit and liquidity risk. In retrospect it is easy to rationalise that this was a natural step waiting to be taken. Take for example the case of an equity-linked bond, breaking this instrument down into its risk components there are several risks that can be identified:

- market risk (equity);
- market risk (bond);
- currency risk;
- holding period risk;
- liquidity risk;
- default risk.

As discussed in Section 1.2, there are well-established exchange-based and OTC instruments available to manage the equity and bond and foreign exchange market risk components. Holding period risk can be identified as option features embedded in a structure that allow a counterparty to undertake a defined action: call, put or convert for example, and these can be taken into account.

Liquidity risk can refer to the ease with which the instrument can be turned back into cash, or the stream of cash flows associated with an instrument including the speed with which the investment is recouped. In the case of a government bond the markets are highly liquid and the stream of cash flows virtually assured. However, with some investments, although a stream of cash flows will provide a good source of revenue over many years, the instrument itself may not be readily turned back into cash. A corporate bond paying a fixed coupon but with an embedded option feature could fall into this category. However, by splitting the bond into its fixed coupon bond and option components the underlying security can be replicated, given that there is likely to be a deeper more liquid market for a bullet bond and the option the synthetic counterpart can be more readily traded.

Example 1.1

A company with an investment grade credit rating has issued a five-year callable bond paying a fixed income coupon. The bond is callable at the end of Year 3 if interest rates have fallen to a level specified in the bond's term sheet.

When interest rates are high the price of the bond is low and the bond is not called. This is illustrated in Figure 1.9(a). However, at a certain pre-specified level of interest (LIBOR) there will be a comparable bond yield that will provide the issuer with an opportunity to call the bond. When this happens the bond is redeemed and the investor will deposit the funds on an interest-bearing account that now reflects the now lower floating rate. Figure 1.9(b) illustrates the interest rate differential between the original fixed coupon and the lower floating rate.

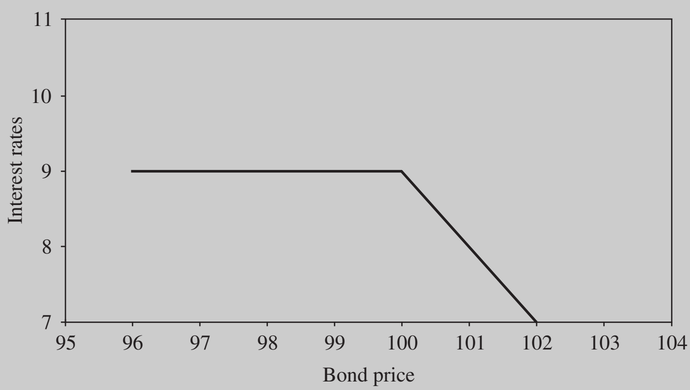

Figure 1.9(a): Callable bond

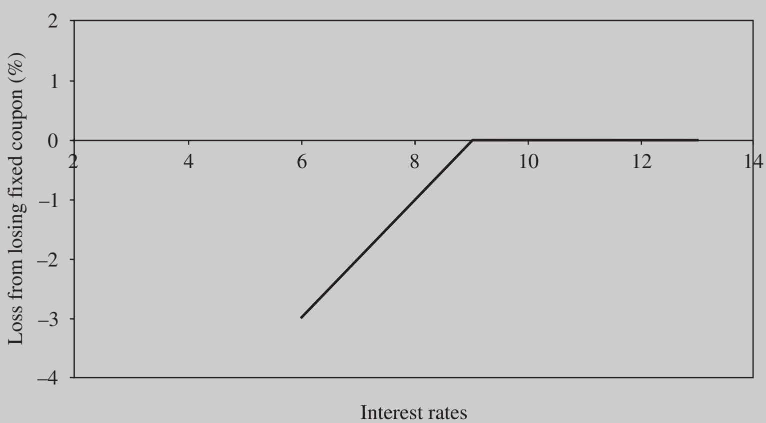

Figure 1.9(b): Bond callable

The profile of the callable bond presented in Figure 1.9(a) and 1.9(b) can be synthesised as follows. Buy a five-year bullet bond with the same coupon, maturity and credit rating as the callable bond. The profile of this instrument is depicted in Figure 1.10(a). The second step, carried out simultaneously, is to write a put option on a swap (a put swaption or receiver swaption) for the appropriate future period and strike price.

If the put swaption is exercised then the writer of the option will pay fixed and receive floating.

There are three scenarios that can be employed to analyse the outcome of the replication process:

1. Interest rates rise and the bond is not called. Replication strategy: the bullet bond pays coupons until maturity, in the absence of default. The put swaption will expire worthless since the holder of the receiver swaption would be contracting to pay higher floating rates and this would be irrational.
2. Interest rates remain unchanged at the point that the bond becomes callable and the bond is not called. Replication strategy: as in 1.

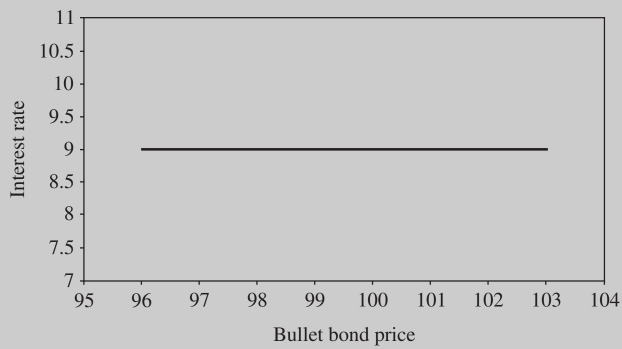

Figure 1.10(a): Bullet bond

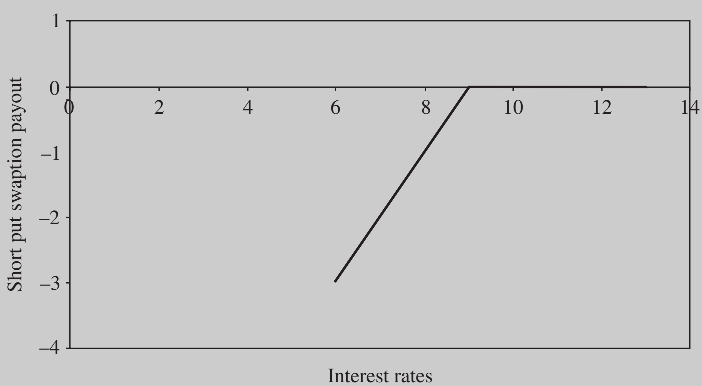

Figure 1.10(b): Swaption payout

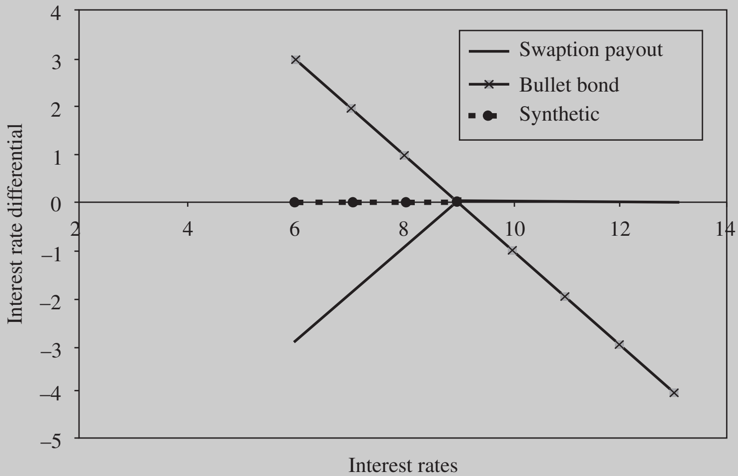

Figure 1.10(c): Synthetic

3. Interest rates fall and the bond is called. Replication strategy: the bullet bond continues to pay coupons, but the put swaption is exercised. The holder of the receiver swaption pays low floating rates and receives the contracted strike rate. This is funded by the bullet bond's fixed coupon.

A swap diagrammatic representation of this is shown in Figure 1.11.

Finally there is the counterparty risk aspect. The risk is that the counterparty to a transaction will fail to meet its obligation to pay coupons when due, to pay the redemption value or will fail to honour both of these obligations. If the instrument in question is a bond, the only way to cover that risk is to sell the bond but, of course, if the market recognises the risk of default the bid price is likely to be low. To this end credit derivatives have been introduced, and, despite initial problems concerning definitions for legal purposes, have gained acceptance in financial markets.

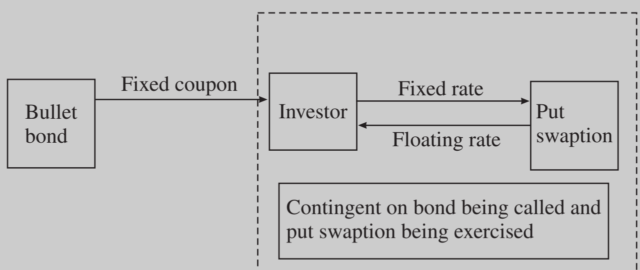

Figure 1.11: Synthetic callable bond

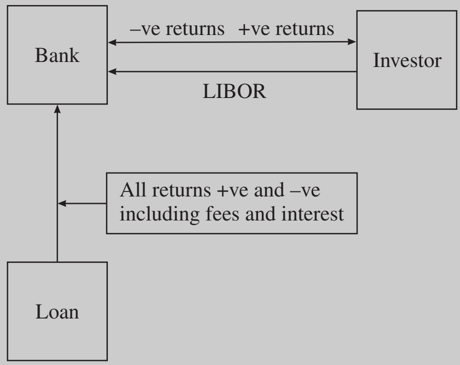

Figure 1.12: Total return swap (TRS)

There are several types of credit derivatives available. The most common are: credit default swaps and options, total return swaps, credit-linked notes, asset swaps and credit spread instruments. It is interesting to observe that a total return swap can also be used as a means of increasing an asset's liquidity. By way of example consider a financial institution with a loan portfolio that is regarded as being too large. For operational reasons, or pragmatic reasons, it is impossible to recall the loans and hence there is both credit and liquidity risk. In order to reduce the identified risks and increase liquidity the institution embarks on a total return swap for one of its large loans. The institution becomes the total return payer and receives a floating interest rate payment, LIBOR, say, in return. Figure 1.12 illustrates the structure of this transaction.

The bank will receive interest payments from the reference loan along with any fees or costs associated with the running of the loan and will pass these on to the investor counterparty so long as they are positive. In any period, if the return on the loan becomes negative, that negative return must be paid by the investor to the bank. In addition the bank will receive LIBOR on a period-by-period basis. Thus by entering this transaction the bank has succeeded in reducing its exposure to counterparty risk and has improved its liquidity position.

These opening pages provide an overview of the major derivative instruments. The chapters that follow will provide a detailed explanation of those derivative instruments and the way they operate in today's markets. Using illustrative examples, the users' derivatives can be put and approaches to valuation will also be explained.

2 Overview of Fixed Income Securities

In this chapter we present an introduction to fixed income analysis, particularly the pricing of default-free zero-coupon and coupon fixed-term bonds. Further reading is given in the bibliography.[1]

2.1 Basic concepts

We are familiar with two types of fixed income security, *zero-coupon bonds*, also known as *discount bonds* or *strips*, and *coupon bonds*. A zero-coupon bond makes a single payment on its maturity date, while a coupon bond makes regular interest payments at regular dates up to and including its maturity date. A coupon bond may be regarded as a set of strips, with each coupon payment and the redemption payment on maturity being equivalent to a zero-coupon bond maturing on that date. This is not a purely academic concept – witness events before the advent of the formal market in US Treasury strips, when a number of investment banks had traded the cash flows of Treasury securities as separate zero-coupon securities.[2]

Bonds are described by their issuer name, coupon rate and term to maturity, and those that have fixed coupons and fixed maturity terms are known as *conventional* or *vanilla* bonds. An example of the basic description of a bond is given in Figure 2.1, the Bloomberg “DES” page.

The literature we review in this section is set in a market of default-free bonds, whether they are zero-coupon bonds or coupon bonds. The market is assumed to be liquid so that bonds may be freely bought and sold. Prices of bonds are determined by the economy-wide supply and demand for the bonds at any time, so they are *macroeconomic* and not set by individual bond issuers or traders.

---

1 Parts of this chapter first appeared in Choudhry (2001), Chapter 3.

2 These banks included Merrill Lynch, Lehman Brothers and Salomon Brothers, among others (Fabozzi, 1993). The term “strips” comes from Separate Trading of Registered Interest and Principal of Securities, the name given when the official market was introduced by the Treasury. The banks would purchase Treasuries which would then be deposited in a safe custody account. Receipts were issued against each cash flow from each Treasury, and these receipts traded as individual zero-coupon securities. The market making banks earned profit due to the arbitrage difference in the price of the original coupon bond and the price at which the individual strips were sold. The US Treasury formalised trading in strips after 1985, after legislation had been introduced that altered the tax treatment of such instruments. The market in UK gilt strips trading began in December 1997. Strips are also traded in France, Germany and the Netherlands, among other countries.

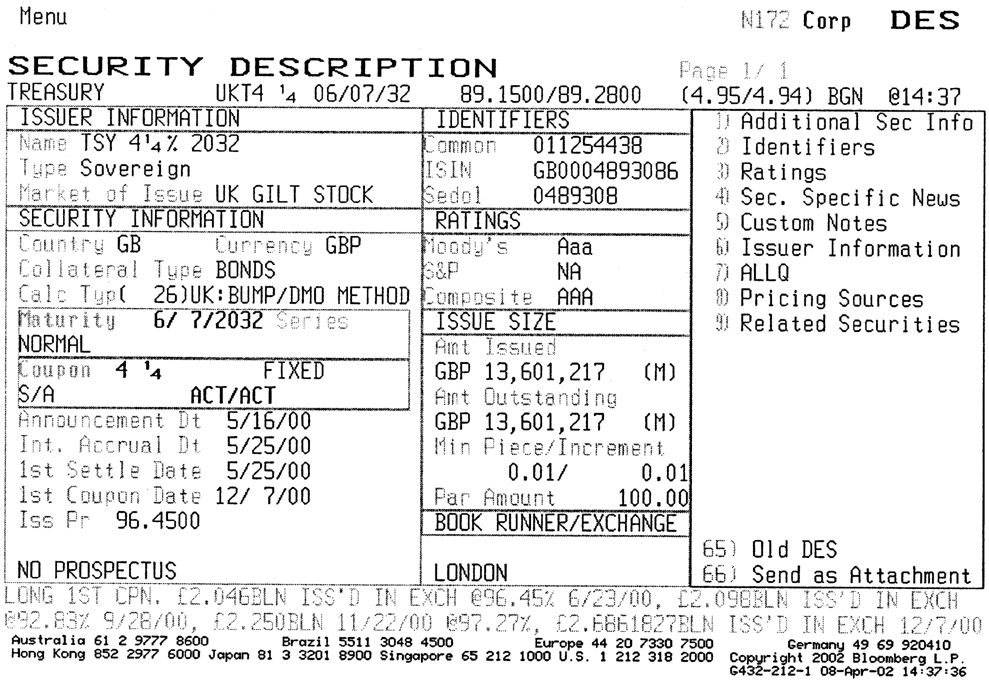
Menu N172 Corp DES
SECURITY DESCRIPTION Page 1/ 1
TREASURY UKT4 ¼ 06/07/32 89.1500/89.2800 (4.95/4.94) BGN @14:37
ISSUER INFORMATION
Name TSY 4¼% 2032
Type Sovereign
Market of Issue UK GILT STOCK
SECURITY INFORMATION
Country GB Currency GBP
Collateral Type BONDS
Calc Typ(26)UK:BUMP/DMO METHOD
Maturity 6/ 7/2032 Series
NORMAL
Coupon 4 ¼ FIXED
S/A ACT/ACT
Announcement Dt 5/16/00
Int. Accrual Dt 5/25/00
1st Settle Date 5/25/00
1st Coupon Date 12/ 7/00
Iss Pr 96.4500
NO PROSPECTUS
IDENTIFIERS
Common 011254438
ISIN GB0004893086
Sedol 0489308
RATINGS
Moody's Aaa
S&P NA
Composite AAA
ISSUE SIZE
Amt Issued
GBP 13,601,217 (M)
Amt Outstanding
GBP 13,601,217 (M)
Min Piece/Increment
0.01/ 0.01
Par Amount 100.00
BOOK RUNNER/EXCHANGE
LONDON
1) Additional Sec Info
2) Identifiers
3) Ratings
4) Sec. Specific News
5) Custom Notes
6) Issuer Information
7) ALLQ
8) Pricing Sources
9) Related Securities
65) Old DES
66) Send as Attachment
LONG 1ST CPN. £2.046BLN ISS'D IN EXCH @96.45% 6/23/00, £2.098BLN ISS'D IN EXCH @92.83% 9/28/00, £2.250BLN 11/22/00 @97.27%, £2.6861827BLN ISS'D IN EXCH 12/7/00
Australia 61 2 9777 8600 Brazil 5511 3048 4500 Europe 44 20 7330 7500 Germany 49 69 920410
Hong Kong 852 2977 6000 Japan 81 3 3201 8900 Singapore 65 212 1000 U.S. 1 212 318 2000 Copyright 2002 Bloomberg L.P.
G432-212-1 08-Apr-02 14:37:36

Figure 2.1: Bloomberg screen DES for United Kingdom gilt, 4¼% Treasury 2032, as at 12 November 2002 © Bloomberg LP. Used with permission.

2.1.1 *Zero-coupon bonds*

A zero-coupon bond is the simplest fixed income security. It is an issue of debt, the issuer promising to pay the face value of the debt to the bondholder on the date the bond matures. There are no coupon payments during the life of the bond, so it is a discount instrument, issued at a price that is below the face or *principal* amount. We denote as $P(t, T)$ the price of a discount bond at time t that matures at time T, with $T \geq t$. The term to maturity of the bond is denoted with n, where $n = T - t$. The price increases over time until the maturity date when it reaches the maturity or *par* value. If the par value of the bond is £1, then the *yield to maturity* of the bond at time t is denoted by $r(t, T)$, where r is actually "one plus the percentage yield" that is earned by holding the bond from t to T. We have

$$P(t,T) = \frac{1}{[r(t,T)]^n}. \tag{2.1}$$

The yield may be obtained from the bond price and is given by

$$r(t,T) = \left[\frac{1}{P(t,T)}\right]^{1/n} \tag{2.2}$$

which is sometimes written as

$$r(t,T) = P(t,T)^{-(1/n)}. \tag{2.3}$$

Analysts and researchers frequently work in terms of logarithms of yields and prices, or continuously compounded rates. One advantage of this is that it converts the non-linear relationship in equation (2.2) into a linear relationship.[3]

The bond price at time t_2 where $t \leq t_2 \leq T$ is given by

$$P(t_2, T) = P(t, T)e^{(t_2-t)r(t,T)} \tag{2.4a}$$

which is natural given that the bond price equation in continuous time is

$$P(t, T) = e^{-r(t,T)(T-t)} \tag{2.4b}$$

so that the yield is given by

$$r(t, T) = -\log\left(\frac{P(t, T)}{n}\right) \tag{2.5}$$

which is sometimes written as

$$\log r(t, T) = -\left(\frac{1}{n}\right)\log P(t, T). \tag{2.6}$$

The equation in (2.4) includes the exponential function, hence the use of the term continuously compounded.

The *term structure of interest rates* is the set of zero-coupon yields at time t for all bonds ranging in maturity from $(t, t+1)$ to $(t, t+m)$ where the bonds have maturities of $\{0, 1, 2, \ldots\ldots, m\}$. A good definition of the term structure of interest rates is given in Sundaresan (1997), who states that it

> "... refers to the relationship between the yield to maturity of default-free zero coupon securities and their maturities."
>
> (page 176)

The *yield curve* is a plot of the set of yields for $r(t, t+1)$ to $r(t, t+m)$ against m at time t. For example, Figures 2.2–2.4 show the log zero-coupon yield curve for US Treasury strips, UK gilt strips and French OAT strips on 27 September 2000. Each of the curves exhibit peculiarities in their shape, although the most common type of curve is gently upward sloping, as is the French curve. The UK curve is *inverted*. We explore further the shape of the yield curve later in this chapter.

---

[3] A linear relationship in X would be a function $Y = f(X)$ in which the X values change via a power or index of 1 only and are not multiplied or divided by another variable or variables. So for example terms such as $X^2, \sqrt{X}$ and other similar functions are not linear in X, nor are terms such as XZ or X/Z where Z is another variable. In econometric analysis, if the value of Y is solely dependent on the value of X, then its rate of change with respect to X, or the derivative of Y with respect to X, denoted dY/dX, is independent of X. Therefore if $Y = 5X$, then $dY/dX = 5$, which is independent of the value of X. However if $Y = 5X^2$, then $dY/dX = 10X$, which is not independent of the value of X. Hence this function is not linear in X. The classic regression function $E(Y|X_i) = \alpha + \beta X_i$ is a linear function with slope β and intercept α and the regression "curve" is represented geometrically by a straight line.

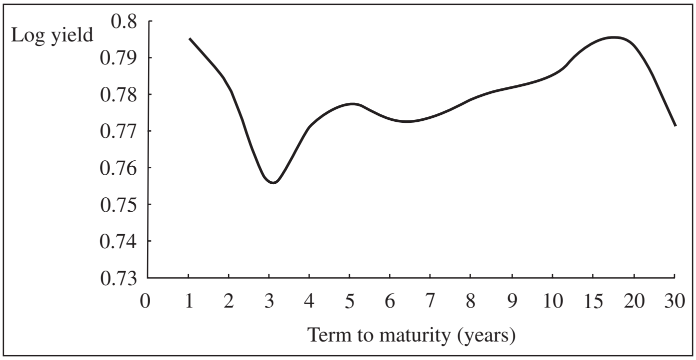

Figure 2.2: US Treasury zero-coupon yield curve in September 2000. (Source: Bloomberg)

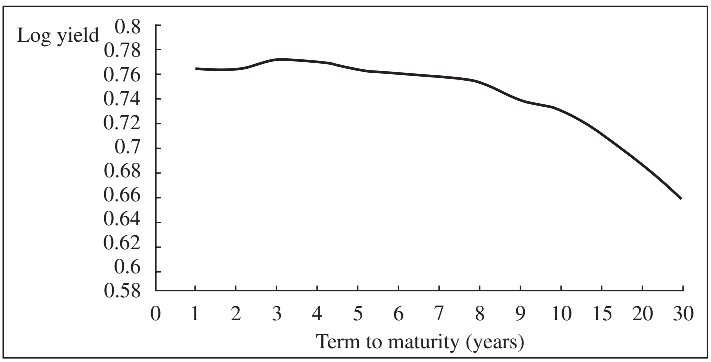

Figure 2.3: UK gilt zero-coupon yield curve. (Source: Bloomberg)

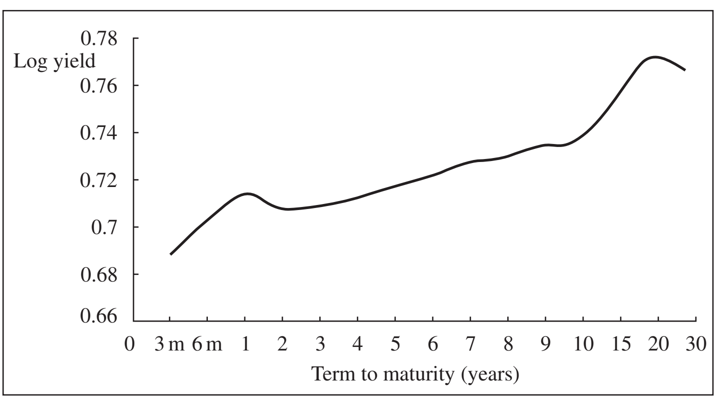

Figure 2.4: French OAT zero-coupon yield curve. (Yield source: Bloomberg)

2.1.2 *Coupon bonds*

The majority of bonds in the market make periodic interest or *coupon* payments during their life, and are known as coupon bonds. We have already noted that such bonds may be viewed as a package of individual zero-coupon bonds. The coupons have a nominal value that is a percentage of the nominal value of the bond itself, with steadily longer maturity dates, while the final redemption payment has the nominal value of the bond itself and is redeemed on the maturity date. We denote a bond issued at time i and maturing at time T as having a w-element vector of payment dates $(t_1, t_2, \ldots.., t_{w-1}, T)$ and matching date payments $C_1, C_2, \ldots.., C_{w-1}, C_w$. In the academic literature these coupon payments are assumed to be made in continuous time, so that the stream of coupon payments is given by a positive function of time $C(t)$, $i < t \leq T$. An investor that purchases a bond at time t that matures at time T pays $P(t, T)$ and will receive the coupon payments as long as she continues to hold the bond.[4]

The yield to maturity at time t of a bond that matures at T is the interest rate that relates the price of the bond to the future returns on the bond, that is, the rate that *discounts* the bond's cash flow stream C_w to its price $P(t, T)$. This is given by

$$P(t, T) = \sum_{t_i > t} C_i e^{-(t_i - t) r(t,T)} \tag{2.7}$$

which says that the bond price is given by the present value of the cash flow stream of the bond, discounted at the rate $r(t, T)$. For a zero-coupon (2.7) reduces to (2.5). In the academic literature where coupon payments are assumed to be made in continuous time, the $\sum$ summation in (2.7) is replaced by the $\int$ integral. We will look at this in a moment.

In some texts the plot of the yield to maturity at time t for the term of the bonds m is described as the term structure of interest rates but it is generally accepted that the term structure is the plot of zero-coupon rates only. Plotting yields to maturity is generally described as graphically depicting the yield curve, rather than the term structure. Of course, given the law of one price, there is a relationship between the yield to maturity yield curve and the zero-coupon term structure, and given the first one can derive the second.

The equation in (2.7) obtains the continuously compounded yield to maturity $r(t, T)$. It is the use of the exponential function that enables us to describe the yield as continuously compounded.

The market frequently uses the measure known as *current yield* which is

$$rc = \frac{C}{P_d} \times 100 \tag{2.8}$$

where P_d is the dirty price of the bond. The measure is also known as the *running yield* or *flat yield*. Current yield is not used to indicate the interest rate or discount rate and therefore should not be mistaken for the yield to maturity.

[4] In theoretical treatment this is the discounted clean price of the bond. For coupon bonds in practice, unless the bond is purchased for value on a coupon date, it will be traded with interest accrued. The interest that has accrued on a pro-rata basis from the last coupon date is added to the clean price of the bond, to give the market "dirty" price that is actually paid by the purchaser.

2.2 Bond price in continuous time[5]

2.2.1 *Fundamental concepts*

In this section we present an introduction to the bond price equation in continuous time. The necessary background on price processes is covered in a number of texts, including Chapter 4 of Choudhry (2001).

Consider a trading environment where bond prices evolve in a w-dimensional process

$$X(t) = [X_1(t),\ X_2(t),\ X_3(t),\ \ldots..,\ X_w(t)], \quad t > 0 \tag{2.9}$$

where the random variables are termed *state variables* that reflect the state of the economy at any point in time. The markets assume that the state variables evolve through a process described as geometric Brownian motion or a Weiner process. It is therefore possible to model the evolution of these variables, in the form of a stochastic differential equation.

The market assumes that the cash flow stream of assets such as bonds and (for equities) dividends is a function of the state variables. A bond is characterised by its coupon process

$$C(t) = \tilde{C}[X_1(t),\ X_2(t),\ X_3(t),\ \ldots..,\ X_w(t), t]. \tag{2.10}$$

The coupon process represents the cash flow that the investor receives during the time that she holds the bond. Over a small incremental increase in time of dt from the time t the investor can purchase $1 + C(t)dt$ units of the bond at the end of the period $t + dt$. Assume that there is a very short-term discount security such as a Treasury bill that matures at $t + dt$, and during this period the investor receives a return of $r(t)$. This rate is the annualised short-term interest rate or *short rate*, which in the mathematical analysis is defined as the rate of interest charged on a loan that is taken out at time t and which matures almost immediately. For this reason the rate is also known as the *instantaneous rate*. The short rate is given by

$$r(t) = r(t, t) \tag{2.11}$$

and

$$r(t) = -\frac{\partial}{\partial T} \log P(t, t). \tag{2.12}$$

5 Readers who wish to explore the concepts discussed in this section to further depth may wish to consider references such as Avellaneda and Laurence (2000), Baxter and Rennie (1996), Neftci (2000), Cambell *et al.* (1997), Ross (1999), and Shiller (1990). These are all excellent texts of very high quality, and strongly recommended. For an accessible and highly readable introduction Ross's book is worth buying for Chapter 4 alone, as is Avellaneda and Laurence's for their Chapter 12. For a general introduction to the main pricing concepts see Campbell *et al.* (1997), Chapter 10. Chapter 3 in Jarrow (1996) is an accessible introduction for discrete-time bond pricing. Sundaresan (1997) is an excellent overview text on the fixed income market as a whole, and is highly recommended. Further recommended references are given in the bibliography at the end of this chapter.

If we continuously reinvest the short-term security such as the T-bill at this short rate, we obtain a cumulative amount that is the original investment multiplied by (2.13).[6]

$$M(t) = \exp\left[\int_0^t r(s)ds\right] \tag{2.13}$$

where M is a money market account that offers a return of the short rate $r(t)$.

If we say that the short rate is constant, making $r(t) = r$, then the price of a risk-free bond that pays £1 on maturity at time T is given by

$$P(t, T) = e^{-r(T-t)}. \tag{2.13a}$$

What equation (2.13a) states is that the bond price is simply a function of the continuously compounded interest rate, with the right-hand side of (2.13a) being the discount factor at time t. At $t = T$ the discount factor will be 1, which is the redemption value of the bond and hence the price of the bond at this time.

Consider the following scenario; a market participant may undertake the following:

- it can invest $e^{-r(T-t)}$ units cash in a money market account today, which will have grown to a sum of £1 at time T;
- it can purchase the risk-free zero-coupon bond today, which has a maturity value of £1 at time T.

The market participant can invest in either instrument, both of which we know beforehand to be risk-free, and both of which have identical payouts at time T and have no cash flow between now and time T. As interest rates are constant, a bond that paid out £1 at T must have the same value as the initial investment in the money market account, which is $e_t^{-r(T-t)}$. Therefore equation (2.13a) must apply. This is a restriction placed on the zero-coupon bond price by the requirement for markets to be arbitrage-free.

If the bond was not priced at this level, arbitrage opportunities would present themselves. Consider if the bond was priced higher than $e_t^{-r(T-t)}$. In this case, an investor could sell short the bond and invest the sale proceeds in the money market account. On maturity at time T, the short position will have a value of −£1 (negative, because the investor is short the bond) while the money market will have accumulated £1, which the investor can use to pay the proceeds on the zero-coupon bond. However the investor will have surplus funds because at time t,

$$P(t, T) - e^{-r(T-t)} > 0$$

and so will have profited from the transaction at no risk to himself.

The same applies if the bond is priced below $e_t^{-r(T-t)}$. In this case the investor borrows $e_t^{-r(T-t)}$ and buys the bond at its price $P(t, T)$. On maturity the bond pays £1 which is used to repay the loan amount, however the investor will gain because

$$e^{-r(T-t)} - P(t, T) > 0.$$

---

[6] This equation uses the integral operator. The integral is the tool used in mathematics to calculate sums of an infinite number of objects, that is where the objects are uncountable. This is different to the $\sum$ operator which is used for a countable number of objects. For a readable and accessible review of the integral and its use in quantitative finance, see Neftci (2000), pp. 59–66. We provide a brief introduction in Appendix 2.1.

Therefore the only price at which no arbitrage profit can be made is if

$$P(t,T) = e^{-r(T-t)}. \tag{2.13b}$$

In the academic literature the price of a zero-coupon bond is given in terms of the evolution of the short-term interest rate, in what is termed the *risk-neutral measure*.[7] The short rate $r(t)$ is the interest rate earned on a money market account or short-dated risk-free security such as the T-bill suggested above, and it is assumed to be continuously compounded. This makes the mathematical treatment simpler. With a zero-coupon bond we assume a payment on maturity of 1 (say \$1 or £1), a one-off cash flow payable on maturity at time T. The value of the zero-coupon bond at time t is therefore given by

$$P(t,T) = \exp\left(-\int_t^T r(s)ds\right) \tag{2.14}$$

which is the redemption value of 1 divided by the value of the money market account, given by equation (2.13).

The bond price for a coupon bond is given in terms of its yield as

$$P(t,T) = \exp(-(T-t)r(T-t)). \tag{2.15}$$

Equation (2.14) is very commonly encountered in the academic literature. Its derivation is not so frequently occurring however; we present it in Appendix 2.2, which is a summary of the description given in Ross (1999). This reference is highly recommended reading. It is also worth referring to Neftci (2000), Chapter 18.

Equation (2.14) represents the zero-coupon bond pricing formula when the spot rate is continuous or *stochastic*, rather than constant. The rate $r(s)$ is the risk-free return earned during the very short or *infinitesimal* time interval $(t, t+dt)$. The rate is used in the equations for the value of a money market account (2.13) and the price of a risk-free zero-coupon bond (2.15).

2.2.2 *Stochastic rates in continuous time*

In the academic literature the bond price given by (2.15) evolves as a *martingale* process under the risk-neutral probability measure $\tilde{P}$. This is an advanced branch of fixed income mathematics, and is outside the scope of this book.[8] However under this analysis the bond price is given as

$$P(t,T) = E_t^{\tilde{P}}\left[e^{-\int_t^T r(s)ds}\right] \tag{2.16}$$

where the right-hand side of equation (2.16) is viewed as the randomly evolved *discount factor* used to obtain the present value of the £1 maturity amount. Equation (2.16) also states that bond prices are dependent on the entire spectrum of short-term interest rates

7 This is part of the *arbitrage pricing theory*. For detail on this see Cox *et al.* (1985), while Duffie (1992) is a fuller treatment for those with a strong grounding in mathematics.

8 Interested readers should consult Neftci (2000), Chapters 2, 17 and 18; another accessible text is Baxter and Rennie (1996) while Duffie (1992) is a leading-edge reference for those with a strong background in mathematics.

$r(s)$ in the future during the period $t < s < T$. This also implies that the term structure at time t contains all the information available on short rates in the future.[9]

From equation (2.16) we say that the function $T \to P_t^T$, $t < T$ is the discount curve (or *discount function*) at time t. Avellaneda and Laurence (2000) note that the markets usually replace the term $(T - t)$ with a term meaning *time to maturity*, so the function becomes

$$\tau \to P_t^{t+\tau}, \tau > 0, \text{ where } \tau = (T - t).$$

Under a constant spot rate, the zero-coupon bond price is given by

$$P(t, T) = e^{-r(t,T)(T-t)}. \tag{2.17}$$

From equations (2.16) and (2.17) we can derive a relationship between the yield $r(t, T)$ of the zero-coupon bond and the short rate $r(t)$, if we equate the two right-hand sides, namely

$$e^{-r(T,t)(T-t)} = E_t^{\tilde{P}}\left[e^{-\int_t^T r(s)ds}\right]. \tag{2.18}$$

Taking the logarithm of both sides we obtain

$$r(t, T) = \frac{-\log E_t^{\tilde{P}}\left[e^{-\int_t^T r(s)ds}\right]}{T - t}. \tag{2.19}$$

This describes the yield on a bond as the average of the spot rates that apply during the life of the bond, and under a constant spot rate the yield is equal to the spot rate.

With a zero-coupon bond and assuming that interest rates are positive, $P(t, T)$ is less than or equal to 1. The yield of the bond is, as we have noted, the continuously compounded interest rate that equates the bond price to the discounted present value of the bond at time t. This is given by

$$r(t, T) = \frac{-\log(P(t, T))}{T - t} \tag{2.20}$$

so we obtain

$$P(t, T) = e^{-(T-t)r(T-t)}. \tag{2.21}$$

In practice this means that an investor will earn $r(t, T)$ if she purchases the bond at t and holds it to maturity.

2.2.3 *Coupon bonds*

Using the same principles as in the previous section, we can derive an equation for the price of a coupon bond in the same terms of a risk-neutral probability measure of the evolution of interest rates. Under this analysis, the bond price is given by

$$P_c = 100 \times E_t^{\tilde{P}}\left(e^{-\int_t^{tN} r(s)ds}\right) + \sum_{n:t_n>t}^{N} \frac{C}{w} E_t^{\tilde{P}}\left(e^{-\int_t^{tn} r(s)ds}\right) \tag{2.22}$$

[9] This is related to the view of the short rate evolving as a martingale process. For a derivation of equation (2.16) see Neftci (2000), page 417.

where

P_c	is the price of a coupon bond;
C	is the bond coupon;
t_n	is the coupon date, with $n \leq N$, and $t = 0$ at the time of valuation;
w	is the coupon frequency.[10]

Here 100 is used as the convention for *principal* or bond nominal value (that is, prices are quoted per cent, or per 100 nominal).

Equation (2.22) is written in some texts as

$$P_c = 100e^{-rN} + \int_n^N Ce^{-rn} dt. \tag{2.23}$$

We can simplify equation (2.22) by substituting *Df* to denote the discount factor part of the equation and assuming an annual coupon, which gives us

$$P = 100 \times Df_N + \sum_{n:t_n \geq t}^{N} C \times Df_n, \tag{2.24}$$

which states that the market value of a risk-free bond on any date is determined by the discount function on that date.

We know from earlier in this chapter that the actual price paid in the market for a bond includes accrued interest from the last coupon date, so that the price given by equation (2.24) is known as the *clean price*, and the traded price, which includes accrued interest, is known as the *dirty price*.

2.3 Forward rates

An investor can combine positions in bonds of differing maturities to guarantee a rate of return that begins at a point in the future. That is, the trade ticket would be written at time t but would cover the period T to $T+1$ where $t < T$ (sometimes written as beginning at T_1 and ending at T_2, with $t < T_1 < T_2$). The interest rate earned during this period is known as the *forward rate*.[11] The mechanism by which this forward rate can be guaranteed is described in the box, following Jarrow (1996) and Campbell *et al.* (1997).

---

[10] Conventional or *plain vanilla* bonds pay coupon on an annual or semi-annual basis. Other bonds, notably certain floating-rate notes and mortgage- and other asset-backed securities also pay coupon on a monthly basis, depending on the structuring of the transaction.

[11] See the footnote on page 639 of Shiller (1990) for a fascinating insight on the origin of the term "forward rate", which Mr Shiller ascribes to John Hicks in his book *Value and Capital* (2nd edition, Oxford University Press 1946).

The forward rate

An investor buys at time t1 unit of a zero-coupon bond maturing at time T, priced at $P(t, T)$ and simultaneously sells $P(t, T)/P(t, T+1)$ bonds that mature at $T+1$. From Table 2.1 we see that the net result of these transactions is a zero cash flow. At time T there is a cash inflow of 1, and then at time $T+1$ there is a cash outflow of $P(t, T)/P(t, T+1)$. These cash flows are identical to a loan of funds made during the period T to $T+1$, contracted at time t. The interest rate on this loan is given by $P(t, T)/P(t, T+1)$, which is therefore the forward rate. That is,

$$f(t, T) = \frac{P(t, T)}{P(t, T+1)}. \tag{2.25}$$

Together with our earlier relationships on bond price and yield, from equation (2.25) we can define the forward rate in terms of yield, with the return earned during the period $(T, T+1)$ being

$$f(t, T, T+1) = \frac{1}{(P(t, T+1)/P(t, T))} = \frac{(r(t, T+1))^{(T+1)}}{r(t, T)^T}. \tag{2.26}$$

From equation (2.26) we can obtain a bond price equation in terms of the forward rates that hold from t to T,

$$P(t, T) = \frac{1}{\prod_{k=t}^{T-1} f(t, k)}. \tag{2.27}$$

A derivation of this equation can be found in Jarrow (1996), Chapter 3. Equation (2.27) states that the price of a zero-coupon bond is equal to the nominal value, here assumed to be 1, receivable at time T after it has been discounted at the set of forward rates that apply from t to T.[12]

Transactions		**Time**	
Buy 1 unit of T-period bond	$t - P(t, T)$	$T+1$	$T+1$
Sell $P(t, T)/P(t, T+1)$ $T+1$-period bonds	$+[P(t, T)/P(t, T+1)]P(t, T+1)$		$-P(t, T)/P(t, T+1)$
Net cash flows	0	+1	$-P(t, T)/P(t, T+1)$

Table 2.1

[12] The symbol $\prod$ means "take the product of", and is defined as $\prod_{i=1}^{n} x_i = x_1 \cdot x_2 \cdot \cdots \cdot x_n$, so that $\prod_{k=t}^{T-1} f(t, k) = f(t, t) \cdot f(t, t+1) \cdot \cdots \cdot f(t, T-1)$ which is the result of multiplying the rates that obtain when the index k runs from t to $T-1$.

When calculating a forward rate, it is as if we are writing an interest rate today that is applicable at the forward start date, in other words we trade a forward contract. The law of one price, or no-arbitrage, is used to calculate the rate. For a loan that begins at T and matures at $T+1$, similarly to the way we described in the box above, consider a purchase of a $T+1$-period bond and a sale of p amount of the T-period bond. The net cash position at t must be zero, so p is given by

$$p = \frac{P(t, T+1)}{P(t, T)} \tag{2.28}$$

and to avoid arbitrage the value of p must be the price of the $T+1$-period bond at time T. Therefore the forward yield is given by

$$f(t, T+1) = -\frac{\log P(t, T+1) - \log P(t, T)}{(T+1) - T}. \tag{2.29}$$

If the period between T and the maturity of the later-dated bond is reduced, so we now have bonds that mature at T and T_2, and $T_2 = T + \Delta t$, then as the incremental change in time Δt becomes progressively smaller we obtain an instantaneous forward rate, which is given by

$$f(t, T) = -\frac{\partial}{\partial T} \log P(t, T). \tag{2.30}$$

This rate is defined as the forward rate and is the price today of forward borrowing at time T. The forward rate for borrowing today where $T = t$ is equal to the instantaneous short rate $r(t)$. At time t the spot and forward rates for the period (t, t) will be identical, at other maturity terms they will differ.

For all points other than at (t, t) the forward rate yield curve will lie above the spot rate curve if the spot curve is positively sloping. The opposite applies if the spot rate curve is downward sloping. Campbell *et al.* (1997, pages 400–401) observe that this property is a standard one for marginal and average cost curves. That is, when the cost of a marginal unit (say, of production) is above that of an average unit, then the average cost will increase with the addition of a marginal unit. This results in the average cost rising when the marginal cost is above the average cost. Equally the average cost per unit will decrease when the marginal cost lies below the average cost.

Example 2.1: The spot and forward yield curves

From the discussion in this section we see that it is possible to calculate bond prices, spot and forward rates provided that one has a set of only one of these parameters. Therefore given the following set of zero-coupon rates, observed in the market, given in Table 2.2, we calculate the corresponding forward rates and zero-coupon bond prices as shown. The initial term structure is upward sloping. The two curves are illustrated in Figure 2.5.

There are technical reasons why the theoretical forward rate has a severe kink at the later maturity, but we shall not go into an explanation of this as it is outside the scope of this book.

Term to maturity (0, T)	Spot rate $r(0, T)^*$	Forward rate $f(0, T)^*$	Bond price $P(0, T)$
0			1
1	1.054	1.054	0.94877
2	1.055	1.056	0.89845
3	1.0563	1.059	0.8484
4	1.0582	1.064	0.79737
5	1.0602	1.068	0.7466
6	1.0628	1.076	0.69386
7	1.06553	1.082	0.64128
8	1.06856	1.0901	0.58833
9	1.07168	1.0972	0.53631
10	1.07526	1.1001	0.48403
11	1.07929	1.1205	0.43198

Table 2.2: Hypothetical zero-coupon yield and forward rates

* Interest rates are given as $(1+r)$

1y	5.4	5.4
2y	5.5	5.6
3y	5.63	5.9
4y	5.82	6.4
5y	6.02	6.8
6y	6.28	7.6
7y	6.553	8.2
8y	6.856	9.01
9y	7.168	9.72
10y	7.526	10.01
11y	7.4	6.8

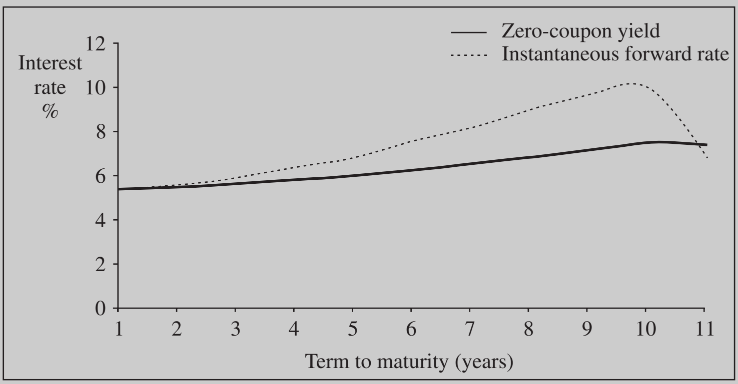

Figure 2.5: Hypothetical zero-coupon and forward yield curves

Appendix 2.1: The integral

The approach used to define integrals begins with an approximation involving a countable number of objects, which is then gradually transformed into an uncountable number of objects. A common form of integral is the Riemann integral.

Given a calculable or *deterministic* function that has been graphed for a period of time, let us say we require the area represented by this graph. The function is $f(t)$ and it is graphed over the period $[0, T]$. The area of the graph is given by the integral

$$\int_0^T f(s)ds \tag{A2.1.1}$$

which is the area represented by the graph. This can be calculated using the Riemann integral, for which the area represented is shown in Figure 2.6.

The general definition

To make the calculation the time interval is separated into n intervals, given by

$$t_0 = 0 < t_1 < t_2 < \ldots.. < t_{n-1} < t_n < T. \tag{A2.1.2}$$

The approximate area under the graph is given by the sum of the area of each of the rectangles, for which we assume each segment outside the graph is compensated by the area under the line that is not captured by any of the rectangles. Therefore we can say that an approximating measure is described by

$$\sum_{i=0}^{n} f\left(\frac{t_i + t_{i-1}}{2}\right)(t_i - t_{i-1}). \tag{A2.1.3}$$

This states that the area under the graph can be approximated by taking the sum of the n rectangles, which are created from the base x-axis which begins from t_0 through to T_n and the y-axis as height, described as

$$f((t_i + t_{i-1})/2).$$

This approximation only works if a sufficiently small base has been used for each interval, and if the function $f(t)$ is a smooth function, that is it does not experience sudden swings or kinks.

The definition of the Riemann integral is, given that

$$\max_i |t_i - t_{i-l}| \to 0,$$

defined by the limit

$$\sum_{i=0}^{n} f\left(\frac{t_i + t_{i-1}}{2}\right)(t_i - t_{i-l}) \to \int_0^T f(s)ds. \tag{A2.1.4}$$

If the approximation is not sufficiently accurate we can adjust it by making the intervals smaller still, which will make the approximation closer. This approach cannot be used if the

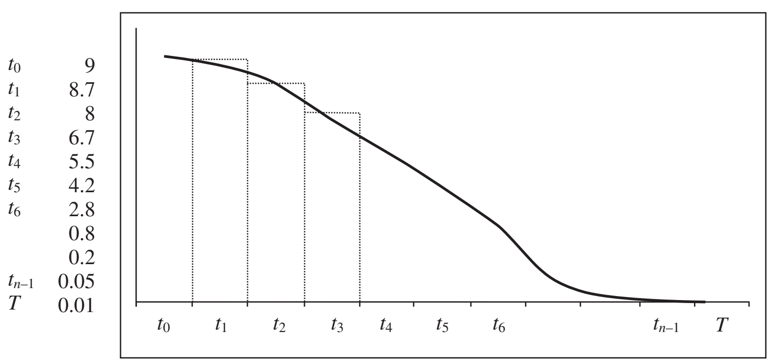

Figure 2.6

function is not smooth; in mathematics this requirement is stated by saying that the function must be integrable or Riemann integrable.

Other integral forms are also used, a good introduction to these is given in Neftci (2000), Chapter 4.

Appendix 2.2: The derivation of the bond price equation in continuous time

This section summarises the section on the bond price equation described in Ross (1999), on pages 54–56. This is an excellent reference, very readable and accessible, and is highly recommended. We replace Ross's use of investment at time 0 for maturity at time t with the terms t and T respectively, which is consistent with the price equations given in the main text. We also use the symbol M for the maturity value of the money market account, again to maintain consistency with the equations used in Chapters 2–4 of this book.

Assume a continuously compounded interest rate $r(s)$ that is payable on a money market account at time s. This is the instantaneous interest rate at time s. Assume further that an investor places x in this account at time s; after a very short time period of time h, the account would contain

$$x_h \approx x(1 + r(s)h). \tag{A2.2.1}$$

Say that $M(T)$ is the amount that will be in the account at time T if an investor deposits £1 at time t. To calculate $M(T)$ in terms of the spot rate $r(s)$, where $t \leq s \leq T$, for an incremental change in time of h, we have

$$M(s + h) \approx M(s)(1 + r(s)h) \tag{A2.2.2}$$

which leads us to

$$M(s + h) - M(s) \approx M(s)r(s)h \tag{A2.2.3}$$

and

$$\frac{M(s + h) - M(s)}{h} \approx M(s)r(s). \tag{A2.2.4}$$

The approximation given by equation (A2.2.4) turns into an equality as the time represented by h becomes progressively smaller. At the limit given as h approaches zero, we say

$$M'(s) = M(s)r(s) \tag{A2.2.5}$$

which can be rearranged to give

$$\frac{M'(s)}{M(s)} = r(s). \tag{A2.2.6}$$

From equation (A2.2.6) we imply that in a continuous time process

$$\int_t^T \frac{M'(s)}{M(s)} ds = \int_t^T r(s) ds \tag{A2.2.7}$$

and that

$$\log(M(T)) - \log(M(t)) = \int_t^T r(s)ds. \tag{A2.2.8}$$

However we deposited £1 at time t, that is $M(t) = 1$, so from equation (A2.2.8) we obtain the value of the money market account at T to be

$$M(T) = \exp\left(\int_t^T r(s)ds\right) \tag{A2.2.9}$$

which was our basic equation shown as equation (2.13).

Let us now introduce a risk-free zero-coupon bond that has a maturity value of £1 when it is redeemed at time T. If the spot rate is constant, then the price at time t of this bond is given by

$$P(t, T) = e^{-r(T-t)} \tag{A2.2.10}$$

where r is the continuously compounded instantaneous interest rate. The right-hand side of equation (A2.2.10) is the equation for the present value of £1, payable at time T, discounted at time t at the continuously compounded, constant interest rate r.

So we say that $P(t, T)$ is the present value at time t of £1 to be received at time T. Since a deposit of $1/M(T)$ at time t will have a value of 1 at time T, we are able to say that

$$P(t, T) = \frac{1}{M(T)} = \exp\left(-\int_t^T r(s)ds\right), \tag{A2.2.11}$$

which is our bond price equation in continuous time.

If we say that the *average* of the spot interest rates from t to T is denoted by $rf(T)$, so we have

$$rf(T) = \frac{1}{T}\int_t^T r(s)ds, \tag{A2.2.12}$$

then the function $rf(T)$ is the term structure of interest rates.

Selected bibliography and references

Adams, K., Van Deventer, D., "Fitting Yield Curves and Forward Rate Curves with Maximum Smoothness", *Journal of Fixed Income* 4, 1994, pp. 52–62

Avellaneda, M., Laurence, P., *Quantitative Modelling of Derivative Securities*, Chapman and Hall/CRC 2000, Chapters 10–12

Baxter, M., Rennie, A., *Financial Calculus*, Cambridge University Press 1996, Chapter 5

Campbell, J., Lo, A., MacKinlay, A., *The Econometrics of Financial Markets*, Princeton University Press 1997, Chapters 10–11

Choudhry, M., *Bond Market Securities*, FT Prentice Hall 2001

Cox, J., Ingersoll, J., Ross, S., "An Inter-Temporal General Equilibrium Model of Asset Prices", *Econometrica* 53, 1985

Deacon, M., Derry, A., "Estimating the Term Structure of Interest Rates", *Bank of England Working Paper Series* No 24, July 1994

Duffie, D., *Dynamic Asset Pricing Theory*, Princeton University Press 1992

Fabozzi, F., *Bond Markets, Analysis and Strategies*, 2nd edition, Prentice Hall 1993, Chapter 5

Fabozzi, F., *Fixed Income Mathematics*, McGraw-Hill 1997

Gujarati, N. D., *Basic Econometrics*, 3rd edition, McGraw-Hill 1995
Hicks, J., *Value and Capital*, 2nd edition, Oxford University Press 1946
Ingersoll, J. Jr., *Theory of Financial Decision Making*, Rowman & Littlefield 1987, Chapter 18
James, J., Webber, N., *Interest Rate Modelling*, Wiley 2000
Jarrow, R., *Modelling Fixed Income Securities and Interest Rate Options*, McGraw-Hill 1996, Chapter 3
Kitter, G., *Investment Mathematics for Finance and Treasury Professionals*, Wiley 1999, Chapters 3, 5
McCulloch, J., "Measuring the Term Structure of Interest Rates", *Journal of Business* 44, 1971, pp. 19–31
McCulloch, J., "The Tax-Adjusted Yield Curve", *Journal of Finance* 30, 1975, pp. 811–830
Neftci, S., *An Introduction to the Mathematics of Financial Derivatives*, 2nd edition, Academic Press 2000, Chapter 18
Nelson, C., Siegel, A., "Parsimonious Modelling of Yield Curves, *Journal of Business* 60 (4), 1987, pp. 473–489
Ross, Sheldon M., *An Introduction to Mathematical Finance*, Cambridge University Press 1999
Schaefer, S., "Measuring a Tax-Specific Term Structure of Interest Rates in the Market for British Government Securities", *Economic Journal* 91, 1981, pp. 415–438
Shiller, R., "The Term Structure of Interest Rates", in Friedman, B., Hanh, F. (eds), *Handbook of Monetary Economics*, Volume 1, North Holland 1990, Chapter 13
Suits, D., Mason, A., Chan, L., "Spline Functions Fitted by Standard Regression Methods", *Review of Economics and Statistics* 60, 1978, pp. 132–139
Sundaresan, S., *Fixed Income Markets and Their Derivatives*, South-Western 1997
Van Deventer, D., Imai, K., *Financial Risk Analytics*, Irwin 1997
Van Horne, J., *Financial Management and Policy*, 10th edition, Prentice Hall 1995
Waggoner, D., "Spline Methods for Extracting Interest Rate Curves from Coupon Bond Prices", *Working Paper, Federal Reserve Bank of Atlanta*, 1997, pp. 97–10
Windas, T., *An Introduction to Option-Adjusted Spread Analysis*, Bloomberg Publications 1993

In addition interested readers may wish to consult the following recommended references on term structure analysis.

Constantinides, G., "A Theory of the Nominal Term Structure of Interest Rates", *The Review of Financial Studies* 5 (4), 1992, pp. 531–552
Cox, J., Ingersoll, J., Ross, S., "A Re-examination of Traditional Hypotheses about the Term Structure of Interest Rates", *Journal of Finance* 36, 1981, pp. 769–799
Cox, J., Ingersoll, J., Ross, S., " A Theory of the Term Structure of Interest Rates", *Econometrica* 53, 1985, pp. 385–407
Culbertson, J., "The Term Structure of Interest Rates", *Quarterly Journal of Economics* LXXI, 1957, pp. 489–504
McCulloch, J. H., "A Reexamination of Traditional Hypotheses about the Term Structure: A Comment", *Journal of Finance* 63 (2), 1993, pp. 779–789
Shiller, R., Campbell, J., Schoenholtz, K., "Forward Rates and Future Policy: Interpreting the Term Structure of Interest Rates", *Brookings Papers on Economic Activity* 1, 1983, pp. 173–223
Stambaugh, R., "The Information in Forward Rates: Implications for Models of the Term Structure", *Journal of Financial Economics* 21, 1988, pp. 41–70

3 Forwards and Futures Valuation

To begin the discussion of interest rate futures, we discuss basic valuation concepts of forward and futures contracts. The discussion follows the approach described in Rubinstein (1999) and is adapted with permission.[1]

3.1 Introduction

A forward contract is an agreement between two parties in which the buyer contracts to purchase from the seller a specified asset, for delivery at a future date, at a price agreed today. The terms are set so that the present value of the contract is zero. For the forthcoming analysis we use the following notation:

P	is the current price of the underlying asset, also known as the *spot* price
P_T	is the price of the underlying asset at the time of delivery
X	is the delivery price of the forward contract
T	is the term to maturity of the contract in years, also referred to as the time-to-delivery
r	is the risk-free interest rate
R	is the return of the payout or its *yield*
F	is the current price of the forward contract

The payoff of a forward contract is therefore given by

$$P_T - X \tag{3.1}$$

with X set at the start so that the present value of $(P_T - X)$ is zero. The payout yield is calculated by obtaining the percentage of the spot price that is paid out on expiry.

3.1.1 *Forwards*

When a forward contract is written, its delivery price is set so that the present value of the payout is zero. This means that the forward price F is then the price on delivery which would make the present value of the payout, on the delivery date, equal to zero. That is, at the start $F = X$. This is the case only on day 1 of the contract however. From then until the contract expiry the value of X is fixed, but the forward price F will fluctuate continuously until delivery. It is the behaviour of this forward price that we wish to examine. For instance, generally as the spot price of the underlying increases, so the price of a forward contract written on the asset also increases; and vice versa.

At this stage it is important to remember that the forward price of a contract is not the same as the value of the contract, and the terms of the agreement are set so that at inception

[1] This excellent book is highly recommended, and for all students and practitioners interested in capital markets, not just those involved with derivative instruments.

10	–10
20	–8
30	–6
40	–4
50	–2
52.5	0
60	2
70	4
80	6
90	8
100	10

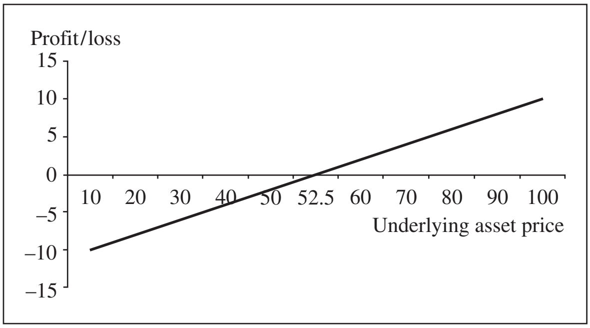

Figure 3.1: Forward contract profit/loss profile

the value is zero. The relationship given above is used to show that an equation can be derived which relates F to P, T, r and R.

Consider first the profit/loss profile for a forward contract. This is shown in Figure 3.1. The price of the forward can be shown to be related to the underlying variables as

$$F = S(r/R)^T \tag{3.2}$$

and for the one-year contract highlighted in Figure 3.1 is 52.5, where the parameters are $S = 50$, $r = 1.05$ and $R = 1.00$.

3.1.2 *Futures*

Forward contracts are tailor-made instruments designed to meet specific individual requirements. Futures contracts on the other hand are standardised contracts that are traded on recognised futures exchanges. Apart from this, the significant difference between them, and the feature that influences differences between forward and futures prices, is that profits or losses that are gained or suffered in futures trading are paid out at the end of the day. This does not occur with forwards. The majority of trading in futures contracts is always closed out, that is the position is netted out to zero before the expiry of the contract. If a position is run into the delivery month, depending on the terms and conditions of the particular exchange, the long future may be delivered into. Settlement is by physical delivery in the case of commodity futures or in cash in the case of certain financial futures. Bond futures are financial futures where any bond that is in the *delivery basket* for that contract will be delivered to the long. With both physical and financial futures, only a very small percentage of contracts are actually delivered into as the majority of trading is undertaken for hedging and speculative purposes.

With futures contracts, as all previous trading profits and losses have been settled, on the day of expiry only the additional change from the previous day needs to be accounted for. With a forward contract all loss or gain is rolled up until the expiry day and handed over as a total amount on this day.[2]

[2] We assume the parties have traded only one forward contract between them. If, as is more accurate to assume, a large number of contracts have been traded across a number of different maturity periods and perhaps instruments, as contracts expire only the net loss or gain is transferred between counterparties.

3.2 Forwards and futures

3.2.1 *Cash flow differences*

We can now look at the cash flow treatment of the two contracts in greater detail. This is illustrated in Table 3.1, which uses F to denote the price of the futures contract as well. The table shows the payoff schedule at the end of each trading day for the two instruments; assuming that they have identical terms. With the forward there is no cash flow on intermediate dates, whereas with the futures contract there is. As with the forward contract, the price of the future fixes the present value of the futures contract at zero. Each day the change in price, which at the end of the day is *marked-to-market* at the *close* price, will have resulted in either a profit or gain,[3] which is handed over or received each day as appropriate. The process of daily settlement of price movements means that the nominal delivery price can be reset each day so that the present value of the contract is always zero. This means that the future and nominal delivery prices of a futures contract are the same at the end of each trading day.

We see in Table 3.1 that there are no cash flows changing hands between counterparties to a forward contract. The price of a futures contract is reset each day; after day 1 this means it is reset from F to F_1. The amount $(F_1 - F)$ if positive, is handed over by the short future to the long future. If this amount is negative, it is paid by the long future to the short. On the expiry day T of the contract the long future will receive a settlement amount equal to $P_T - F_{T-1}$ which expresses the relationship between the price of the future and the price of the underlying asset. As significant, the daily cash flows transferred when holding a futures contract cancel each other out, so that on expiry the value of the contract is (at this stage) identical to that for a forward, that is $(P_T - F)$.

With exchange-traded contracts all market participants are deemed to conduct their trading with a central counterparty, the exchange's clearing house. This eliminates counterparty

Time	Forward contract	Futures contract
0	0	0
1	0	$F_1 - F$
2	0	$F_2 - F_1$
3	0	$F_3 - F_2$
4	0	$F_4 - F_3$
5	0	$F_5 - F_4$
…	0	…
…	0	…
…	0	…
$T - 1$	0	$F_{T-1} - F_{T-2}$
T	$P_T - F$	$P_T - F_{T-1}$
Total	$P_T - F$	$P_T - F$

Table 3.1: Cash flow process for forwards and futures contracts

[3] Or no profit or gain if the closing price is unchanged from the previous day's closing price, a *doji* as technical traders call it.

risk in all transactions, and the clearing house is able to guarantee each bargain because all participants are required to contribute to its clearing fund. This is by the process of *margin*, by which each participant deposits an *initial margin* and then, as its profits or losses are recorded, deposits further *variation margin* on a daily basis. The marking-to-market of futures contracts is an essential part of this margin process. A good description of the exchange clearing process is contained in Galitz (1995).

This is the key difference between future and forward contracts. If holding a futures position that is recording a daily profit, the receipt of this profit on a daily basis is advantageous because the funds can be reinvested while the position is still maintained. This is not available with a forward. Equally, losses are suffered on a daily basis that are not suffered by the holder of a loss-making forward position.

3.2.2 Relationship between forward and future price

Continuing with the analysis contained in Rubinstein, we wish to illustrate that under certain specified assumptions, the price of futures and forwards written with identical terms must be the same.

This can be shown in the following way. Consider two trading strategies of identical term to maturity and written on the same underlying asset; one strategy uses forward contracts while the other uses futures. Both strategies require no initial investment and are *self-financing*. The assumptions are:

- the absence of risk-free arbitrage opportunities;
- the existence of an economist's perfect market;
- certainty of returns.

Under these conditions it can be shown that the forward and future price must be identical. In this analysis the return r is the daily return (or instantaneous money market rate) and T is the maturity term in days. Let us look further at the strategies.

For the strategy employing forwards, we buy r^T forward contracts. The start forward price is $F = X$ but of course there is no cash outlay at the start, and the payoff on expiry is $r^T(P_T - F)$.

The futures strategy is more involved, due to the daily margin cash flows that are received or paid during the term of the trade. On day 1 we buy r contracts each priced at F. After the close we receive $F_1 - F$. The position is closed out and the cash received is invested at the daily rate r up to the expiry date. The return on this cash is r^{T-1} which means that on expiry we will receive an amount of $r(F_1 - F)r^{T-1}$.

The next day we purchase r^2 futures contracts at the price of F_1 and at the close the cash flow received of $F_2 - F_1$ is invested at the close of trading at r^{T-2}. Again we will receive on expiry a sum equal to $r^2(F_2 - F_1)r^{T-2}$.

This process is repeated until the expiry date, which we assume to be the delivery date. What is the net effect of following this strategy? We will receive on the expiry date a set of maturing cash flows that have been invested daily from the end of day 1. The cash sums will be

$$r^T(F_1 - F) + r^T(F_2 - F_1) + r^T(F_3 - F_2) + \cdots\cdots + r^T(P_T - F_{T-1})$$

which nets to

$$r^T(P_T - F)$$

which is also the payoff from the forward contract strategy. Both strategies have a zero cash outlay and are self-financing. The key point is that if indeed we are saying that

$$r^T(P_T - F)_{forward} = r^T(P_T - F)_{future}, \tag{3.3}$$

for the assumption of no arbitrage to hold, then $F_{forward} = F_{future}$.

3.3 The forward–spot parity

We can use the forward strategy to imply the forward price provided we know the current price of the underlying and the money market interest rate. A numerical example of the forward strategy is given in Figure 3.2, with the same parameters given earlier. We assume no-arbitrage and a perfect frictionless market.

What Figure 3.2 is saying is that it is possible to replicate the payoff profile we observed in Figure 3.1 by a portfolio composed of one unit of the underlying asset, which purchase is financed by borrowing a sum that is equal to the present value of the forward price. This borrowing is repaid on maturity and is equal to $(F/1.05) \times 1.05$ which is in fact F. In the absence of arbitrage opportunity the cost of forming the portfolio will be identical to that of the forward itself. However we have set the current cost of the forward contract at zero, which gives us

$$-50 + F/1.05 = 0.$$

We solve this equation to obtain F and this is 52.50.

So to recap the price of the forward contract is 52.50 although the present value of the forward contract when it is written is zero. Following Rubinstein we prove this in Figure 3.3.

What Figure 3.3 states is that the payoff profile for the forward can be replicated precisely by setting up a portfolio that holds R^{-T} units of the underlying asset, which is funded through borrowing a sum equal to the present value of the forward price. This borrowing is repaid at maturity, this amount being equal to

$$\left(Fr^{-T}\right) \times r^T = F.$$

The portfolio has an identical payoff profile (by design) to the forward, this being $(P_T - F)$. In a no-arbitrage environment the cost of setting up the portfolio must be equal to the

	Cash flows	
	Start date	*Expiry date*
Buy forward contract	0	$P_T - F$
Buy one unit of the underlying asset	–50	P_T
Borrow zero present-value of forward price	$F/1.05$	F
Total	$-50 + F/1.05$	$P_T - F$
Result		
Set $-50 + F/1.05$ equal to zero (no-arbitrage condition)		
Therefore $F = 52.5$		

Figure 3.2: Forward strategy

	Start date	Expiry date
Buy forward contract	0	$P_T - F$
Buy R^{-T} units of the underlying asset	$-PR^{-T}$	P_T
Borrow zero present-value of forward price	Fr^{-T}	$-F$
Total	$-PR^{-T} + Fr^{-T}$	$P_T - F$
Set		
$-PR^{-T} + Fr^{-T} = 0$		
Therefore		
$F = P(r/R)^T$		

Figure 3.3: Algebraic proof of forward price

current price of the forward, as they have identical payoffs and if one was cheaper than the other, there would be a risk-free profit for a trader who bought the cheap instrument and shorted the dear one. However we set the current cost of the forward (its present value) as zero, which means the cost of constructing the duplicating portfolio must therefore be zero as well. This gives us

$$-PR^{-T} + Fr^{-T} = 0$$

which allows us to solve for the forward price F.

The significant aspect for the buyer of a forward contract is that the payoff of the forward is identical to that of a portfolio containing an equivalent amount of the underlying asset, which has been constructed using borrowed funds. The portfolio is known as the *replicating portfolio*. The price of the forward contract is a function of the current underlying spot price, the risk-free or money market interest rate, the payoff and the maturity of the contract.

To recap then the forward–spot parity states that

$$F = P(r/R)^T. \tag{3.4}$$

It can be shown that neither of the possibilities $F > P(r/R)^T$ or $F < P(r/R)^T$ will hold unless arbitrage possibilities are admitted. The only possibility is equation (3.4), at which the futures price is *fair value*.

3.4 The basis and implied repo rate

For later analysis, we introduce now some terms used in the futures markets.

The difference between the price of a futures contract and the current underlying spot price is known as the *basis*. For bond futures contracts, which are written not on a specific bond but a *notional* bond that can in fact be represented by any bond that fits within the contract terms, the size of the basis is given by equation (3.5):

$$Basis = P_{bond} - (P_{fut} \times CF) \tag{3.5}$$

where the basis is the *gross basis* and *CF* is the *conversion factor* for the bond in question. All delivery-eligible bonds are said to be in the *delivery basket*. The conversion factor equalises each deliverable bond to the futures price.[4] The size of the gross basis represents the cost of carry associated with the bond from today to the delivery date. The bond with the lowest basis associated with it is known as the *cheapest-to-deliver* bond.

The magnitude of the basis changes continuously and this uncertainty is termed *basis risk*. Generally the basis declines over time as the maturity of the contract approaches, and converges to zero on the expiry date. The significance of basis risk is greatest for market participants who use futures contracts for hedging positions held in the underlying asset. The basis is positive or negative according to the type of market in question, and is a function of issues such as *cost of carry*. When the basis is positive, that is $F > P$, the situation is described as a *contango*, and is common in precious metals markets. A negative basis $P < F$ is described as *backwardation* and is common in oil contracts and foreign currency markets.

The hedging of futures and the underlying asset requires a keen observation of the basis. To hedge a position in a futures contract, one could run an opposite position in the underlying. However running such a position incurs the cost of carry referred to above, which depending on the nature of the asset may include storage costs, opportunity cost of interest foregone, funding costs of holding the asset and so on. The futures price may be analysed in terms of the forward–spot parity relationship and the risk-free interest rate. If we say that the risk-free rate is

$$r - 1$$

and the forward–spot parity is

$$F = P(r/R)^T$$

we can set

$$r - 1 = R(F/P)^{1/T} - 1 \tag{3.6}$$

which must hold because of the no-arbitrage assumption.

This interest rate is known as the *implied repo rate*, because it is similar to a repurchase agreement carried out with the futures market. Generally a relatively high implied repo rate is indicative of high futures prices, and the same for low implied repo rates. The rates can be used to compare contracts with each other, when these have different terms to maturity and even underlying assets. The implied repo rate for the contract is more stable than the basis; as maturity approaches the level of the rate becomes very sensitive to changes in the futures price, spot price and (by definition) time to maturity.

We present an illustration of the basis convergence in Figure 3.4. The implied repo rate is constant during this time. If we start with

$$F = P(r/R)^T$$

---

4 For a description and analysis of bond futures contracts, the basis, implied repo and the cheapest-to-deliver bond, see Choudhry (2002), an accessible account of the analysis of the gilt bond basis. Plona (1997) is also a readable treatment of the European government bond basis.

we can derive

$$r = R(F/P)^{1/T}$$

where r is the implied repo rate. The cost of carry is defined as the net cost of holding the underlying asset from trade date to expiry date (or delivery date).

The case study shows the price histories for the LIFFE[5] September 2000 long gilt futures contract.

Case study: LIFFE September 2000 long gilt futures contract

In theory, the futures contract represents the price for forward delivery of the underlying asset. Therefore one should observe a convergence in the price of the future and the price of the underlying as the contract approaches maturity.[6] The gilt price is for the UK Treasury 5¾% 2009, the cheapest-to-deliver gilt for the September 2000 contract throughout the contract's life. However note that there is in fact no actual convergence of prices; for a bond futures contract, the convergence is best analysed using the basis, as Figure 3.4 shows (the x-axis lists the contract days to maturity, which are in fact the dates 29 June 2000 to 27 September 2000, the contract expiry date). The associated

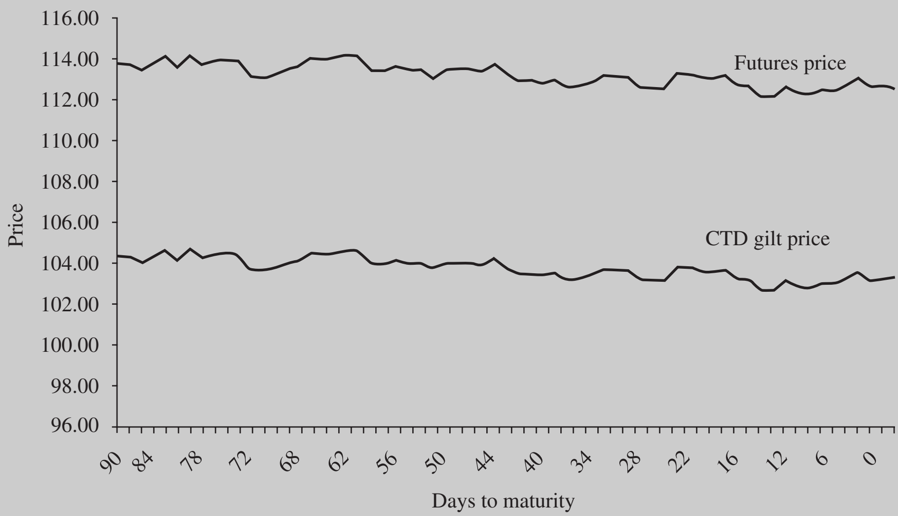

Figure 3.4

---

[5] The London International Financial Futures and Options Exchange (LIFFE), pronounced "life". The terms of this contract are available at the exchange's website, www.liffe.com which also contains other useful information on contracts traded at the exchange.

[6] For example, see the diagram on page 105 of Rubinstein (1999).

yields are shown in Figure 3.5, and do demonstrate slight convergence on the very last day of trading. However the liquidity of bond futures is very low in the delivery month, and decreases further in the last three days of trading up to the expiry Wednesday.

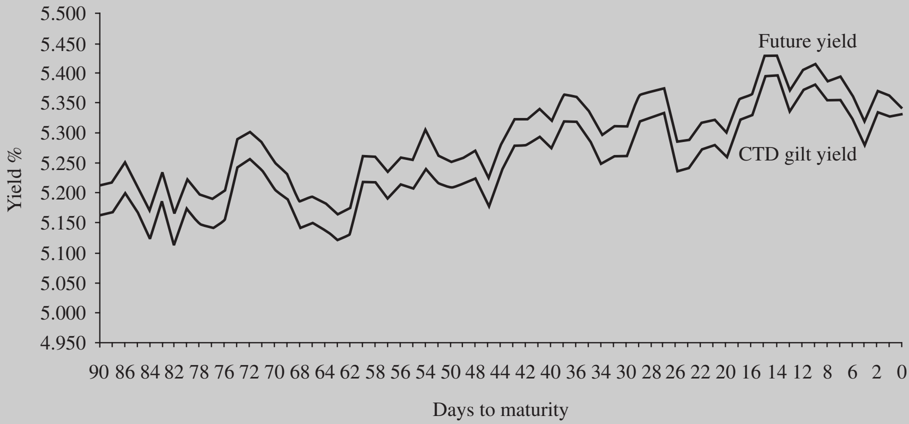

Figure 3.5

The basis history is shown in Figure 3.6, which confirms the convergence, as the contract approaches expiry; the dates match those used for Figure 3.4. Figure 3.7 shows the stability of the implied repo rate, confirming the analysis we suggested earlier. The spike towards maturity also illustrates the sensitivity of the implied repo rate to very small changes in cash or futures price.

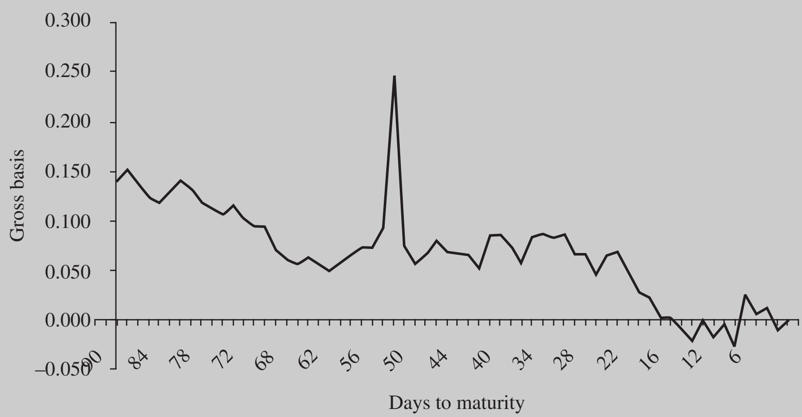

Figure 3.6

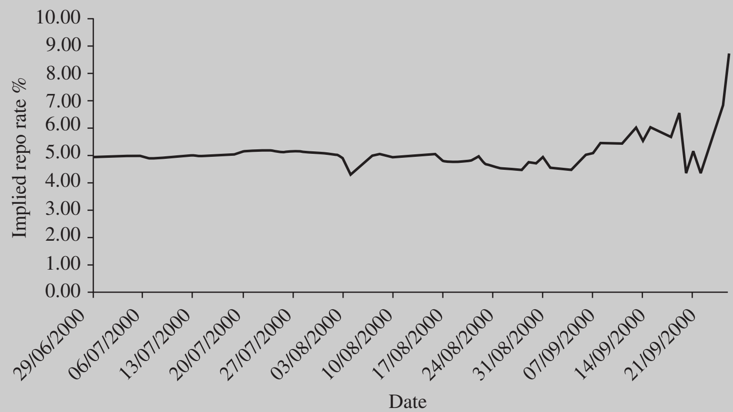

Figure 3.7

Selected bibliography and references

Choudhry, M., *The Repo Handbook*, Butterworth-Heinemann 2002

French, K., "A Comparison of Futures and Forwards Prices", *Journal of Financial Economics* 12, November 1983, pp. 311–342

Galitz, L., *Financial Engineering*, FT Pitman 1995, Chapters 3, 4, 6–8

Grinblatt, M., Jegadeesh, N., "Futures vs. Forward Prices: Implications for Swap Pricing and Derivatives Valuation", in Jegadeesh, N., Tuckman, B. (eds), *Advanced Fixed-Income Valuation Tools*, Wiley 2000

Hull, J., *Options, Futures and Other Derivatives*, 4th edition, Prentice Hall 2000

Jarrow, R., Oldfield, G., "Forward Contracts and Futures Contracts", *Journal of Financial Economics* 9, December 1981, pp. 373–382

Kolb, R., *Futures, Options and Swaps*, 3rd edition, Blackwell 2000

Park, H., Chen, A., "Differences between Futures and Forward Prices: A Further Investigation of Marking to Market Effects", *Journal of Futures Markets* 5, February 1985, pp. 77–88

Plona, C., *The European Bond Basis*, McGraw-Hill 1997

Rubinstein, M., *Rubinstein on Derivatives*, RISK Books 1999

4 FRAs and Interest Rate Futures

The market in short-term interest rate derivatives is a large and liquid one, and the instruments involved are used for a variety of purposes. In this chapter we review the two main contracts used in money markets trading, the short-term *interest rate future* and the *forward rate agreement*. Money market derivatives are priced on the basis of the forward rate, and are flexible instruments for hedging against or speculating on forward interest rates. The FRA and the exchange-traded interest rate future both date from around the same time, and although initially developed to hedge forward interest rate exposure, they now have a variety of uses. In this chapter the instruments are introduced and analysed, and there is a review of their main uses.

4.1 Forward rate agreements

A *forward rate agreement* (FRA) is an OTC derivative instrument that trades as part of the money markets. It is essentially a forward-starting loan, but with no exchange of principal, so that only the difference in interest rates is traded. An FRA is a forward-dated loan, dealt at a fixed rate, but with no exchange of principal – only the interest applicable on the notional amount between the rate dealt and the actual rate prevailing at the time of settlement changes hands. So FRAs are *off-balance sheet* (OBS) instruments. By trading today at an interest rate that is effective at some point in the future, FRAs enable banks and corporates to hedge interest rate exposure. They may also be used to speculate on the level of future interest rates.

4.1.1 *Definition*

An FRA is an agreement to borrow or lend a *notional* cash sum for a period of time lasting up to twelve months, starting at any point over the next twelve months, at an agreed rate of interest (the FRA rate). The "buyer" of an FRA is borrowing a notional sum of money while the "seller" is lending this cash sum. Note how this differs from all other money market instruments. In the cash market, the party buying a CD or bill, or bidding for stock in the repo market, is the lender of funds. In the FRA market, to "buy" is to "borrow". Of course, we use the term "notional" because with an FRA no borrowing or lending of cash actually takes place, as it is an off-balance sheet product. The notional sum is simply the amount on which interest payment is calculated.

So when an FRA is traded, the buyer is borrowing (and the seller is lending) a specified notional sum at a fixed rate of interest for a specified period, the "loan" to commence at an agreed date in the future. The *buyer* is the notional borrower, and so if there is a rise in interest rates between the date that the FRA is traded and the date that the FRA comes into effect, she will be protected. If there is a fall in interest rates, the buyer must pay the difference between the rate at which the FRA was traded and the actual rate, as a percentage of the notional sum. The buyer may be using the FRA to hedge an actual exposure, that is an actual borrowing of money, or simply speculating on a rise in interest rates. The counterparty to the transaction, the *seller* of the FRA, is the notional lender of funds, and has fixed the rate for lending funds. If there is a fall in interest rates the seller will gain, and if there is a rise in rates the seller will pay. Again, the seller may have an actual loan of cash to hedge or be a speculator.

In FRA trading only the payment that arises as a result of the difference in interest rates changes hands. There is no exchange of cash at the time of the trade. The cash payment that does arise is the difference in interest rates between that at which the FRA was traded and the actual rate prevailing when the FRA matures, as a percentage of the notional amount. FRAs are traded by both banks and corporates and between banks. The FRA market is very liquid in all major currencies and rates are readily quoted on screens by both banks and brokers. Dealing is over the telephone or over a dealing system such as Reuters.

The terminology quoting FRAs refers to the borrowing time period and the time at which the FRA comes into effect (or matures). Hence if a buyer of an FRA wished to hedge against a rise in rates to cover a three-month loan starting in three months' time, she would transact a "three-against-six month" FRA, or more usually a 3×6 or 3-v-6 FRA. This is referred to in the market as a "threes-sixes" FRA, and means a three-month loan beginning in three months' time. So a "ones-fours" FRA (1-v-4) is a three-month loan in one month's time, and a "threes-nines" FRA (3-v-9) is six-month money in three months' time.

Note that when one buys an FRA one is "borrowing" funds. This differs from cash products such as CD or repo, as well as interest rate futures, where "buying" is lending funds.

Example 4.1
A company knows that it will need to borrow £1 million in three months' time for a twelve-month period. It can borrow funds today at LIBOR + 50 basis points. LIBOR rates today are at 5% but the company's treasurer expects rates to go up to about 6% over the next few weeks. So the company will be forced to borrow at higher rates unless some sort of hedge is transacted to protect the borrowing requirement. The treasurer decides to buy a 3-v-15 ("threes-fifteens") FRA to cover the twelve-month period beginning three months from now. A bank quotes 5½% for the FRA which the company buys for a notional £1 million. Three months from now rates have indeed gone up to 6%, so the treasurer must borrow funds at 6½% (the LIBOR rate plus spread), however she will receive a settlement amount which will be the difference between the rate at which the FRA was bought and today's twelve-month LIBOR rate (6%) as a percentage of £1 million, which will compensate for some of the increased borrowing costs.

4.2 FRA mechanics

In virtually every market FRAs trade under a set of terms and conventions that are identical. The British Bankers' Association (BBA) has compiled standard legal documentation to cover FRA trading. The following standard terms are used in the market.

- **Notional sum**: The amount for which the FRA is traded.
- **Trade date**: The date on which the FRA is dealt.

- **Settlement date**: The date on which the notional loan or deposit of funds becomes effective, that is, is said to begin. This date is used, in conjunction with the notional sum, for calculation purposes only as no actual loan or deposit takes place.
- **Fixing date**: This is the date on which the *reference rate* is determined, that is, the rate to which the FRA dealing rate is compared.
- **Maturity date**: The date on which the notional loan or deposit expires.
- **Contract period**: The time between the settlement date and maturity date.
- **FRA rate**: The interest rate at which the FRA is traded.
- **Reference rate**: This is the rate used as part of the calculation of the settlement amount, usually the LIBOR rate on the fixing date for the contract period in question.
- **Settlement sum**: The amount calculated as the difference between the FRA rate and the reference rate as a percentage of the notional sum, paid by one party to the other on the settlement date.

These terms are illustrated in Figure 4.1.

The spot date is usually two business days after the trade date, however it can by agreement be sooner or later than this. The settlement date will be the time period after the spot date referred to by the FRA terms, for example a 1×4 FRA will have a settlement date one calendar month after the spot date. The fixing date is usually two business days before the settlement date. The settlement sum is paid on the settlement date, and as it refers to an amount over a period of time that is paid up front, at the start of the contract period, the calculated sum is discounted present value. This is because a normal payment of interest on a loan/deposit is paid at the end of the time period to which it relates; because an FRA makes this payment at the *start* of the relevant period, the settlement amount is a discounted present value sum.

With most FRA trades the reference rate is the LIBOR fixing on the fixing date.

The settlement sum is calculated after the fixing date, for payment on the settlement date. We may illustrate this with a hypothetical example. Consider a case where a corporate has bought £1 million notional of a 1-v-4 FRA, and dealt at 5.75%, and that the market rate is 6.50% on the fixing date. The contract period is 90 days. In the cash market the extra interest charge that the corporate would pay is a simple interest calculation, and is:

$$\frac{6.50 - 5.75}{100} \times 1{,}000{,}000 \times \frac{91}{365} = £1869.$$

This extra interest that the corporate is facing would be payable with the interest payment for the loan, which (as it is a money market loan) is when the loan matures. Under an FRA then, the settlement sum payable should, if it was paid on the same day as the cash market

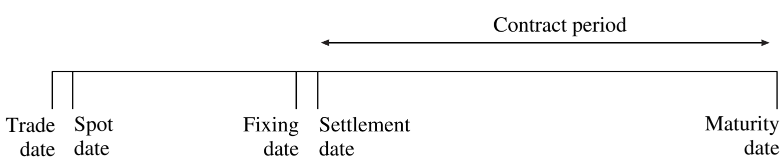

Figure 4.1: Key dates in an FRA trade

interest charge, be exactly equal to this. This would make it a perfect hedge. As we noted above though, FRA settlement value is paid at the start of the contract period, that is, the beginning of the underlying loan and not the end. Therefore the settlement sum has to be adjusted to account for this, and the amount of the adjustment is the value of the interest that would be earned if the unadjusted cash value was invested for the contract period in the money market. The settlement value is given by equation (4.1).

$$\text{Settlement} = \frac{(r_{ref} - r_{FRA}) \times M \times \frac{n}{B}}{1 + \left(r_{ref} \times \frac{n}{B}\right)} \tag{4.1}$$

where

r_{ref}	is the reference interest fixing rate;
r_{FRA}	is the FRA rate or *contract rate*;
M	is the notional value;
n	is the number of days in the contract period;
B	is the day-count base (360 or 365).

Equation (4.1) simply calculates the extra interest payable in the cash market, resulting from the difference between the two interest rates, and then discounts the amount because it is payable at the start of the period and not, as would happen in the cash market, at the end of the period.

In our hypothetical illustration, as the fixing rate is higher than the dealt rate, the corporate buyer of the FRA receives the settlement sum from the seller. This then compensates the corporate for the higher borrowing costs that he would have to pay in the cash market. If the fixing rate had been lower than 5.75%, the buyer would pay the difference to the seller, because the cash market rates will mean that he is subject to a lower interest rate in the cash market. What the FRA has done is hedge the interest rate, so that whatever happens in the market, it will pay 5.75% on its borrowing.

A market maker in FRAs is trading short-term interest rates. The settlement sum is the value of the FRA. The concept is exactly as with trading short-term interest rate futures; a trader who buys an FRA is running a long position, so that if on the fixing date $r_{ref} > r_{FRA}$, the settlement sum is positive and the trader realises a profit. What has happened is that the trader, by buying the FRA, "borrowed" money at an interest rate, which subsequently rose. This is a gain, exactly like a *short* position in an interest rate future, where if the price goes down (that is, interest rates go up), the trader realises a gain. Equally a "short" position in an FRA, put on by selling an FRA, realises a gain if on the fixing date $r_{ref} < r_{FRA}$.

4.2.1 *FRA pricing*

As their name implies, FRAs are forward rate instruments and are priced using the forward rate principles we established earlier. Consider an investor who has two alternatives, either a six-month investment at 5% or a one-year investment at 6%. If the investor wishes to invest for six months and then roll over the investment for a further six months, what rate is required for the rollover period such that the final return equals the 6% available from the one-year investment? If we view an FRA rate as the breakeven forward rate between the two periods, we simply solve for this forward rate and that is our approximate FRA rate. This rate is sometimes referred to as the interest rate "gap" in the money markets (not to be confused with an interbank desk's *gap risk*, the interest rate exposure arising from the net maturity position of its assets and liabilities).

We can use the standard forward-rate breakeven formula to solve for the required FRA rate; we established this relationship earlier when discussing the calculation of forward rates that are arbitrage-free. The relationship given in equation (4.2) connects simple (bullet) interest rates for periods of time up to one year, where no compounding of interest is required. As FRAs are money market instruments we are not required to calculate rates for periods in excess of one year,[1] where compounding would need to be built into the equation. This is given by equation (4.2):

$$(1 + r_2 t_2) = (1 + r_1 t_1)(1 + r_f t_f) \tag{4.2}$$

where

- r_2 is the cash market interest rate for the long period;
- r_1 is the cash market interest rate for the short period;
- r_f is the forward rate for the gap period;
- t_2 is the time period from today to the end of the long period;
- t_1 is the time period from today to the end of the short period;
- t_f is the forward gap time period, or the contract period for the FRA.

This is illustrated diagrammatically in Figure 4.2.

The time period t_1 is the time from the dealing date to the FRA settlement date, while t_2 is the time from the dealing date to the FRA maturity date. The time period for the FRA (contract period) is t_2 minus t_1. We can replace the symbol t for time period with n for the actual number of days in the time periods themselves. If we do this and then rearrange the equation to solve for r_{FRA} the FRA rate, we obtain (4.3):

$$r_{FRA} = \frac{r_2 n_2 - r_1 n_1}{n_{fra}\left(1 + r_1 \dfrac{n_1}{365}\right)} \tag{4.3}$$

where

- n_1 is the number of days from the dealing date or spot date to the settlement date;
- n_2 is the number of days from the dealing date or spot date to the maturity date;
- r_1 is the spot rate to the settlement date;
- r_2 is the spot rate from the spot date to the maturity date;
- n_{fra} is the number of days in the FRA contract period;
- r_{FRA} is the FRA rate.

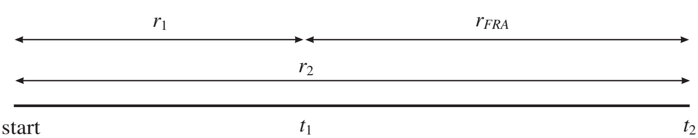

Figure 4.2: Rates used in FRA pricing

[1] Although it is of course possible to trade FRAs with contract periods greater than one year, for which a different pricing formula must be used.

If the formula is applied to say the US money markets, the 365 in the equation is replaced by 360, the day count base for that market.

In practice FRAs are priced off the exchange-traded short-term interest rate future for that currency, so that sterling FRAs are priced off LIFFE short sterling futures. Traders normally use a spreadsheet pricing model that has futures prices directly fed into it. FRA positions are also usually hedged with other FRAs or short-term interest rate futures.

Example 4.2: Hedging an FRA position

An FRA market maker sells a EUR 100 million 3-v-6 FRA, that is, an agreement to make a notional deposit (without exchange of principal) for three months in three months' time, at a rate of 7.52%. He is exposed to the risk that interest rates will have risen by the FRA settlement date in three months' time.

Date	14 December
3-v-6 FRA rate	7.52%
March futures price	92.50%
Current spot rate	6.85%

Action

The dealer first needs to calculate a precise hedge ratio. This is a three-stage process:

1. Calculate the nominal value of a basis point move in LIBOR on the FRA settlement payment;

$$BPV = FRA_{nom} \times 0.01\% \times \frac{n}{360}.$$

Therefore: €100,000,000 × 0.01% × 90/360 = €2500.

2. Find the present value of 1. by discounting it back to the transaction date using the FRA and spot rates;

Present value of a basis point move =

$$\frac{\text{Nominal value of basis point}}{\left(1 + \text{spot rate} \times \dfrac{\text{Days in hedge period}}{360}\right) \times \left(1 + \text{FRA rate} \times \dfrac{\text{Days in hedge period}}{360}\right)}.$$

Therefore:

$$\frac{€2500}{\left(1 + 6.85\% \times \dfrac{90}{360}\right) \times \left(1 + 7.52\% \times \dfrac{90}{360}\right)}.$$

3. Determine the correct hedge ratio by dividing 2. by the futures tick value.

$$\text{Hedge ratio} = \frac{2412}{25} = 96.48.$$

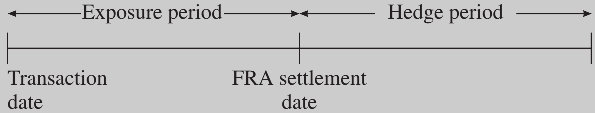

Figure 4.3: FRA dates

The appropriate number of contracts for the hedge of a EUR 100,000,000 3-v-6 FRA would therefore be 96 or 97, as the fraction is under one-half, 96 is correct. To hedge the risk of an increase in interest rates, the trader sells 96 ECU three months' futures contracts at 92.50. Any increase in rates during the hedge period should be offset by a gain realised on the futures contracts through daily variation margin receipts.

Outcome

Date	15 March
Three month LIBOR	7.625%
March EDSP	92.38

The hedge is lifted upon expiry of the March futures contracts. Three-month LIBOR on the FRA settlement date has risen to 7.625% so the trader incurs a loss of EUR 25,759 on his FRA position (i.e., EUR 26,250 discounted back over the three month FRA period at current LIBOR rate), calculated as follows:

$$\frac{(\text{LIBOR} - \text{FRA rate}) \times (\text{days in FRA period}/360) \times \text{Contract Nominal Amount}}{1 + \text{LIBOR rate} \times (\text{days in FRA period}/360)}.$$

Therefore:

$$\frac{26{,}250^{*}}{1 + 7.625\% \times \dfrac{90}{360}} = \text{EUR } 25{,}759$$

$$^{*}\text{i.e., } 0.105\% \times \frac{90}{360} \times \text{EUR } 100{,}000{,}000$$

Futures P/L: 12 ticks (92.50 – 92.38) × €25 × 96 contracts = EUR 28,800.

Conclusion

The EUR 25,759 loss on the FRA position is more than offset by the EUR 28,800 profit on the futures position when the hedge is lifted.

If the dealer has sold 100 contracts his futures profit would have been EUR 30,000, and, accordingly, a less accurate hedge. The excess profit in the hedge position can mostly be attributed to the arbitrage profit realised by the market maker (i.e., the market maker has sold the FRA for 7.52% and in effect bought it back in the futures market by selling futures at 92.50 or 7.50% for a 2 tick profit.)

4.3 Forward contracts

A forward contract is an OTC instrument with terms set for delivery of an underlying asset at some point in the future. That is, a forward contract fixes the price and the conditions now for an asset that will be delivered in the future. As each contract is tailor-made to suit user requirements, a forward contract is not as liquid as an exchange-traded futures contract with standardised terms.

The theoretical textbook price of a forward contract is the spot price of the underlying asset plus the funding cost associated with holding the asset until forward expiry date, when the asset is delivered. More formally we may write the price of a forward contract (written on an underlying asset that pays no dividends, such as a zero-coupon bond), as equation (4.5):

$$P_{fwd} = P_{spot}e^{rn} \tag{4.5}$$

where

P_{spot}	is the price of the underlying asset of the forward contract;
r	is the continuously compounded risk-free interest rate for a period of maturity n;
n	is the term to maturity of the forward contract in days.

The rule of no-arbitrage pricing states that equation (4.5) must be true. If $P_{fwd} < P_{und}e^{rn}$ then a trader could buy the cheaper instrument, the forward contract, and simultaneously sell the underlying asset. The proceeds from the short sale could be invested at r for n days; on expiry the short position in the asset is closed out at the forward price P_{fwd} and the trader will have generated a profit of $P_{und}e^{rn} - P_{fwd}$. In the opposite scenario, where $P_{fwd} > P_{und}e^{rn}$, a trader could put on a long position in the underlying asset, funded at the risk-free interest rate r for n days, and simultaneously sell the forward contract. On expiry the asset is sold under the terms of the forward contract at the forward price and the proceeds from the sale used to close out the funding initially taken on to buy the asset. Again a profit would be generated, which would be equal to the difference between the two prices.

The relationship described here is used by the market to assume that forward rates implied by the price of short-term interest rate futures contracts are equal to forward rates given by a same-maturity forward contract. Although this assumption holds good for futures contracts with a maturity of up to three or four years, it breaks down for longer-dated futures and forwards.

4.4 Short-term interest rate futures

A *futures* contract is a transaction that fixes the price today for a commodity that will be delivered at some point in the future. Financial futures fix the price for interest rates, bonds, equities and so on, but trade in the same manner as commodity futures. Contracts for futures are standardised and traded on recognised exchanges. In London the main futures exchange is LIFFE, although other futures are also traded on for example, the International Petroleum Exchange and the London Metal Exchange. The money markets trade short-term interest rate futures, which fix the rate of interest on a notional fixed term deposit of money (usually for 90 days or three months) for a specified period in the future. The sum is notional because no actual sum of money is deposited when buying or selling futures; the instrument is off-balance sheet. Buying such a contract is equivalent to making a notional deposit, while selling a contract is equivalent to borrowing a notional sum.

The three-month interest rate future is the most widely used instrument used for hedging interest rate risk.

Name	90-day sterling LIBOR interest rate future
Contract size	£500,000
Delivery months	March, June, September, December
Delivery date	First business day after the last trading day
Last trading day	Third Wednesday of delivery month
Price	100 minus interest rate
Tick size	0.005
Tick value	£6.25
Trading hours	LIFFE CONNECT™ 0805–1800 hours

Table 4.1: Description of LIFFE short sterling future contract. Source: LIFFE

The LIFFE exchange in London trades short-term interest rate futures for major currencies including sterling, euros, yen and Swiss francs. Table 4.1 summarises the terms for the short sterling contract as traded on LIFFE.

The original futures contracts related to physical commodities, which is why we speak of *delivery* when referring to the expiry of financial futures contracts. Exchange-traded futures such as those on LIFFE are set to expire every quarter during the year. The short sterling contract is a deposit of cash, so as its price refers to the rate of interest on this deposit, the price of the contract is set as $P = 100 - r$ where P is the price of the contract and r is the rate of interest at the time of expiry implied by the futures contract. This means that if the price of the contract rises, the rate of interest implied goes down, and vice versa. For example the price of the June 1999 short sterling future (written as Jun99 or M99, from the futures identity letters of H, M, U and Z for contracts expiring in March, June, September and December respectively) at the start of trading on 13 March 1999 was 94.880, which implied a three-month LIBOR rate of 5.12% on expiry of the contract in June. If a trader bought 20 contracts at this price and then sold them just before the close of trading that day, when the price had risen to 94.96, an implied rate of 5.04%, she would have made 16 ticks profit or £2000. That is, a 16 tick upward price movement in a long position of 20 contracts is equal to £2000. This is calculated as follows:

$$\text{Profit} = \text{Ticks gained} \times \text{Tick value} \times \text{Number of contracts}$$
$$\text{Loss} = \text{Ticks lost} \times \text{Tick value} \times \text{Number of contracts}.$$

The tick value for the short sterling contract is straightforward to calculate, since we know that the contract size is £500,000, there is a minimum price movement (tick movement) of 0.005% and the contract has a three-month "maturity":

$$\text{Tick value} = 0.005\% \times £500{,}000 \times \frac{3}{12} = £6.25.$$

The profit made by the trader in our example is logical because if we buy short sterling futures we are depositing (notional) funds; if the price of the futures rises, it means the interest rate has fallen. We profit because we have "deposited" funds at a higher rate beforehand. If we expected sterling interest rates to rise, we would sell short sterling futures, which is equivalent to borrowing funds and locking in the loan rate at a lower level.

Note how the concept of buying and selling interest rate futures differs from FRAs: if we buy an FRA we are borrowing notional funds, whereas if we buy a futures contract we are depositing notional funds. If a position in an interest rate futures contract is held to expiry, cash settlement will take place on the delivery day for that contract.

Short-term interest rate contracts in other currencies are similar to the short sterling contract and trade on exchanges such as Eurex in Frankfurt and MATIF in Paris.

4.4.1 *Pricing interest rate futures*

The price of a three-month interest rate futures contract is the implied interest rate for that currency's three-month rate at the time of expiry of the contract. Therefore there is always a close relationship and correlation between futures prices, FRA rates (which are derived from futures prices) and cash market rates. On the day of expiry the price of the future will be equal to the LIBOR rate as fixed that day. This is known as the exchange delivery settlement price (EDSP) and is used in the calculation of the delivery amount. During the life of the contract its price will be less closely related to the actual three-month LIBOR rate *today*, but closely related to the *forward rate* for the time of expiry.

Equation (4.1) was our basic forward rate formula for money market maturity forward rates, which we adapted to use as our FRA price equation. If we incorporate some extra terminology to cover the dealing dates involved it can also be used as our futures price formula. Let us say that:

T_0 is the trade date;
T_M is the contract expiry date;
T_{CASH} is the value date for cash market deposits traded on T_0;
T_1 is the value date for cash market deposits traded on T_M;
T_2 is the maturity date for a three-month cash market deposit traded on T_M.

We can then use equation (4.2) as our futures price formula to obtain P_{fut}, the futures price for a contract up to the expiry date.

$$P_{fut} = 100 - \frac{r_2 n_2 - r_1 n_1}{n_f \left(1 + r_1 \frac{n_1}{365}\right)} \tag{4.6}$$

where

P_{fut} is the futures price;
r_1 is the cash market interest rate to T_1;
r_2 is the cash market interest rate to T_2;
n_1 is the number of days from T_{CASH} to T_1;
n_2 is the number of days from T_{CASH} to T_2;
n_f is the number of days from T_1 to T_2.

The formula uses a 365 day count convention which applies in the sterling money markets; where a market uses a 360-day base this must be used in the equation instead.

In practice the price of a contract at any one time will be close to the theoretical price that would be established by equation (4.6) above. Discrepancies will arise for supply and demand reasons in the market, as well as because LIBOR rates is often quoted only to the nearest sixteenth or 0.0625. The price between FRAs and futures is correlated very closely, in fact banks will often price FRAs using futures, and use futures to hedge their FRA books.

When hedging an FRA book with futures, the hedge is quite close to being exact, because the two prices track each other almost tick for tick.[2] However the tick value of a futures contract is fixed, and uses (as we saw above) a 3/12 basis, while FRA settlement values use a 360 or 365-day base. The FRA trader will be aware of this when putting on her hedge.

In our discussion of forward rates in Chapter 3 we emphasised that they were the market's view on future rates using all information available today. Of course a futures price today is very unlikely to be in line with the actual three-month interest rate that is prevailing at the time of the contract's expiry. This explains why prices for futures and actual cash rates will differ on any particular day. Up until expiry the futures price is the implied forward rate; of course there is always a discrepancy between this forward rate and the cash market rate *today*. The gap between the cash price and the futures price is known as the *basis*. This is defined as:

$$\text{Basis} = \text{Cash price} - \text{Futures price}.$$

At any point during the life of a futures contract prior to final settlement – at which point futures and cash rates converge – there is usually a difference between current cash market rates and the rates implied by the futures price. This is the difference we've just explained; in fact the difference between the price implied by the current three-month interbank deposit and the futures price is known as *simple basis*, but it is what most market participants refer to as the basis. Simple basis consists of two separate components, *theoretical basis* and *value basis*. Theoretical basis is the difference between the price implied by the current three-month interbank deposit rate and that implied by the theoretical fair futures price based on cash market forward rates, given by equation (4.6) above. This basis may be either positive or negative depending on the shape of the yield curve.

The value basis is the difference between the theoretical fair futures price and the actual futures price. It is a measure of how under- or over-valued the futures contract is relative to its fair value. Value basis reflects the fact that a futures contract does not always trade at its mathematically calculated theoretical price, due to the impact of market sentiment and demand and supply. The theoretical and value bases can and do move independently of one another and in response to different influences. Both however converge to zero on the last trading day when final cash settlement of the futures contract is made.

Futures contracts do not in practice provide a precise tool for locking into cash market rates today for a transaction that takes place in the future, although this is what they are in theory designed to do. Futures do allow a bank to lock in a rate for a transaction to take place in the future, and this rate is the *forward rate*. The basis is the difference between today's cash market rate and the forward rate on a particular date in the future. As a futures contract approaches expiry, its price and the rate in the cash market will converge (the process is given the name *convergence*). As we noted earlier this is given by the EDSP and the two prices (rates) will be exactly in line at the exact moment of expiry.

---

2 That is, the basis risk is minimised.

Example 4.3: The Eurodollar futures contract

The Eurodollar futures contract is traded on the Chicago Mercantile Exchange. The underlying asset is a deposit of US dollars in a bank outside the United States, and the contract is on the rate on dollar 90-day LIBOR. The Eurodollar future is cash settled on the second business day before the third Wednesday of the delivery month (London business day). The final settlement price is used to set the price of the contract, given by:

$$10{,}000(100 - 0.25r)$$

where r is the quoted Eurodollar rate at the time. This rate is the actual 90-day Eurodollar deposit rate.

The longest-dated Eurodollar contract has an expiry date of 10 years. The market assumes that futures prices and forward prices are equal: In practice it also holds for short-dated futures contracts, but is inaccurate for longer-dated futures contracts. Therefore using futures contracts with a maturity greater than five years to calculate zero-coupon rates or implied forward rates will produce errors in results, which need to be taken into account if the derived rates are used to price other instruments such as swaps.

4.4.2 *Hedging using interest rate futures*

Banks use interest rate futures to hedge interest rate risk exposure in cash and OBS instruments. Bond trading desks also often use futures to hedge positions in bonds of up to two or three years' maturity, as contracts are traded up to three years' maturity. The liquidity of such "far month" contracts is considerably lower than for near month contracts and the "front month" contract (the current contract, for the next maturity month). When hedging a bond with a maturity of say two years, the trader will put on a *strip* of futures contracts that matches as near as possible the expiry date of the bond.

The purpose of a hedge is to protect the value of a current or anticipated cash market or OBS position from adverse changes in interest rates. The hedger will try to offset the effect of the change in interest rate on the value of their cash position with the change in value of their hedging instrument. If the hedge is an exact one the loss on the main position should be compensated by a profit on the hedge position. If the trader is expecting a fall in interest rates and wishes to protect against such a fall they will buy futures, known as a long hedge, and will sell futures (a short hedge) if wishing to protect against a rise in rates.

Bond traders also use three-month interest rate contracts to hedge positions in short-dated bonds; for instance, a market maker running a short-dated bond book would find it more appropriate to hedge his book using short-dated futures rather than the longer-dated bond futures contract. When this happens it is important to accurately calculate the correct number of contracts to use for the hedge. To construct a bond hedge it will be necessary to use a strip of contracts, thus ensuring that the maturity date of the bond is covered by the longest-dated futures contract. The hedge is calculated by finding the sensitivity of each cash flow to changes in each of the relevant forward rates. Each cash flow is considered individually and the hedge values are then aggregated and rounded to the nearest whole number of contracts.

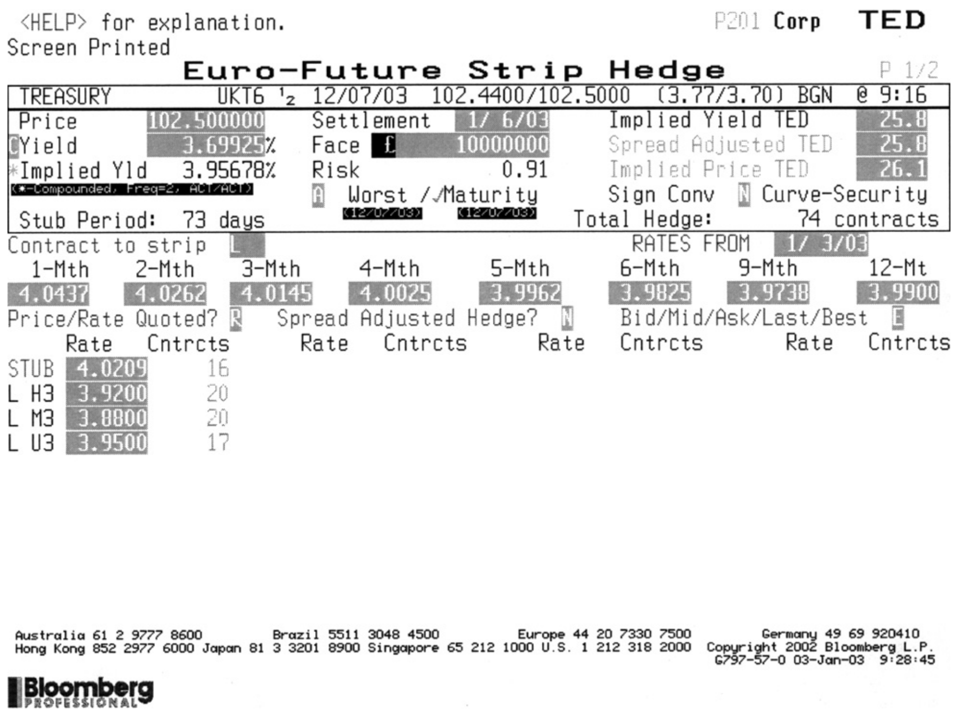

Figure 4.4: Bloomberg screen TED screen, showing futures strip hedge for position in UK gilt 6½% 2003, on 3 January 2003 ©Bloomberg LP. Reproduced with permission.

Figure 4.4 is a reproduction of page TED on a Bloomberg terminal, which calculates the strip hedge for short-dated bonds.[3] The example shown is for short sterling contract hedge for a position in the UK 6½% 2003 short-dated gilt, for settlement on 6 January 2003. The screen shows the number of each contract that must be bought (or sold) to hedge the position, which in the example is a holding of £10 million of the bond. The "stub" requirement is met using the near month contract. A total of 73 contracts are required.

The following examples illustrate hedging with short-term interest rate contracts.

Example 4.4: Hedging a forward three-month lending requirement

On 1 June a corporate treasurer is expecting a cash inflow of £10 million in three months' time (1 December), which they will then invest for three months. The treasurer expects that interest rates will fall over the next few weeks and wishes to protect themself against such a fall. This can be done using short sterling futures. Market rates on 1 June are as follows:

3-mo LIBOR	6½%
Sep futures price	93.220

[3] This screen was introduced in 1995, the author recalls being in message communication with Bloomberg in Princeton discussing the screen before it went live!

The treasurer buys 20 September short sterling futures at 93.220, this number being exactly equivalent to a sum of £10 million. This allows them to lock in a forward *lending* rate of 6.78%, if we assume there is no bid–offer quote spread.

$$\begin{aligned}\text{Expected lending rate} &= \text{rate implied by futures price}\\ &= 100 - 93.220\\ &= 6.78\%.\end{aligned}$$

On 1 September market rates are as follows:

3-mo LIBOR	6¼%
Set futures price	93.705

The treasurer unwinds the hedge at this price.

$$\begin{aligned}\text{Futures p/l} &= 97 \text{ ticks}(93.705 - 93.22) \text{ or } 0.485\%\\ \text{Effective lending rate} &= \text{3-mo LIBOR futures profit}\\ &= 6.25\% + 0.485\%\\ &= 6.735\%.\end{aligned}$$

The treasurer was quite close to achieving their target lending rate of 6.78% and the hedge has helped to protect against the drop in LIBOR rates from 6½% to 6¼%, due to the profit from the futures transaction.

In the real world the cash market bid–offer spread will impact the amount of profit/loss from the hedge transaction. Futures generally trade and settle near the offered side of the market rate (LIBOR) whereas lending, certainly by corporates, will be nearer the LIBID rate.

Example 4.5: Hedging a forward six-month borrowing requirement[4]

A treasury dealer has a six-month borrowing requirement for DEM30 million in three months' time, on 16 September. She expects interest rates to rise by at least ½% before that date and would like to lock in a future borrowing rate. The scenario is detailed below.

Date:	16 June
Three-month LIBOR	6.0625%
Six-month LIBOR	6.25
Sep futures contract	93.66
Dec futures contract	93.39

In order to hedge a six-month DEM30 million exposure the dealer needs to use a total of 60 futures contracts, as each has a nominal value of DEM1 million, and corresponds

[4] This example predates the introduction of the euro.

to a three-month notional deposit period. The dealer decides to sell 30 September futures contracts and 30 December futures contracts, which is referred to as a *strip* hedge. The expected forward borrowing rate that can be achieved by this strategy, where the expected borrowing rate is r_f, is calculated as follows:

$$1 + r_f \times \frac{\text{days in period}}{360} = \left(1 + \text{Sep implied rate} \times \frac{\text{Sep days period}}{360}\right).$$

Therefore we have:

$$1 + r_f \times \frac{180}{360} = \left(1 + 0.0634 \times \frac{90}{360}\right) \times \left(1 + 0.0661 \times \frac{90}{360}\right) \rightarrow r_f = 6.53\%.$$

The rate r_f is sometimes referred to as the "strip rate".

The hedge is unwound upon expiry of the September futures contract. Assume the following rates now prevail:

Three-month LIBOR	6.4375%
Six-month LIBOR	6.8125
Sep futures contract	93.56
Dec futures contract	92.93

The futures profit-and-loss is:

September contract:	+10 ticks
December contract:	+46 ticks

This represents a 56 tick or 0.56% profit in three-month interest rate terms, or 0.28% in six-month interest rate terms. The effective borrowing rate is the six-month LIBOR rate minus the futures profit, or:

$$6.8125\% - 0.28\% \text{ or } 6.5325\%.$$

In this case the hedge has proved effective because the dealer has realised a borrowing rate of 6.5325%, which is close to the target strip rate of 6.53%.

The dealer is still exposed to the basis risk when the December contracts are bought back from the market at the expiry date of the September contract. If for example, the future was bought back at 92.73, the effective borrowing rate would be only 6.4325%, and the dealer would benefit. Of course the other possibility is that the futures contract could be trading 20 ticks more expensive, which would give a borrowing rate of 6.6325%, which is 10 basis points above the target rate. If this happened, the dealer may elect to borrow in the cash market for three months, and maintain the December futures position until the December contract expiry date, and roll over the borrowing at that time. The profit (or loss) on the December futures position will compensate for any change in three-month rates at that time.

Selected bibliography and references

Cox, J., Ingersoll, J., Ross, S., "The Relationship between Forward Prices and Futures Prices", *Journal of Financial Economics* 9, December 1981, pp. 321–346

French, K., "A Comparison of Futures and Forwards Prices", *Journal of Financial Economics* 12, November 1983, pp. 311–342

Hull, J., *Options, Futures and Other Derivatives*, 4th edition, Prentice Hall 1999

Jarrow, R., Oldfield, G., "Forward Contracts and Futures Contracts", *Journal of Financial Economics* 9, December 1981, pp. 373–382

Kolb, R., *Futures, Options and Swaps*, 3rd edition, Blackwell 2000

5 Bond Futures

A widely used risk management instrument in the debt capital markets is the government bond futures contract. This is an exchange-traded standardised contract that fixes the price today at which a specified quantity and quality of a bond will be delivered at a date during the expiry month of the futures contract. Unlike short-term interest rate futures, which only require cash settlement, bond futures require the actual physical delivery of a bond when they are settled.

In this chapter we review bond futures contracts and their use for trading and hedging purposes.

5.1 Introduction

A *futures contract* is an agreement between two counterparties that fixes the terms of an exchange that will take place between them at some future date. They are standardised agreements as opposed to OTC ones, when traded on an exchange, so they are also referred to as *exchange traded futures*. In the UK financial futures are traded on LIFFE, the London International Financial Futures Exchange which opened in 1982. LIFFE is the biggest financial futures exchange in Europe in terms of volume of contracts traded. There are four classes of contract traded on LIFFE: short-term interest rate contracts, long-term interest rate contracts (bond futures), currency contracts and stock index contracts.

We discussed short-term interest rate futures contracts, which generally trade as part of the money markets, in Chapter 4. In this section we will look at bond futures contracts, which are an important part of the bond markets; they are used for hedging and speculative purposes. Most futures contracts on exchanges around the world trade at three-month maturity intervals, with maturity dates fixed at March, June, September and December each year. This includes the contracts traded on LIFFE. Therefore at pre-set times during the year a contract for each of these months will *expire*, and a final *settlement* price is determined for it. The further out one goes the less liquid the trading is in that contract. It is normal to see liquid trading only in the *front month* contract (the current contract, so that if we are trading in April 2002 the front month is the June 2002 future), and possibly one or two of the next contracts, for most bond futures contracts. The liquidity of contracts diminishes the further one trades out in the maturity range.

When a party establishes a position in a futures contract, it can either run this position to maturity or close out the position between trade date and maturity. If a position is closed out the party will have either a profit or loss to book. If a position is held until maturity, the party who is long futures will take delivery of the underlying asset (bond) at the settlement price; the party who is short futures will deliver the underlying asset. This is referred to as *physical settlement* or sometimes, confusingly, as *cash settlement*.

There is no counterparty risk associated with trading exchange-traded futures, because of the role of the *clearing house*, such as the London Clearing House. This is the body through which contracts are settled. A clearing house acts as the buyer to all contracts sold on the exchange, and the seller to all contracts that are bought. So in the London market the LCH acts as the counterparty to all transactions, so that settlement is effectively guaranteed. The LCH requires all exchange participants to deposit *margin* with it, a cash sum that is the cost of conducting business (plus brokers' commissions). The size of the margin depends on the size of a party's net *open* position in contracts (an open position is a position in a contract that is held overnight and not closed out). There are two types of margin, *maintenance margin* and *variation margin*. Maintenance margin is the minimum level required to be held at the clearing house; the level is set by the exchange. Variation margin is the additional amount that must be deposited to cover any trading losses and as the size of the net open positions increases. Note that this is not like margin in say, a repo transaction. Margin in repo is a safeguard against a drop in value of collateral that has been supplied against a loan of cash. The margin deposited at a futures exchange clearing house acts essentially as "good faith" funds, required to provide comfort to the exchange that the futures trader is able to satisfy the obligations of the futures contract.

5.1.1 *Contract specifications*

We have noted that futures contracts traded on an exchange are standardised. This means that each contract represents exactly the same commodity, and it cannot be tailored to meet individual customer requirements. In this section we describe two very liquid and commonly traded contracts, starting with the US T-bond contract traded on the Chicago Board of Trade (CBOT). The details of this contract are given in Table 5.1.

The terms of this contract relate to a US Treasury bond with a minimum maturity of 15 years and a *notional* coupon of 8%. We introduced the concept of the notional bond in the chapter on repo markets. A futures contract specifies a notional coupon to prevent delivery and liquidity problems that would arise if there was shortage of bonds with exactly the coupon required, or if one market participant purchased a large proportion of all the bonds in issue with the required coupon. For exchange-traded futures, a short future can deliver any bond that fits the maturity criteria specified in the contract terms. Of course a long future would like to deliver a high-coupon bond with significant accrued interest,

Unit of trading	US Treasury bond with notional value of $100,000 and a coupon of 8%
Deliverable grades	US T-bonds with a minimum maturity of 15 years from first day of delivery month
Delivery months	March, June, September, December
Delivery date	Any business day during the delivery month
Last trading day	12:00 noon, seventh business day before last business day of delivery month
Quotation	Per cent of par expressed as points and thirty-seconds of a point, e.g., 108 – 16 is 108 16/32 or 108.50
Minimum price movement	1/32
Tick value	$31.25
Trading hours	07:20–14:00 (trading pit) 17:20–20:05 22:30–06:00 hours (screen trading)

Table 5.1: CBOT US T-bond futures contract. Source: CBOT

while the short future would want to deliver a low-coupon bond with low interest accrued. In fact this issue does not arise because of the way the *invoice amount* (the amount paid by the long future to purchase the bond) is calculated. The invoice amount on the expiry date is given as equation (5.1):

$$Inv_{amt} = P_{fut} \times CF + AI \tag{5.1}$$

where

Inv_{amt}	is the invoice amount;
P_{fut}	is the price of the futures contract;
CF	is the conversion factor;
AI	is the bond accrued interest.

Any bond that meets the maturity specifications of the futures contract is said to be in the *delivery basket*, the group of bonds that are eligible to be delivered into the futures contract. Every bond in the delivery basket will have its own *conversion factor*, which is used to equalise coupon and accrued interest differences of all the delivery bonds. The exchange will announce the conversion factor for each bond before trading in a contract begins; the conversion factor for a bond will change over time, but remains fixed for one individual contract. That is, if a bond has a conversion factor of 1.091252, this will remain fixed for the life of the contract. If a contract specifies a bond with a notional coupon of 7%, like the long gilt future on LIFFE, then the conversion factor will be less than 1.0 for bonds with a coupon lower than 7% and higher than 1.0 for bonds with a coupon higher than 7%. A formal definition of conversion factor is given below.

Conversion factor

The conversion factor (or price factor) gives the price of an individual cash bond such that its yield to maturity on the delivery day of the futures contract is equal to the notional coupon of the contract. The product of the conversion factor and the futures price is the forward price available in the futures market for that cash bond (plus the cost of funding, referred to as the gross basis).

Although conversion factors equalise the yield on bonds, bonds in the delivery basket will trade at different yields, and for this reason they are not "equal" at the time of delivery. Certain bonds will be cheaper than others, and one bond will be the *cheapest-to-deliver* bond. The cheapest-to-deliver bond is the one that gives the greatest return from a strategy of buying a bond and simultaneously selling the futures contract, and then closing out positions on the expiry of the contract. This so-called *cash-and-carry trading* is actively pursued by proprietary trading desks in banks. If a contract is purchased and then held to maturity the buyer will receive, via the exchange's clearing house the cheapest-to-deliver gilt. Traders sometimes try to exploit arbitrage price differentials between the future and the cheapest-to-deliver gilt, known as *basis trading*.

The mathematical calculation of the conversion factor for the gilt future is given in Appendix 5.1.

We summarise the contract specification of the long gilt futures contract traded on LIFFE in Table 5.2. There is also a medium gilt contract on LIFFE, which was introduced in 1998 (having been discontinued in the early 1990s). This trades a notional five-year gilt, with eligible gilts being those of four to seven years' maturity.

Unit of trading	UK gilt bond having a face value of £100,000, a notional coupon of 7% and a notional maturity of 10 years (changed from contract value of £50,000 from the September 1998 contract)
Deliverable grades	UK gilts with a maturity ranging from $8\frac{3}{4}$ to 13 years from the first day of the delivery month (changed from 10–15 years from the December 1998 contract)
Delivery months	March, June, September, December
Delivery date	Any business day during the delivery month
Last trading day	11:00 hours, two business days before last business day of delivery month
Quotation	Per cent of par expressed as points and hundredths of a point, for example 114.56 (changed from ticks and 1/32nds of a point, as in 117 – 17 meaning 114 17/32 or 114.53125, from the June 1998 contract)
Minimum price movement	0.01 of one point (one tick)
Tick value	£10
Trading hours	08:00–18:00 hours All trading conducted electronically on LIFFE CONNECT™ platform

Table 5.2: LIFFE long gilt future contract specification. Source: LIFFE

5.2 Futures pricing

5.2.1 *Theoretical principle*

Although it may not appear so on first trading, floor trading on a futures exchange is probably the closest one gets to an example of the economist's perfect and efficient market. The immediacy and liquidity of the market will ensure that at virtually all times the price of any futures contract reflects fair value. In essence because a futures contract represents an underlying asset, albeit a synthetic one, its price cannot differ from the actual cash market price of the asset itself. This is because the market sets futures prices such that they are arbitrage-free. We can illustrate this with a hypothetical example.

Let us say that the benchmark 10-year bond, with a coupon of 8% is trading at par. This bond is the underlying asset represented by the long bond futures contract; the front month contract expires in precisely three months. If we also say that the three-month LIBOR rate (the repo rate) is 6%, what is fair value for the front month futures contract?

For the purpose of illustration let us start by assuming the futures price to be 105. We could carry out the following arbitrage-type trade:

- buy the bond for £100;
- simultaneously sell the future at £105;
- borrow £100 for three months at the repo rate of 6%.

As this is a leveraged trade we have borrowed the funds with which to buy the bond, and the loan is fixed at three months because we will hold the position to the futures contract expiry, which is in exactly three months' time. At expiry, as we are short futures we will deliver the underlying bond to the futures clearing house and close out the loan. This strategy will result in cash flows for us as shown below.

Futures settlement cash flows
Price received for bond = 105.00
Bond accrued = 2.00 (8% coupon for three months)
Total proceeds = 107.00

Loan cash flows
Repayment of principal = 100.00
Loan interest = 1.500 (6% repo rate for three months)
Total outlay = 101.50

The trade has resulted in a profit of £5.50, and this profit is guaranteed as we have traded the two positions simultaneously and held them both to maturity. We are not affected by subsequent market movements. The trade is an example of a pure arbitrage, which is risk-free. There is no cash outflow at the start of the trade because we borrowed the funds used to buy the bond. In essence we have locked in the forward price of the bond by trading the future today, so that the final settlement price of the futures contract is irrelevant. If the situation described above were to occur in practice it would be very short-lived, precisely because arbitrageurs would buy the bond and sell the future to make this profit. This activity would force changes in the prices of both bond and future until the profit opportunity was removed.

So in our illustration the price of the future was too high (and possibly the price of the bond was too low as well) and not reflecting fair value because the price of the synthetic asset was out of line with the cash asset.

What if the price of the future was too low? Let us imagine that the futures contract is trading at 95.00. We could then carry out the following trade:

- sell the bond at £100;
- simultaneously buy the future for £95;
- lend the proceeds of the short sale (£100) for three months at 6%.

This trade has the same procedure as the first one with no initial cash outflow, except that we have to cover the short position in the repo market, through which we invest the sale proceeds at the repo rate of 6%. After three months we are delivered a bond as part of the futures settlement, and this is used to close out our short position. How has our strategy performed?

Futures settlement cash flows
Clean price of bond = 95.00
Bond accrued = 2.00
Total cash outflow = 97.00

Loan cash flows
Principal on loan maturity = 100.00
Interest from loan = 1.500
Total cash inflow = 101.500

The profit of £4.50 is again a risk-free arbitrage profit. Of course our hypothetical world has ignored considerations such as bid–offer spreads for the bond, future and repo rates, which would apply in the real world and impact on any trading strategy. Yet again however the futures price is out of line with the cash market and has provided opportunity for arbitrage profit.

Given the terms and conditions that apply in our example, there is one price for the futures contract at which no arbitrage profit opportunity is available. If we set the future price at 99.5, we would see that both trading strategies, buying the bond and selling the future or selling the bond and buying the future, yield a net cash flow of zero. There is no profit to be made from either strategy. So at 99.5 the futures price is in line with the cash market, and it will only move as the cash market price moves; any other price will result in an arbitrage profit opportunity.

5.2.2 *Arbitrage-free futures pricing*

The previous section demonstrated how we can arrive at the fair value for a bond futures contract provided we have certain market information. The market mechanism and continuous trading will ensure that the fair price *is* achieved, as arbitrage profit opportunities are eliminated. We can determine the bond future's price given:

- the coupon of the underlying bond, and its price in the cash market;
- the interest rate for borrowing or lending funds, from the trade date to the maturity date of the futures contract. This is known as the *repo* rate.

For the purpose of deriving this pricing model we can ignore bid–offer spreads and borrowing and lending spreads. We set the following:

r	is the repo rate;
rc	is the bond's running yield;
P_{bond}	is the price of the cash bond;
P_{fut}	is the price of the futures contract;
t	is the time to the expiry of the futures contract.

We can substitute these symbols into the cash flow profile for our earlier trade strategy, that of buying the bond and selling the future. This gives us:

Futures settlement cash flows
Clean price for bond $= P_{fut}$
Bond accrued $= rc \times t \times P_{bond}$
Total proceeds $= P_{fut} + (rc \times t \times P_{bond})$

Loan cash flows
Repayment of loan principal $= P_{bond}$
Loan interest $= r \times t \times P_{bond}$
Total outlay $= P_{bond} + (r \times t \times P_{bond})$

The profit from the trade would be the difference between the proceeds and outlay, which we can set as follows:

$$\text{Profit} = P_{fut} + rc \times t \times P_{bond} - P_{bond} + r \times t \times P_{bond}. \qquad (5.2)$$

We have seen how the futures price is at fair value when there is no profit to be gained from carrying out this trade, so if we set profit at zero, we obtain the following:

$$0 = P_{fut} + rc \times t \times P_{bond} - P_{bond} + r \times t \times P_{bond}.$$

Solving this expression for the futures price P_{fut} gives us:

$$P_{fut} = P_{bond} + P_{bond}t(r - rc).$$

Rearranging this we get:

$$P_{fut} = P_{bond}(1 + t(r - rc)). \tag{5.3}$$

If we repeat the procedure for the other strategy, that of selling the bond and simultaneously buying the future, and set the profit to zero, we will obtain the same equation for the futures price as given in equation (5.3).

It is the level of the repo rate in the market, compared to the running yield on the underlying bond, that sets the price for the futures contract. From the examples used at the start of this section we can see that it is the cost of funding compared to the repo rate that determines if the trade strategy results in a profit. The equation $(r - rc)$ from (5.3) is the net financing cost in the arbitrage trade, and is known as the *cost of carry*. If the running yield on the bond is higher than the funding cost (the repo rate) this is positive funding or *positive carry*. Negative funding (*negative carry*) is when the repo rate is higher than the running yield. The level of $(r - rc)$ will determine whether the futures price is trading above the cash market price or below it. If we have positive carry (when $rc > r$) then the futures price will trade below the cash market price, known as trading at a *discount*. Where $r > rc$ and we have negative carry then the futures price will be at a premium over the cash market price. If the net funding cost was zero, such that we had neither positive nor negative carry, then the futures price would be equal to the underlying bond price.

The cost of carry related to a bond futures contract is a function of the yield curve. In a positive yield curve environment the three-month repo rate is likely to be lower than the running yield on a bond so that the cost of carry is likely to be positive. As there is generally only a liquid market in long bond futures out to contracts that mature up to one year from the trade date, with a positive yield curve it would be unusual to have a short-term repo rate higher than the running yield on the long bond. So in such an environment we would have the future trading at a discount to the underlying cash bond. If there is a negative sloping yield curve the futures price will trade at a premium to the cash price. It is in circumstances of changes in the shape of the yield curve that opportunities for relative value and arbitrage trading arise, especially as the bond that is cheapest-to-deliver for the futures contract may change with large changes in the curve.

A trading strategy that involved simultaneous and opposite positions in the cheapest-to-deliver bond (CTD) and the futures contract is known as *cash-and-carry* trading or *basis trading*. However, by the law of no-arbitrage pricing, the payoff from such a trading strategy should be zero. If we set the profit from such a trading strategy as zero, we can obtain a pricing formula for the fair value of a futures contract, which summarises the discussion above, and states that the fair value futures price is a function of the cost of carry on the underlying bond. This is given as equation (5.4):

$$P_{fut} = \frac{(P_{bond} + AI_0) \times (1 + rt) - \sum_{i=1}^{N} C_i(1 + rt_{i,del}) - AI_{del}}{CF} \tag{5.4}$$

where

AI_0	is the accrued interest on the underlying bond today;
AI_{del}	is the accrued interest on the underlying bond on the expiry or delivery date (assuming the bond is delivered on the final day, which will be the case if the running yield on the bond is above the money market rate);
C_i	is the *i*th coupon;
N	is the number of coupons paid from today to the expiry or delivery date;
r	is the repo rate;
t	is the time period (in years) over which the trade takes place;
CF	is the bond conversion factor;
$t_{i,del}$	is the period from receipt of the *i*th coupon to delivery.

5.3 Hedging using bond futures

5.3.1 *Introduction*

Bond futures are used for a variety of purposes. Much of one day's trading in futures will be speculative, that is, a punt on the direction of the market. Another main use of futures is to hedge bond positions. In theory if hedging a cash bond position with a bond futures contract, if cash and futures prices move together then any loss from one position will be offset by a gain from the other. When prices move exactly in lock-step with each other, the hedge is considered perfect. In practice the price of even the cheapest-to-deliver bond (which one can view as being the bond being traded – implicitly – when one is trading the bond future) and the bond future will not move exactly in line with each other over a period of time. The difference between the cash price and the futures price is called the *basis*. The risk that the basis will change in an unpredictable way is known as *basis risk*.

Futures are a liquid and straightforward way of hedging a bond position. By hedging a bond position the trader or fund manager is hoping to balance the loss on the cash position by the profit gained from the hedge. However the hedge will not be exact for all bonds except the cheapest-to-deliver (CTD) bond, which we can assume is the futures contract underlying bond. The basis risk in a hedge position arises because the bond being hedged is not identical to the CTD bond. The basic principle is that if the trader is long (or net long, where the desk is running long and short positions in different bonds) in the cash market, an equivalent number of futures contracts will be sold to set up the hedge. If the cash position is short the trader will buy futures. The hedging requirement can arise for different reasons. A market maker will wish to hedge positions arising out of client business, when they are unsure when the resulting bond positions will be unwound. A fund manager may, for example, know that they need to realise a cash sum at a specific time in the future to meet fund liabilities, and sell bonds at that time. The market maker will want to hedge against a drop in value of positions during the time the bonds are held. The fund manager will want to hedge against a rise in interest rates between now and the bond sale date, to protect the value of the portfolio.

When putting on the hedge position the key is to trade the correct number of futures contracts. This is determined by using the *hedge ratio* of the bond and the future, which is a function of the volatilities of the two instruments. The amount of contracts to trade is calculated using the hedge ratio, which is given by:

$$\text{Hedge ratio} = \frac{\text{Volatility of bond to be hedged}}{\text{Volatility of hedging instrument}}.$$

Therefore one needs to use the volatility values of each instrument. We can see from the calculation that if the bond is more volatile than the hedging instrument, then a greater amount of the hedging instrument will be required. Let us now look in greater detail at the hedge ratio.

There are different methods available to calculate hedge ratios. The most common ones are the conversion factor method, which can be used for deliverable bonds (also known as the *price factor* method) and the modified duration method (also known as the basis point value method).

Where a hedge is put on against a bond that is in the futures delivery basket it is common for the conversion factor to be used to calculate the hedge ratio. A conversion factor hedge ratio is more useful as it is transparent and remains constant, irrespective of any changes in the price of the cash bond or the futures contract. The number of futures contracts required to hedge a deliverable bond using the conversion factor hedge ratio is determined using equation (5.5):

$$\text{Number of contracts} = \frac{M_{bond} \times CF}{M_{fut}} \tag{5.5}$$

where M is the nominal value of the bond or futures contract.

The conversion factor method may only be used for bonds in the delivery basket. It is important to ensure that this method is only used for one bond. It is an erroneous procedure to use the ratio of conversion factors of two different bonds when calculating a hedge ratio. This will be considered again later.

Unlike the conversion factor method, the modified duration hedge ratio may be used for all bonds, both deliverable and non-deliverable. In calculating this hedge ratio the modified duration is multiplied by the dirty price of the cash bond to obtain the *basis point value* (BPV). The BPV represents the actual impact of a change in the yield on the price of a specific bond. The BPV allows the trader to calculate the hedge ratio to reflect the different price sensitivity of the chosen bond (compared to the CTD bond) to interest rate movements. The hedge ratio calculated using BPVs must be constantly updated, because it will change if the price of the bond and/or the futures contract changes. This may necessitate periodic adjustments to the number of lots used in the hedge.

The number of futures contracts required to hedge a bond using the BPV method is calculated using the following:

$$\text{Number of contracts} = \frac{M_{bond}}{M_{fut}} \times \frac{BPV_{bond}}{BPV_{fut}} \tag{5.6}$$

where the BPV of a futures contract is defined with respect to the BPV of its CTD bond, as given by equation (5.7):

$$BPV_{fut} = \frac{BPV_{CTDbond}}{CF_{CTDbond}}. \tag{5.7}$$

The simplest hedge procedure to undertake is one for a position consisting of only one bond, the cheapest-to-deliver bond. The relationship between the futures price and the price of the CTD given by equation (5.4) indicates that the price of the future will move for moves in the price of the CTD bond, so therefore we may set:

$$\Delta P_{fut} \simeq \frac{\Delta P_{bond}}{CF} \tag{5.8}$$

where *CF* is the CTD conversion factor.

The price of the futures contract, over time, does not move tick-for-tick with the CTD bond (although it may on an intra-day basis) but rather by the amount of the change divided by the conversion factor. It is apparent therefore that to hedge a position in the CTD bond we must hold the number of futures contracts equivalent to the value of bonds held multiplied by the conversion factor. Obviously if a conversion factor is less than one, the number of futures contracts will be less than the equivalent nominal value of the cash position; the opposite is true for bonds that have a conversion factor greater than one. However the hedge is not as simple as dividing the nominal value of the bond position by the nominal value represented by one futures contract.

To measure the effectiveness of the hedge position, it is necessary to compare the performance of the futures position with that of the cash bond position, and to see how much the hedge instrument mirrored the performance of the cash instrument. A simple calculation is made to measure the effectiveness of the hedge, given by equation (5.9), which is the percentage value of the hedge effectiveness:

$$\text{Hedge effectiveness} = -\left(\frac{Fut\ p/l}{Bond\ p/l}\right) \times 100. \tag{5.9}$$

5.3.2 *Hedging a bond portfolio*

The principles established above may be applied when hedging a portfolio containing a number of bonds. It is more realistic to consider a portfolio holding not just bonds that are outside the delivery basket, but are also not government bonds. In this case we need to calculate the number of futures contracts to put on as a hedge based on the volatility of each bond in the portfolio compared to the volatility of the CTD bond. Note that in practice, there is usually more than one futures contract that may be used as the hedge instrument. For example in the sterling market it would be more sensible to use LIFFE's medium gilt contract, whose underlying bond has a notional maturity of four to seven years, if hedging a portfolio of short- to medium-dated bonds. However for the purposes of illustration we will assume that only one contract, the long bond, is available.

To calculate the number of futures contracts required to hold as a hedge against any specific bond, we use equation (5.10).

$$\text{Hedge} = \frac{M_{bond}}{M_{fut}} \times Vol_{bond/CTD} \times Vol_{CTD/fut} \tag{5.10}$$

where

M	is the nominal value of the bond or future;
$Vol_{bond/CTD}$	is the relative volatility of the bond being hedged compared to that of the CTD bond;
$Vol_{CTD/fut}$	is the relative volatility of the CTD bond compared to that of the future.

It is not necessarily straightforward to determine the relative volatility of a bond vis-à-vis the CTD bond. If the bond being hedged is a government bond, we can calculate the relative volatility using the two bonds' modified duration. This is because the yields of both may be safely assumed to be strongly positively correlated. If however the bond being hedged is a corporate bond and/or non-vanilla bond, we must obtain the relative volatility using regression analysis, as the yields between the two bonds may not be strongly positively correlated. This is apparent when one remembers that the yield spread of corporate bonds over government bonds is not constant, and will fluctuate with changes in government bond yields. To use regression analysis to determine relative volatilities, historical price data on the bond is required; the daily price moves in the target bond and the CTD bond are then analysed to assess the slope of the regression line. In this section we will restrict the discussion to a portfolio of government bonds.

If we are hedging a portfolio of government bonds we can use equation (5.11) to determine relative volatility values, based on the modified duration of each of the bonds in the portfolio.

$$Vol_{bond/CTD} = \frac{\Delta P_{bond}}{\Delta P_{CTD}} = \frac{MD_{bond} \times P_{bond}}{MD_{CTD} \times P_{CTD}} \tag{5.11}$$

where *MD* is the modified duration of the bond being hedged or the CTD bond, as appropriate.[1]

Once we have calculated the relative volatility of the bond being hedged, equation (5.12) (obtained from (5.8) and (5.11)) tells us that the relative volatility of the CTD bond to that of the futures contract is approximately the same as its conversion factor. We are then in a position to calculate the futures hedge for each bond in a portfolio.

$$Vol_{CTD/fut} = \frac{\Delta P_{CTD}}{\Delta P_{fut}} \approx CF_{CTD}. \tag{5.12}$$

Table 5.3 shows a portfolio of five UK gilts on 20 October 1999. The nominal value of the bonds in the portfolio is £200 million, and the bonds have a market value excluding accrued interest of £206.84 million. Only one of the bonds is a deliverable bond, the $5\frac{3}{4}$% 2009 gilt which is in fact the CTD bond. For the Dec '99 futures contract the bond had a conversion factor of 0.9124950. The fact that this bond is the CTD explains why it has a relative volatility of 1. We calculate the number of futures contracts required to hedge each position, using

1 In certain textbooks and practitioner research documents, it is suggested that the ratio of the conversion factors of the bond being hedged (if it is in the delivery basket) and the CTD bond can be used to determine the relative volatility of the target bond. This is a specious argument. The conversion factor of a deliverable bond is the price factor that will set the yield of the bond equal to the notional coupon of the futures contract on the delivery date, and it is a function mainly of the coupon of the deliverable bond. The price volatility of a bond on the other hand, is a measure of its modified duration, which is a function of the bond's duration (that is, the weighted average term to maturity). Therefore using conversion factors to measure volatility levels will produce erroneous results. It is important not to misuse conversion factors when arranging hedge ratios.

CTD		5.75% 2009	Modified duration			7.234565567	
Conversion factor		**0.9124950**	**Price**			**99.84**	
Bond	**Nominal amount (£m)**	**Price**	**Yield %**	**Duration**	**Modified duration**	**Relative volatility**	**Number of contracts**
UKT 8% 2000	12	102.17	5.972	1.072	1.011587967	0.143090242	15.67
UKT 7% 2002	5	101.50	6.367	2.388	2.245057208	0.315483336	14.39
UKT 5% 2004	38	94.74	6.327	4.104	3.859791022	0.50626761	175.55
UKT 5.75% 2009	100	99.84	5.770	7.652	7.234565567	1.00	912.50
UKT 6% 2028	45	119.25	4.770	15.031	14.34666412	2.368603078	972.60
Total	200						2090.71

Table 5.3: Bond futures hedge for hypothetical gilt portfolio, 20 October 1999

the equations listed above. For example, the hedge requirement for the position in the 7% 2002 gilt was calculated as follows:

$$\frac{5{,}000{,}000}{100{,}000} \times \frac{2.245 \times 101.50}{7.235 \times 99.84} \times 0.9124950 = 14.39.$$

The volatility of all the bonds is calculated relative to the CTD bond, and the number of futures contracts determined using the conversion factor for the CTD bond. The bond with the highest volatility is not surprisingly the 6% 2028, which has the longest maturity of all the bonds and hence the highest modified duration. We note from Table 5.3 that the portfolio requires a hedge position of 2091 futures contracts. This illustrates how a "rough-and-ready" estimate of the hedging requirement, based on nominal values, would be insufficient as that would suggest a hedge position of only 2000 contracts.

The effectiveness of the hedge must be monitored over time. No hedge will be completely perfect however, and the calculation illustrated above, as it uses modified duration value, does not take into account the convexity effect of the bonds. The reason why a futures hedge will not be perfect is because in practice, the price of the futures contract will not move tick-for-tick with the CTD bond, at least not over a period of time. This is the basis risk that is inherent in hedging cash bonds with futures. In addition, the calculation of the hedge is only completely accurate for a parallel shift in yields, as it is based on modified duration, so as the yield curve changes around pivots, the hedge will move out of line. Finally, the long gilt future is not the appropriate contract to use to hedge three of the bonds in the portfolio, or over 25% of the portfolio by nominal value. This is because these bonds are short- or medium-dated, and so their price movements will not track the futures price as closely as longer-dated bonds. In this case, the more appropriate futures contract to use would have been the medium gilt contract, or (for the first bond, the 8% 2000) a strip of short sterling contracts. Using shorter-dated instruments would reduce some of the basis risk contained in the portfolio hedge.

5.4 The margin process

Institutions buying and selling futures on an exchange deal with only one counterparty at all times, the exchange clearing house. The clearing house is responsible for the settlement of all contracts, including managing the delivery process. A central clearing mechanism eliminates counterparty risk for anyone dealing on the exchange, because the clearing house guarantees the settlement of all transactions. The clearing house may be owned by the exchange itself, such as the one associated with the Chicago Mercantile Exchange (the CME Clearinghouse) or it may be a separate entity, such as the London Clearing House, which settles transactions on LIFFE. The LCH is also involved in running clearing systems for swaps and repo products in certain currencies.

One of the key benefits to the market of the clearing house mechanism is that counterparty risk, as it is transferred to the clearing house, is virtually eliminated. The mechanism that enables the clearing house to accept the counterparty risk is the *margining* process that is employed at all futures exchanges. A bank or local trader must deposit margin before commencing dealing on the exchange; each day a further amount must be deposited or will be returned, depending on the results of the day's trading activity.

The exchange will specify the level of margin that must be deposited for each type of futures contract that a bank wishes to deal in. The *initial margin* will be a fixed sum per lot, so for example if the margin was £1000 per lot an opening position of 100 lots would require margin of £100,000. Once initial margin has been deposited, there is a mark-to-market of all positions at the close of business; exchange-traded instruments are the most transparent products in the market, and the closing price is not only known to everyone, it is also indisputable. The closing price is also known as the *settlement price*. Any losses suffered by a trading counterparty, whether closed out or run overnight, are entered as a debit on the party's account and must be paid the next day. Trading profits are credited and may be withdrawn from the margin account the next day. This daily process is known as *variation margining*. Thus the margin account is updated on a daily basis and the maximum loss that must be made up on any morning is the maximum price movement that occurred the previous day. It is a serious issue if a trading party is unable to meet a margin call. In such a case, the exchange will order it to cease trading, and will also liquidate all its open positions; any losses will be met out of the firm's margin account. If the level of funds in the margin account is insufficient, the losses will be made good from funds paid out of a general fund run by the clearing house, which is maintained by all members of the exchange.

Payment of margin is made by electronic funds transfer between the trading party's bank account and the clearing house. Initial margin is usually paid in cash, although clearing houses will also accept high-quality securities such as T-bills or certain government bonds, to the value of the margin required. Variation margin is always cash. The advantage of depositing securities rather than cash is that the depositing firm earns interest on its margin. This is not available on a cash margin, and the interest foregone on a cash margin is effectively the cost of trading futures on the exchange. However if securities are used, there is effectively no cost associated with trading on the exchange (we ignore of course infrastructure costs and staff salaries).

The daily settlement of exchange-traded futures contracts, as opposed to when the contract expires or the position is closed out, is the main reason why futures prices are not equal to forward prices for long-dated instruments.

Appendix 5.1: The conversion factor for the long gilt future

Here we describe the process used for the calculation of the *conversion factor* or *price factor* for deliverable bonds of the long gilt contract. The contract specifies a bond of maturity $8\frac{3}{4}$–13 years and a notional coupon of 7%. For each bond that is eligible to be in the delivery basket, the conversion factor is given by the following equation: $P(7)/100$ where the numerator $P(7)$ is equal to the price per £100 nominal of the deliverable gilt at which it has a gross redemption yield of 7%, calculated as at the first day of the delivery month, less the accrued interest on the bond on that day. This calculation uses the formulae given in equation (A5.1.1) and the equation used to calculate accrued interest.

The numerator $P(7)$ is given by equation (A5.1.1):

$$P(7) = \frac{1}{1.035^{t/s}}\left(c_1 + \frac{c_2}{1.035} + \frac{C}{0.07}\left(\frac{1}{1.035} - \frac{1}{1.035^n}\right) + \frac{100}{1.035^n}\right) - AI \tag{A5.1.1}$$

where

- c_1 is the cash flow due on the following quasi-coupon date, per £100 nominal of the gilt. c_1 will be zero if the first day of the delivery month occurs in the ex-dividend period or if the gilt has a long first coupon period and the first day of the delivery month occurs in the first full coupon period. c_1 will be less than c_2 if the first day of the delivery month falls in a short first coupon period. c_1 will be greater than c_2 if the first day of the delivery month falls in a long first coupon period and the first day of the delivery month occurs in the second full coupon period;
- c_2 is the cash flow due on the next but one quasi-coupon date, per £100 nominal of the gilt. c_1 will be greater than c_2 if the first day of the delivery month falls in a long first coupon period and in the first full coupon period. In all other cases, $c_2 = C/2$;
- C is the annual coupon of the gilt, per £100 nominal;
- t is the number of calendar days from and including the first day of the delivery month up to but excluding the next quasi-coupon date;
- s is the number of calendar days in the full coupon period in which the first day of the delivery month occurs;
- n is the number of full coupon periods between the following quasi-coupon date and the redemption date;
- AI is the accrued interest per £100 nominal of the gilt.

The accrued interest used in the formula above is given according to the following procedures.

If the first day of the delivery month occurs in a standard coupon period, and the first day of the delivery month occurs on or before the ex-dividend date, then:

$$AI = \frac{t}{s} \times \frac{C}{2}. \tag{A5.1.2}$$

If the first day of the delivery month occurs in a standard coupon period, and the first day of the delivery month occurs after the ex-dividend date, then:

$$AI = \left(\frac{t}{s} - 1\right) \times \frac{C}{2} \tag{A5.1.3}$$

where

- t is the number of calendar days from and including the last coupon date up to but excluding the first day of the delivery month;
- s is the number of calendar days in the full coupon period in which the first day of the delivery month occurs.

If the first day of the delivery month occurs in a short first coupon period, and the first day of the delivery month occurs on or before the ex-dividend date, then:

$$AI = \frac{t^*}{s} \times \frac{C}{2}. \tag{A5.1.4}$$

If the first day of the delivery month occurs in a short first coupon period, and the first day of the delivery month occurs after the ex-dividend date, then:

$$AI = \left(\frac{t^* - n}{s}\right) \times \frac{C}{2} \tag{A5.1.5}$$

where

- t^* is the number of calendar days from and including the issue date up to but excluding the first day of the delivery month;
- n is the number of calendar days from and including the issue date up to but excluding the next quasi-coupon date.

If the first day of the delivery month occurs in a long first coupon period, and during the first full coupon period, then:

$$AI = \frac{u}{s_1} \times \frac{C}{2}. \tag{A5.1.6}$$

If the first day of the delivery month occurs in a long first coupon period, and during the second full coupon period and on or before the ex-dividend date, then:

$$AI = \left(\frac{p_1}{s_1} + \frac{p_2}{s_2}\right) \times \frac{C}{2}. \tag{A5.1.7}$$

If the first day of the delivery month occurs in a long first coupon period, and during the second full coupon period and after the ex-dividend date, then:

$$AI = \left(\frac{p_2}{s_2} - 1\right) \times \frac{C}{2} \tag{A5.1.8}$$

where

- u is the number of calendar days from and including the issue date up to but excluding the first day of the delivery month;
- s_1 is the number of calendar days in the full coupon period in which the issue date occurs;
- s_2 is the number of days in the next full coupon period after the full coupon period in which the issue date occurs;
- p_1 is the number of calendar days from and including the issue date up to but excluding the next quasi-coupon date;

p_2 is the number of calendar days from and including the quasi-coupon date after the issue date up to but excluding the first day of the delivery month which falls in the next full coupon period after the full coupon period in which the issue date occurs.

Selected bibliography and references

Fabozzi, F., *Fixed Income Mathematics*, Probus 1993
Galitz, L., *Financial Engineering*, FT Pitman 1995
Kolb, R., *Futures, Options and Swaps*, Blackwell 2000

6 Swaps

Swaps are off-balance sheet instruments involving combinations of two or more basic building blocks. Most swaps currently traded in the market involve combinations of cash market securities, for example a fixed interest rate security combined with a floating interest rate security, possibly also combined with a currency transaction. However the market has also seen swaps that involve a futures or forward component, as well as swaps that involve an option component. The market in say, dollar, euro and sterling interest rate swaps is very large and very liquid. The main types of swap are interest rate swaps, asset swaps, basis swaps, fixed-rate currency swaps and currency coupon swaps. The market for swaps is organised by the International Swap Dealers Association.

Swaps are now one of the most important and useful instruments in the debt capital markets. They are used by a wide range of institutions, including banks, mortgage banks and building societies, corporates and local authorities. The demand for them has grown as the continuing uncertainty and volatility of interest rates and exchange rates has made it ever more important to hedge their exposures. As the market has matured the instrument has gained wider acceptance, and is regarded as a "plain vanilla" product in the debt capital markets. Virtually all commercial and investment banks will quote swap prices for their customers, and as they are OTC instruments, dealt over the telephone, it is possible for banks to tailor swaps to match the precise requirements of individual customers. There is also a close relationship between the bond market and the swap market, and corporate finance teams and underwriting banks keep a close eye on the government yield curve and the swap yield curve, looking out for possibilities regarding new issue of debt.

It is not proposed to cover the historical evolution of the swaps market, which is abundantly covered in existing literature, nor the myriad of swap products which can be traded today. In this chapter we review the use of interest rate swaps from the point of view of the bond market participant; this includes pricing and valuation and its use as a hedging tool. The bibliography lists further reading on important topics such as pricing, valuation and credit risk.[1]

6.1 Interest rate swaps

6.1.1 *Introduction*

Interest rate swaps are the most important type of swap in terms of volume of transactions. They are used to manage and hedge interest rate risk and exposure, while market makers will also take positions in swaps that reflect their view on the direction of interest rates. An interest rate swap is an agreement between two counterparties to make periodic interest payments to one another during the life of the swap, on a predetermined set of dates, based on a *notional* principal amount. One party is the fixed-rate payer, and this rate is

[1] This chapter is abridged from Choudhry (2001), Chapter 11.

agreed at the time of trade of the swap; the other party is the floating-rate payer, the floating rate being determined during the life of the swap by reference to a specific market index. The principal or notional amount is never physically exchanged, hence the term "off-balance sheet", but is used merely to calculate the interest payments. The fixed-rate payer receives floating-rate interest and is said to be "long" or to have "bought" the swap. The long side has conceptually purchased a floating-rate note (because it receives floating-rate interest) and issued a fixed coupon bond (because it pays out fixed interest at intervals), that is, it has in principle borrowed funds. The floating-rate payer is said to be "short" or to have "sold" the swap. The short side has conceptually purchased a coupon bond (because it receives fixed-rate interest) and issued a floating-rate note (because it pays floating-rate interest). So an interest rate swap is:

- an agreement between two parties
- to exchange a stream of cash flows
- calculated as a percentage of a *notional* sum
- and calculated on different interest bases.

For example in a trade between Bank A and Bank B, Bank A may agree to pay fixed semi-annual coupons of 10% on a notional principal sum of £1 million, in return for receiving from Bank B the prevailing six-month sterling LIBOR rate on the same amount. The known cash flow is the fixed payment of £50,000 every six months by Bank A to Bank B.

Interest rate swaps trade in a secondary market so their value moves in line with market interest rates, in exactly the same way as bonds. If a five-year interest rate swap is transacted today at a rate of 5%, and five-year interest rates subsequently fall to 4.75%, the swap will have decreased in value to the fixed-rate payer, and correspondingly increased in value to the floating-rate payer, who has now seen the level of interest payments fall. The opposite would be true if five-year rates moved to 5.25%. Why is this? Consider the fixed-rate payer in an IR swap to be a borrower of funds; if she fixes the interest rate payable on a loan for five years, and then this interest rate decreases shortly afterwards, is she better off? No, because she is now paying above the market rate for the funds borrowed. For this reason a swap contract decreases in value to the fixed-rate payer if there is a fall in rates. Equally a floating-rate payer gains if there is a fall in rates, as he can take advantage of the new rates and pay a lower level of interest; hence the value of a swap increases to the floating-rate payer if there is a fall in rates.

A bank swaps desk will have an overall net interest rate position arising from all the swaps it has traded that are currently on the book. This position is an interest rate exposure at all points along the term structure, out to the maturity of the longest-dated swap. At the close of business each day all the swaps on the book will be *marked-to-market* at the interest rate quote for that day.

A swap can be viewed in two ways, either as a bundle of forward or futures contracts, or as a bundle of cash flows arising from the "sale" and "purchase" of cash market instruments. If we imagine a strip of futures contracts, maturing every three or six months out to three years, we can see how this is conceptually similar to a three-year interest rate swap. However in the author's view it is better to visualise a swap as being a bundle of cash flows arising from cash instruments.

Cashflows resulting from long position in FRN and short position in gilt			
Period (6 mo)	**FRN**	**Gilt**	**Net cash flow**
0	$-$£100 m	+£100 m	£0
1	+(LIBOR × 100)/2	−3	+(LIBOR × 100)/2 − 3.0
2	+(LIBOR × 100)/2	−3	+(LIBOR × 100)/2 − 3.0
3	+(LIBOR × 100)/2	−3	+(LIBOR × 100)/2 − 3.0
4	+(LIBOR × 100)/2	−3	+(LIBOR × 100)/2 − 3.0
5	+(LIBOR × 100)/2	−3	+(LIBOR × 100)/2 − 3.0
6	+[(LIBOR × 100)/2] + 100	−103	+(LIBOR × 100)/2 − 3.0

Table 6.1: Three-year cash flows. The LIBOR rate is the six-month rate prevailing at the time of the setting, for instance the LIBOR rate at period 4 will be the rate actually prevailing at period 4

Let us imagine we have only two positions on our book:

- a long position in £100 million of a three-year FRN that pays six-month LIBOR semi-annually, and is trading at par;
- a short position in £100 million of a three-year gilt with coupon of 6% that is also trading at par.

Being short a bond is the equivalent to being a borrower of funds. Assuming this position is kept to maturity, the resulting cash flows are shown in Table 6.1.

There is no net outflow or inflow at the start of these trades, as the £100 million purchase of the FRN is netted with receipt of £100 million from the sale of the gilt. The resulting cash flows over the three-year period are shown in the last column of Table 6.1. This net position is exactly the same as that of a fixed-rate payer in an IR swap. As we had at the start of the trade, there is no cash inflow or outflow on maturity. For a floating-rate payer, the cash flow would mirror exactly a long position in a fixed-rate bond and a short position in an FRN. Therefore the fixed-rate payer in a swap is said to be short in the bond market, that is a borrower of funds; the floating-rate payer in a swap is said to be long in the bond market.

6.1.2 *Market terminology*

Virtually all swaps are traded under the legal terms and conditions stipulated in the ISDA standard documentation. The trade date for a swap is, not surprisingly, the date on which the swap is transacted. The terms of the trade include the fixed interest rate, the maturity and notional amount of the swap, and the payment bases of both legs of the swap. The date from which floating interest payments are determined is the *setting date*, which may also be the trade date. Most swaps fix the floating-rate payments to LIBOR, although other reference rates that are used include the US prime rate, euribor, the Treasury bill rate and the commercial paper rate. In the same way as for FRA and eurocurrency deposits, the rate is fixed two business days before the interest period begins. The second (and subsequent) setting date will be two business days before the beginning of the second (and subsequent) swap periods. The *effective date* is the date from which interest on the swap is calculated, and this is typically two business days after the trade date. In a *forward*-start swap the effective date will be at some point in the future, specified in the swap terms. The floating interest rate for each period is fixed at the start of the period, so that the interest payment amount is known in advance by both parties (the fixed rate is known of course, throughout the swap by both parties).

Although for the purposes of explaining swap structures, both parties are said to pay interest payments (and receive them), in practice only the net difference between both payments changes hands at the end of each interest period. This eases the administration associated with swaps and reduces the number of cash flows for each swap. The counterparty that is the net payer at the end of each period will make a payment to the counterparty. The first payment date will occur at the end of the first interest period, and subsequent payment dates will fall at the end of successive interest periods. The final payment date falls on the maturity date of the swap. The calculation of interest is given by equation (6.1).

$$I = M \times r \times \frac{n}{B} \tag{6.1}$$

where

I	is the interest amount;
M	is the nominal amount of the swap;
B	is the interest day-base for the swap.

Dollar and euro-denominated swaps use an actual/360 day count, similar to other money market instruments in those currencies, while sterling swaps use an actual/365 day count basis.

The cash flows resulting from a vanilla interest rate swap are illustrated in Figure 6.1, using the normal convention where cash inflows are shown as an arrow pointing up, while cash outflows are shown as an arrow pointing down. The counterparties in a swap transaction only pay across net cash flows however, so at each interest payment date only one actual cash transfer will be made, by the net payer. This is shown as Figure 6.1(c).

6.1.3 *Swap spreads and the swap yield curve*

In the market banks will quote two-way swap rates, on screens and on the telephone or via a dealing system such as Reuters. Brokers will also be active in relaying prices in the market. The convention in the market is for the swap market maker to set the floating leg at LIBOR and then quote the fixed rate that is payable for that maturity. So for a five-year swap a bank's swap desk might be willing to quote the following:

Floating-rate payer:	pay 6 mo LIBOR
	receive fixed rate of 5.19%
Fixed-rate payer:	pay fixed rate of 5.25%
	receive 6 mo LIBOR

In this case the bank is quoting an offer rate of 5.25%, which the fixed-rate payer will pay, in return for receiving LIBOR flat. The bid price quote is 5.19% which is what a floating-rate payer will receive fixed. The bid–offer spread in this case is therefore 6 basis points. The fixed-rate quotes are always at a spread above the government bond yield curve. Let us assume that the five-year gilt is yielding 4.88%; in this case then the five-year swap bid rate is 31 basis points above this yield. So the bank's swap trader could quote the swap rates as

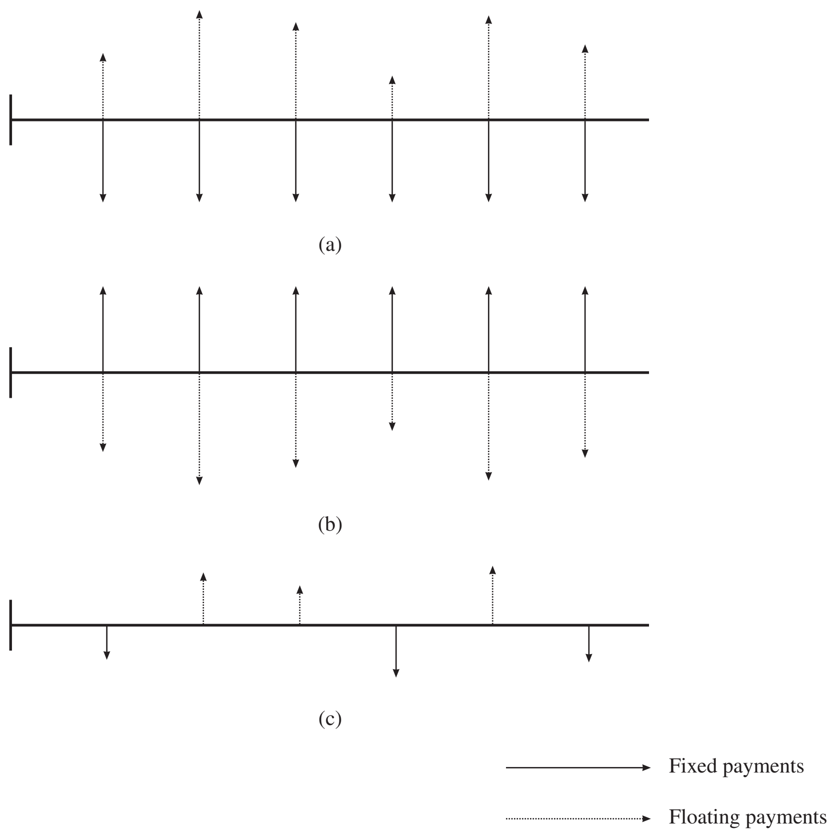

Figure 6.1: Cash flows for typical interest rate swap. (a) Cash flows for fixed-rate payer. (b) Cash flows for floating-rate payer. (c) Net cash flows

a spread above the benchmark bond yield curve, say $37 - 31$, which is her swap spread quote. This means that the bank is happy to enter into a swap paying fixed 31 basis points above the benchmark yield and receiving LIBOR, and receiving fixed 37 basis points above the yield curve and paying LIBOR. The bank's screen on say, Bloomberg or Reuters might look something like Table 6.2, which quotes the swap rates as well as the current spread over the government bond benchmark.

1 year	4.50	4.45	+17
2 years	4.69	4.62	+25
3 years	4.88	4.80	+23
4 years	5.15	5.05	+29
5 years	5.25	5.19	+31
10 years	5.50	5.40	+35

Table 6.2: Swap quotes

The swap spread is a function of the same factors that influence the spread over government bonds for other instruments. For shorter duration swaps, say up to three years, there are other yield curves that can be used in comparison, such as the cash market curve or a curve derived from futures prices. For longer-dated swaps the spread is determined mainly by the credit spreads that prevail in the corporate bond market. Because a swap is viewed as a package of long and short positions in fixed- and floating-rate bonds, it is the credit spreads in these two markets that will determine the swap spread. This is logical; essentially it is the premium for greater credit risk involved in lending to corporates that dictates that a swap rate will be higher than the same maturity government bond yield. Technical factors will be responsible for day-to-day fluctuations in swap rates, such as the supply of corporate bonds and the level of demand for swaps, plus the cost to swap traders of hedging their swap positions.

We can summarise by saying that swap spreads over government bonds reflect the supply and demand conditions of both swaps and government bonds, as well as the market's view on the credit quality of swap counterparties. There is considerable information content in the swap yield curve, much like that in the government bond yield curve. During times of credit concerns in the market, such as the corrections in Asian and Latin American markets in summer 1998, and the possibility of default by the Russian government regarding its long-dated US dollar bonds, the swap spread will increase, more so at higher maturities. To illustrate this let us consider the sterling swap spread in 1998/99. The UK swap spread widened from the second half of 1998 onwards, a reaction to bond market volatility around the world. At such times investors embark on a "flight to quality" that results in yield spreads widening. In the swap market, the spread between two-year and ten-year swaps also increased, reflecting market concern with credit and counterparty risk. The spreads narrowed in the first quarter 1999, as credit concerns brought about by market corrections in 1998 declined. The change in swap spreads is shown in Figure 6.2.

Figure 6.2: Sterling 2-year and 10-year swap spreads 1998/99

6.2 Generic swap valuation

Banks generally use *par* swap or *zero-coupon* swap pricing. We will look at this method in Section 6.2.2; first however we will introduce an intuitive swap valuation method.

6.2.1 *Generic swap pricing*

Assume we have a vanilla interest rate swap with a notional principal of N that pays n payments during its life, to a maturity date of T. The date of each payment is on t_i with $i = 1, \ldots, n$. The present value today of a future payment is denoted by $PV(0, t)$. If the swap rate is r, the value of the fixed-leg payments is given by equation (6.2) below:

$$PV_{fixed} = N \sum_{i=1}^{n} PV(0, t_i) \times \left[r \times \left(\frac{t_i - t_{i-1}}{B} \right) \right] \tag{6.2}$$

where B is the money market day base. The term $t_i - t_{i-1}$ is simply the number of days between the ith and the $i-1$th payments.

The value of the floating-leg payments at the date t_1 for an existing swap is given by:

$$PV_{float} = N \times \left[rl \times \frac{t_1}{B} \right] + N - [N \times PV(t_1, t_n)] \tag{6.3}$$

where rl is the LIBOR rate that has been set for the next interest payment. We set the present value of the floating-rate payment at time 0 as follows:

$$PV(0, t_1) = \frac{1}{1 + rl(t_1)\left(\dfrac{t_1}{B}\right)}. \tag{6.4}$$

For a new swap the value of the floating payments is given by:

$$PV_{float} = N \left[rl \times \frac{t_1}{B} + 1 \right] \times PV(0, t_1) - PV(0, t_n). \tag{6.5}$$

The swap valuation is then given by $PV_{fixed} - PV_{float}$. The swap rate quoted by a market making bank is that which sets $PV_{fixed} = PV_{float}$ and is known as the par or zero-coupon swap rate. We consider this next.

6.2.2 *Zero-coupon swap pricing*

So far we have discussed how vanilla swap prices are often quoted as a spread over the benchmark government bond yield in that currency, and how this swap spread is mainly a function of the credit spread required by the market over the government (risk-free) rate. This method is convenient and also logical because banks use government bonds as the main instrument when hedging their swap books. However because much bank swap trading is now conducted in non-standard, tailor-made swaps, this method can sometimes be unwieldy as each swap needs to have its spread calculated to suit its particular characteristics. Therefore banks use a standard pricing method for all swaps known as *zero-coupon* swap pricing.

In Chapter 2 we referred to zero-coupon bonds and zero-coupon interest rates. Zero-coupon rates, or *spot rates* are true interest rates for their particular term to maturity. In zero-coupon swap pricing, a bank will view all swaps, even the most complex, as a series of cash flows. The zero-coupon rates that apply now for each of the cash flows in a swap can be used to value these cash flows. Therefore to value and price a swap, each of the swap's cash flows are present valued using known spot rates; the sum of these present values is the value of the swap.

In a swap the fixed-rate payments are known in advance and so it is straightforward to present value them. The present value of the floating-rate payments is usually estimated in two stages. First the implied forward rates can be calculated using equation (6.6). We are quite familiar with this relationship from our reading of the earlier chapter.

$$rf_i = \left(\frac{df_i}{df_{i+1}} - 1\right)N \tag{6.6}$$

where

rf_i	is the one-period forward rate starting at time i;
df_i	is the discount factor for the maturity period i;
df_{i+1}	is the discount factor for the period $i+1$;
N	is the number of times per year that coupons are paid.

By definition the floating-payment interest rates are not known in advance, so the swap bank will predict what these will be, using the forward rates applicable to each payment date. The forward rates are those that are currently implied from spot rates. Once the size of the floating-rate payments has been estimated, these can also be valued by using the spot rates. The total value of the fixed and floating legs is the sum of all the present values, so the value of the total swap is the net of the present values of the fixed and floating legs.

While the term *zero-coupon* refers to an interest rate that applies to a discount instrument that pays no coupon and has one cash flow (at maturity), it is not necessary to have a functioning zero-coupon bond market in order to construct a zero-coupon yield curve. In practice most financial pricing models use a combination of the following instruments to construct zero-coupon yield curves:

- money market deposits;
- interest rate futures;
- FRAs;
- government bonds.

Frequently an overlap in the maturity period of all instruments is used; FRA rates are usually calculated from interest rate futures so it is necessary to use only one of either FRA or futures rates.

Once a zero-coupon yield curve (*term structure*) is derived, this may be used to value a future cash flow maturing at any time along the term structure. This includes swaps: to price an interest rate swap, we calculate the present value of each of the cash flows using the zero-coupon rates and then sum all the cash flows. As we noted above, while the fixed-rate payments are known in advance, the floating-rate payments must be estimated, using the forward rates implied by the zero-coupon yield curve. The net present value of the swap is the net difference between the present values of the fixed- and floating-rate legs.

6.2.3 *Calculating the forward rate from the spot rate discount factors*

Remember that one way to view a swap is as a long position in a fixed-coupon bond that was funded at LIBOR, or against a short position in a floating-rate bond. The cash flows from such an arrangement would be paying floating rate and receiving fixed rate. In the former arrangement, where a long position in a fixed-rate bond is funded with a floating-rate loan, the cash flows from the principals will cancel out, as they are equal and opposite (assuming the price of the bond on purchase was par), leaving a collection of cash flows that mirror an interest rate swap that pays floating and receives fixed. Therefore, as the fixed rate on an interest rate swap is the same as the coupon (and yield) on a bond priced at par, calculating the fixed rate on an interest rate swap is the same as calculating the coupon for a bond that we wish to issue at par.

The price of a bond paying semi-annual coupons is given by equation (6.7), which may be rearranged for the coupon rate r to provide an equation that enables us to determine the par yield, and hence the swap rate r, given by equation (6.8):

$$P = \frac{r_n}{2} df_1 + \frac{r_n}{2} df_2 + \cdots\cdots + \frac{r_n}{2} df_n + M df_n \tag{6.7}$$

where r_n is the coupon on an n-period bond with n coupons and M is the maturity payment. It can be shown then that

$$\begin{aligned} r_n &= \frac{1 - df_n}{\dfrac{df_1}{2} + \dfrac{df_2}{2} + \cdots\cdots + \dfrac{df_n}{2}} \\ &= \frac{1 - df_n}{\sum\limits_{i=1}^{n} \dfrac{df_i}{2}}. \end{aligned} \tag{6.8}$$

For annual coupon bonds there is no denominator for the discount factor, while for bonds paying coupons on a frequency of N we replace the denominator 2 with N.[2] Equation (6.8) may be rearranged again, using F for the coupon frequency, to obtain an equation which may be used to calculate the nth discount factor for an n-period swap rate, given as equation (6.9):

$$df_n = \frac{1 - r_n \sum\limits_{i=1}^{n-1} \dfrac{df_i}{N}}{1 + \dfrac{r_n}{N}}. \tag{6.9}$$

Equation (6.9) is the general equation for the *boot-strapping* process of constructing a zero-coupon curve. Essentially, to calculate the n-year discount factor we use the discount factors for the years 1 to $n-1$, and the n-year swap rate or zero-coupon rate. If we have the discount factor for any period, we may use equation (6.9) to determine the same period zero-coupon rate, after rearranging it, shown as equation (6.10):

$$rs_n = t_n \sqrt{\frac{1}{df_n}} - 1. \tag{6.10}$$

2 The equation also assumes an actual/365 day count basis. If any other day count convention is used, the $1/N$ factor must be replaced by a fraction made up of the actual number of days as the numerator and the appropriate year base as the denominator.

Discount factors for spot rates may also be used to calculate forward rates. We know that

$$df_1 = \frac{1}{\left(1 + \frac{rs_1}{N}\right)} \tag{6.11}$$

where *rs* is the zero-coupon rate. If we know the *forward rate* we may use this to calculate a second discount rate, shown by equation (6.12):

$$df_2 = \frac{df_1}{\left(1 + \frac{rf_1}{N}\right)} \tag{6.12}$$

where rf_1 is the forward rate. This is no use in itself, however we may derive from it an equation to enable us to calculate the discount factor at any point in time between the previous discount rate and the given forward rate for the period n to $n+1$, shown as equation (6.13), which may then be rearranged to give us the general equation to calculate a forward rate, given as equation (6.14).

$$df_{n+1} = \frac{df_n}{\left(1 + \frac{rf_n}{N}\right)} \tag{6.13}$$

$$rf_n = \left(\frac{df_n}{df_{n+1}} - 1\right)N. \tag{6.14}$$

The general equation for an n-period discount rate at time n from the previous period forward rates is given by equation (6.15).

$$df_n = \frac{1}{\left(1 + \frac{rf_{n-1}}{N}\right)} \times \frac{1}{\left(1 + \frac{rf_{n-2}}{N}\right)} \times \cdots\cdots \times \frac{1}{\left(1 + \frac{rf_n}{N}\right)} \tag{6.15}$$

$$df_n = \prod_{i=0}^{n-1} \left[\frac{1}{\left(1 + \frac{rf_i}{N}\right)}\right].$$

From the above we may combine equations (6.8) and (6.14) to obtain the general expression for an n-period swap rate and zero-coupon rate, given by equations (6.16) and (6.17) respectively.

$$r_n = \frac{\sum_{i=1}^{n} \frac{rf_{i-1} df_i}{N}}{\sum_{i=1}^{n} \frac{df_i}{F}} \tag{6.16}$$

$$1 + rs_n = t_n \sqrt{\prod_{i=0}^{n-1} \left(1 + \frac{rf_i}{N}\right)}. \tag{6.17}$$

The swap rate, which we have denoted as r_n is shown by equation (6.16) to be the weighted average of the forward rates. Consider that a strip of FRAs would constitute an interest rate swap, that is a swap rate for a continuous period could be covered by a strip of FRAs. Therefore an average of the FRA rates would be the correct swap rate. As FRA rates are forward rates, we may be comfortable with equation (6.16), which states that the n-period swap rate is the average of the forward rates from rf_0 to rf_n. To be accurate we must weight the forward rates, and these are weighted by the discount factors for each period. Note that although swap rates are derived from forward rates, interest payments under a swap are paid in the normal way at the end of an interest period, while payments for an FRA are made at the beginning of the period and must be discounted.

Equation (6.17) states that the zero-coupon rate is calculated from the geometric average of (one plus) the forward rates. Again, this is apparent from a reading of the case study example in Chapter 3. The n-period forward rate is obtained using the discount factors for periods n and $n-1$. The discount factor for the complete period is obtained by multiplying the individual discount factors together, and exactly the same result would be obtained by using the zero-coupon interest rate for the whole period to obtain the discount factor.[3]

6.2.4 *Illustrating the principles for an interest rate swap*

The rate charged on a newly-transacted interest rate swap is the one that gives its net present value as zero. The term *valuation* of a swap is used to denote the process of calculating the net present value of an existing swap, when marking-to-market the swap against current market interest rates. Therefore when we price a swap, we set its net present value to zero, while when we value a swap we set its fixed rate at the market rate and calculate the net present value.

To illustrate the basic principle, we price a plain vanilla interest rate swap with the terms set out below; for simplicity we assume that the annual fixed-rate payments are the same amount each year, although in practice there would be slight differences. Also assume we already have our zero-coupon yields as shown in Table 6.3.

We use the zero-coupon rates to calculate the discount factors, and then use the discount factors to calculate the forward rates. This is done using equation (6.14). These forward rates are then used to predict what the floating-rate payments will be at each

Period	Zero-coupon rate %	Discount factor	Forward rate %	Fixed payment	Floating payment	PV fixed payment	PV floating payment
1	5.5	0.947867298	5.5	689,625	550000	653672.9858	521327.0142
2	6	0.88999644	6.502369605	689,625	650236.9605	613763.7949	578708.58
3	6.25	0.833706493	6.751770257	689,625	675177.0257	574944.8402	562899.4702
4	6.5	0.777323091	7.253534951	689,625	725353.4951	536061.4366	563834.0208
5	7	0.712986179	9.023584719	689,625	902358.4719	491693.094	643369.1194
		4.161879501				2870137	2870137

Table 6.3: Generic interest rate swap

[3] Zero-coupon and forward rates are also related in another way. If we change the zero-coupon rate rs_n and the forward rate rf_i into their continuously compounded equivalent rates, given by $\ln(1+rs_n)$ and $\ln(1+rf_i)$, we may obtain an equation for the continuously compounded zero-coupon rate as being the simple average of the continuously compounded forward rates, given by:

$$rs'_n = \frac{1}{t_n}\sum_{i=0}^{n-1}\frac{rf'_i}{F}.$$

interest period. Both fixed-rate and floating-rate payments are then present valued at the appropriate zero-coupon rate, which enables us to calculate the net present value.

The fixed rate for the swap is calculated using equation (6.8) to give us:

$$\frac{1 - 0.71298618}{4.16187950}$$

or 6.8963%.

Nominal principal	£10,000,000
Fixed rate	6.8963%
Day count fixed	Actual/365
Day count floating	Actual/365
Payment frequency fixed	Annual
Payment frequency floating	Annual
Trade date	31 January 2000
Effective date	2 February 2000
Maturity date	2 February 2005
Term	Five years

For reference the Microsoft Excel® formulae are shown in Table 6.4. It is not surprising that the net present value is zero, because the zero-coupon curve is used to derive the discount factors which are then used to derive the forward rates, which are used to value the swap. As with any financial instrument, the fair value is its breakeven price or hedge cost, and in this case the bank that is pricing the five-year swap shown in Table 6.3 could hedge the swap with a series of FRAs transacted at the forward rates shown. If the bank is paying fixed and receiving floating, the value of the swap to it will rise if there is a rise in market rates, and fall if there is a fall in market rates. Conversely, if the bank was receiving fixed and paying floating, the swap value to it would fall if there was a rise in rates, and vice versa.

This method is used to price any interest rate swap, even exotic ones.

6.2.5 *Valuation using final maturity discount factor*

A short-cut to valuing the floating-leg payments of an interest rate swap involves using the discount factor for the final maturity period. This is possible because, for the purposes of valuation, an exchange of principal at the beginning and end of the swap is conceptually the same as the floating-leg interest payments. This holds because, in an exchange of principal, the interest payments earned on investing the initial principal would be uncertain, as they are floating rate, while on maturity the original principal would be returned. The net result is a floating-rate level of receipts, exactly similar to the floating-leg payments in a swap. To value the principals then, we need only the final maturity discount rate.

To illustrate, consider Table 6.3, where the present value of both legs was found to be £2,870,137. The same result is obtained if we use the five-year discount factor, as shown below.

$$PV_{floating} = (10{,}000{,}000 \times 1) - (10{,}000{,}000 \times 0.71298618) = 2{,}870{,}137.$$

The first term is the principal multiplied by the discount factor 1; this is because the present value of an amount valued immediately is unchanged (or rather, it is multiplied by the immediate payment discount factor, which is 1.0000).

CELL	C	D	E	F	G	H	I	J
21			10000000					
22								
23	**Period**	**Zero-coupon rate %**	**Discount factor**	**Forward rate %**	**Fixed payment**	**Floating payment**	**PV fixed payment**	**PV floating payment**
24	1	5.5	0.947867298	5.5	689,625	″(F24*10000000)/100	″G24/1.055	″H24/(1.055)
25	2	6	0.88999644	″((E24/E25)−1)*100	689,625	″(F25*10000000)/100	″G24/(1.06)^2	″H25/(1.06)^2
26	3	6.25	0.833706493	″((E25/E26)−1)*100	689,625	″(F26*10000000)/100	″G24/(1.0625)^3	″H26/(1.0625)^3
27	4	6.5	0.777323091	″((E26/E27)−1)*100	689,625	″(F27*10000000)/100	″G24/(1.065)^4	″H27/(1.065)^4
28	5	7	0.712986179	″((E27/E28)−1)*100	689,625	″(F28*10000000)/100	″G24/(1.07)^5	″H28/(1.07)^5
			″SUM(E24:E28)				2,870,137	2,870,137

Table 6.4: Generic interest rate swap (Excel formulae)

Therefore we may use the principal amount of a swap if we wish to value the swap. This is of course for valuation only, as there is no actual exchange of principal in a swap.

6.2.6 *Summary of IR swap*

Let us summarise the chief characteristics of swaps. A plain vanilla swap has the following characteristics:

- one leg of the swap is fixed-rate interest, while the other will be floating-rate, usually linked to a standard index such as LIBOR;
- the fixed rate is fixed through the entire life of the swap;
- the floating rate is set in advance of each period (quarterly, semi-annually or annually) and paid in arrears;
- both legs have the same payment frequency;
- the maturity can be standard whole years up to 30 years, or set to match customer requirements;
- the notional principal remains constant during the life of the swap.

Of course to meet customer demand banks can set up swaps that have variations on any or all of the above standard points. Some of the more common variations are discussed in Section 6.3.

6.3 Non-vanilla interest rate swaps

The swap market is very flexible and instruments can be tailor-made to fit the requirements of individual customers. A wide variety of swap contracts have been traded in the market. Although the most common reference rate for the floating leg of a swap is six-month LIBOR, for a semi-annual paying floating leg, other reference rates that have been used include three-month LIBOR, the prime rate (for dollar swaps), the one-month commercial paper rate, the Treasury bill rate and the municipal bond rate (again, for dollar swaps). The term of a swap need not be fixed; swaps may be *extendable* or *putable*. In an extendable swap, one of the parties has the right but not the obligation to extend the life of the swap beyond the fixed maturity date, while in a putable swap one party has the right to terminate the swap ahead of the specified maturity date. It is also possible to transact options on swaps, known as *swaptions*. A swaption is the right to enter into a swap agreement at some point in the future, during the life of the option. Essentially a swaption is an option to exchange a fixed-rate bond cash flow for a floating-rate bond cash flow structure. As a floating-rate bond is valued on its principal value at the start of a swap, a swaption may be viewed as the value on a fixed-rate bond, with a strike price that is equal to the face value of the floating-rate bond.

- A *constant maturity swap* is a swap in which the parties exchange a LIBOR rate for a fixed swap rate. For example, the terms of the swap might state that six-month LIBOR is exchanged for the five-year swap rate on a semi-annual basis for the next five years, or for the five-year government bond rate. In the US market the second type of constant maturity swap is known as a *constant maturity Treasury swap*.

- *Accreting and amortising swaps*. In a plain vanilla swap the notional principal remains unchanged during the life of the swap. However it is possible to trade a swap where the notional principal varies during its life. An accreting (or *step-up*) swap is one in which the principal starts off at one level and then increases in amount over time. The opposite, an amortising swap, is one in which the notional reduces in size over time. An accreting swap would be useful where, for instance a funding liability that is being hedged increases over time. The amortising swap might be employed by a borrower hedging a bond issue that featured sinking fund payments, where a part of the notional amount outstanding is paid off at set points during the life of the bond. If the principal fluctuates in amount, for example increasing in one year and then reducing in another, the swap is known as a *roller-coaster swap*. Another application for an amortising swap is as a hedge for a loan that is itself an amortising one. Frequently this is combined with a forward-starting swap, to tie in with the cash flows payable on the loan. The pricing and valuation of an amortising swap is no different in principle to a vanilla interest rate swap; a single swap rate is calculated using the relevant discount factors, and at this rate the net present value of the swap cash flows will equal zero at the start of the swap.
- *LIBOR-in-arrears swap*. In this type of swap (also known as a *back-set swap*) the setting date is just before the end of the accrual period for the floating rate setting and not just before the start. Such a swap would be attractive to a counterparty who had a different view on interest rates compared to the market consensus. For instance in a rising yield curve environment, forward rates will be higher than current market rates, and this will be reflected in the pricing of a swap. A LIBOR-in-arrears swap would be priced higher than a conventional swap. If the floating-rate payer believed that interest rates would in fact rise more slowly than forward rates (and the market) were suggesting, she may wish to enter into an arrears swap as opposed to a conventional swap.
- *Basis swap*. In a conventional swap one leg comprises fixed-rate payments and the other floating-rate payments. In a basis swap both legs are floating rate, but linked to different money market indices. One leg is normally linked to LIBOR, while the other might be linked to the CD rate say, or the commercial paper rate. This type of swap would be used by a bank in the US that had made loans that paid at the prime rate, and financed its loans at LIBOR. A basis swap would eliminate the *basis risk* between the bank's income and expense cash flows. Other basis swaps have been traded where both legs are linked to LIBOR, but at different maturities; for instance one leg might be at three-month LIBOR and the other at six-month LIBOR. In such a swap the basis is different and so is the payment frequency: one leg pays out semi-annually while the other would be paying on a quarterly basis. Note that where the payment frequencies differ, there is a higher level of counterparty risk for one of the parties. For instance, if one party is paying out on a monthly basis but receiving semi-annual cash flows, it would have made five interest payments before receiving one in return.
- *Margin swap*. It is common to encounter swaps where there is a margin above or below LIBOR on the floating leg, as opposed to a floating leg of LIBOR flat. If a bank's borrowing is financed at LIBOR + 25 bps, it may wish to receive LIBOR + 25 bps in the swap so that its cash flows match exactly. The fixed-rate quote for a swap must be adjusted correspondingly to allow for the margin on the floating side, so in our example if the fixed-rate quote is say, 6.00%, it would be adjusted to around 6.25%; differences in

the margin quoted on the fixed leg might arise if the day count convention or payment frequency were to differ between fixed and floating legs. Another reason why there may be a margin is if the credit quality of the counterparty demanded it, so that highly rated counterparties may pay slightly below LIBOR, for instance.

- *Off-market swap*. When a swap is transacted its fixed rate is quoted at the current market rate for that maturity. Where the fixed rate is different to the market rate, this is an off-market swap, and a compensating payment is made by one party to the other. An off-market rate may be used for particular hedging requirements for example, or when a bond issuer wishes to use the swap to hedge the bond as well as to cover the bond's issue costs.
- *Differential swap*. A differential swap is a basis swap but with one of the legs calculated in a different currency. Typically one leg is floating rate, while the other is floating rate but with the reference index rate for another currency, but denominated in the domestic currency. For example, a differential swap may have one party paying six-month sterling LIBOR, in sterling, on a notional principal of £10 million, and receiving euro-LIBOR, minus a margin, payable in sterling and on the same notional principal. Differential swaps are not very common and are the most difficult for a bank to hedge. The hedging is usually carried out using what is known as a *quanto* option.
- *Forward-start swap*. A forward-start swap is one where the *effective date* is not the usual one or two days after the trade date but a considerable time afterwards, for instance say six months after trade date. Such a swap might be entered into where one counterparty wanted to fix a hedge or cost of borrowing now, but for a point some time in the future. Typically this would be because the party considered that interest rates would rise or the cost of hedging would rise. The swap rate for a forward-starting swap is calculated in the same way as that for a vanilla swap.

6.4 Currency swaps

So far we have discussed swap contracts where the interest payments are both in the same currency. A *cross-currency* swap, or simply *currency swap*, is similar to an interest rate swap, except that the currencies of the two legs are different. Like interest rate swaps, the legs are usually fixed and floating rate, although again it is common to encounter both fixed-rate or both floating-rate legs in a currency swap. On maturity of the swap there is an exchange of principals, and usually (but not always) there is an exchange of principals at the start of the swap. Where currencies are exchanged at the start of the swap, at the prevailing spot exchange rate for the two currencies, the exact amounts are exchanged back on maturity. During the time of the swap, the parties make interest payments in the currency that they have *received* where principals are exchanged. It may seem that exchanging the same amount on maturity gives rise to some sort of currency risk, in fact it is this feature that removes any element of currency risk from the swap transaction.

Currency swaps are widely used in association with bond issues by borrowers who seek to tap opportunities in different markets but have no requirement for that market's currency. By means of a currency swap, a company can raise funds in virtually any market and swap the proceeds into the currency that it requires. Often the underwriting bank that is responsible for the bond issue will also arrange for the currency swap to be transacted. In a currency swap therefore, the exchange of principal means that the value of the principal amounts must be accounted for, and is dependent on the prevailing spot exchange rate between the two currencies.

6.4.1 *Valuation of currency swaps*

The same principles we established for the pricing and valuation of interest rate swaps may be applied to currency swaps. A generic currency swap with fixed-rate payment legs would be valued at the fair value swap rate for each currency, which would give a net present value of zero. The cash flows are illustrated in Figure 6.3. This shows that the two swap rates in a fixed–fixed currency swap would be identical to the same-maturity swap rate for each currency interest rate swap. So the swap rates for a fixed–fixed five-year sterling/dollar currency swap would be the five-year sterling swap rate and the five-year dollar swap rate.

A floating–floating currency swap may be valued in the same way, and for valuation purposes the floating-leg payments are replaced with an exchange of principals, as we observed for the floating leg of an interest rate swap. A fixed–floating currency swap is therefore valued at the fixed-rate swap rate for that currency for the fixed leg, and at LIBOR or the relevant reference rate for the floating leg.

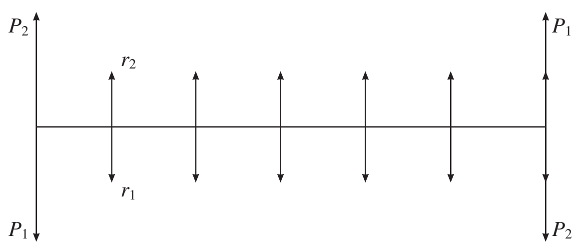

Figure 6.3: Fixed–fixed rate currency swap

Example 6.1: Bond issue and associated cross-currency swap

We illustrate currency swaps with an example from the market. A subsidiary of a US bank that invests in projects in the United States issues paper in markets around the world, in response to investor demand worldwide. The company's funding requirement is in US dollars, however it is active in issuing bonds in various currencies, according to where the most favourable conditions can be obtained. When an issue of debt is made in a currency other than dollars, the proceeds must be swapped into dollars for use in the USA, and interest payable on the swapped (dollar) proceeds. To facilitate this the issuer will enter into a currency swap. One of the bank's issues was a Swiss franc step-up bond, part of an overall Euro-MTN programme. The details of the bond are summarised below.

Issue date	March 1998
Maturity	March 2003
Size	CHF 15 million
Coupon	2.40% to 25/3/99
	2.80% to 25/3/00
	3.80% to 25/3/01
	4.80% to 25/3/02

The bond was also callable on each anniversary from March 1999 onwards, and in fact was called by the issuer at the earliest opportunity. The issuing bank entered into a currency swap that resulted in the exchange of principals and the Swiss franc interest payments to be made by the swap counterparty; in return it paid US dollar three-month LIBOR during the life of the swap. At the prevailing spot rate on the effective date, CHF 15 million was exchanged for $10.304 million; these exact same amounts would be exchanged back on the maturity of the swap. When the issue was called the swap was cancelled and the swap counterparty paid a cancellation fee. The interest payment dates on the fixed leg of the swap matched the coupon dates of the bond exactly, as shown above. The floating leg of the swap paid US dollar LIBOR on a quarterly basis, as required by the bond issuer.

The structure is shown in Figure 6.4. A currency swap structure enables a bank or corporate to borrow money in virtually any currency in which a liquid swap market exists, and swap this into a currency that is required. In our example the US bank was able to issue a bond that was attractive to investors. The swap mechanism also hedged the interest rate exposure on the Swiss franc note. The liability remaining for the issuer was quarterly floating-rate interest on US dollars as part of the swap transaction.

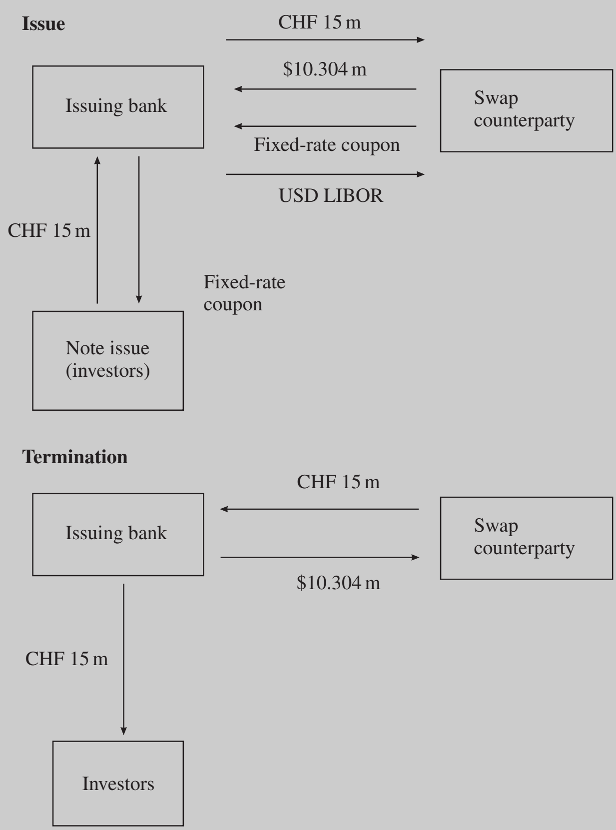

Figure 6.4: Bond issue with currency swap structure

6.5 Swaptions

6.5.1 *Description*

Swaptions are options on swaps. The buyer of a swaption has the right but not the obligation to enter into an interest rate swap agreement during the life of the option. The terms of the swaption will specify whether the buyer is the fixed- or floating-rate payer; the seller of the option (the *writer*) becomes the counterparty to the swap if the option is exercised. The convention is that if the buyer has the right to exercise the option as the fixed-rate payer, he has traded a *call swaption*, also known as a *payer swaption*, while if by exercising the buyer of the swaption becomes the floating-rate payer he has bought a *put swaption*, also known as a *receiver swaption*. The writer of the swaption is the party to the other leg.

Swaptions are similar to forward-start swaps up to a point, but the buyer has the *option* of whether or not to commence payments on the effective date. A bank may purchase a call swaption if it expects interest rates to rise, and will exercise the option if indeed rates do rise as the bank has expected. A company will use swaptions as part of an interest rate hedge for a future exposure. For example, assume that a company will be entering into a five-year bank loan in three months' time. Interest on the loan is charged on a floating-rate basis, but the company intends to swap this to a fixed-rate liability after they have entered into the loan. As an added hedge, the company may choose to purchase a swaption that gives it the right to receive LIBOR and pay a fixed rate, say 10%, for a five-year period beginning in three months' time. When the time comes for the company to take out a swap and exchange its interest rate liability in three months' time (having entered into the loan), if the five-year swap rate is below 10%, the company will transact the swap in the normal way and the swaption will expire worthless. However if the five-year swap rate is above 10%, the company will instead exercise the swaption, giving it the right to enter into a five-year swap and paying a fixed rate of 10%. Essentially the company has taken out protection to ensure that it does not have to pay a fixed rate of more than 10%. Hence swaptions can be used to guarantee a maximum swap rate liability. They are similar to forward-starting swaps, but do not commit a party to enter into a swap on fixed terms. The swaption enables a company to hedge against unfavourable movements in interest rates but also to gain from favourable movements, although there is of course a cost associated with this, which is the premium paid for the swaption.

As with conventional put and call options, swaptions turn in-the-money under opposite circumstances. A call swaption increases in value as interest rates rise, and a put swaption becomes more valuable as interest rates fall. Consider a one-year European call swaption on a five-year semi-annual interest rate swap, purchased by a bank counterparty. The notional value is £10 million and the "strike price" is 6%, against LIBOR. Assume that the price (premium) of the swaption is 25 basis points, or £25,000. On expiry of the swaption, the buyer will either exercise it, in which case she will enter into a five-year swap paying 6% and receiving floating-rate interest, or elect to let the swaption expire with no value. If the five-year swap rate for a counterparty of similar credit quality to the bank is above 6%, the swaption holder will exercise the swaption, while if the rate is below 6% the buyer will not exercise. The principle is the same for a put swaption, only in reverse.

6.5.2 *Valuation*

Swaptions are typically priced using the Black–Scholes or Black 76 option pricing models. These are used to value a European option on a swap, assuming that the appropriate swap rate at the expiry date of the option is lognormal. Consider a swaption with the following general terms:

Swap rate on expiry	rs
Swaption strike rate	rX
Maturity	T
Start date	t
Pay basis	F (say quarterly, semi-annual or annual)
Notional principal	M

If the actual swap rate on the maturity of the swaption is rs, the payoff from the swaption is given by:

$$\frac{M}{F}\max(r - r_n, 0).$$

The value of a swaption is essentially the difference between the strike rate and the swap rate at the time it is being valued. If a swaption is exercised, the payoff at each interest date is given by $(rs - rX) \times M \times F$. As a call swaption is only exercised when the swap rate is higher than the strike rate (that is, $rs > rX$), the option payoff on any interest payment in the swap is given by:

$$Swaption_{InterestPayment} = \mathrm{Max}[0, (rs - rX) \times M \times F]. \tag{6.18}$$

It can then easily be shown that the value of a call swaption on expiry is given by:

$$PV_{Swaption} = \sum_{n=1}^{N} Df_{(0,n)}(rs - rX) \times M \times F \tag{6.19}$$

where $Df_{(0,n)}$ is the spot rate discount factor for the term beginning now and ending at time t. By the same logic the value of a put swaption is given by the same equation except that $(rX - rs)$ is substituted at the relevant point above.

Pressing on, a swaption can be viewed as a collection of calls or puts on interest deposits or LIBOR, enabling us to use the Black 76 model when valuing it. This means that we value each call or put for a single payment in the swap, and then sum these payments to obtain the value of the swaption. The main assumption made when using this model is that the LIBOR rate follows a lognormal distribution over time, with constant volatility.

Consider a call swaption being valued at time t that matures at time T. We begin by valuing a single payment under the swap (assuming the option is exercised) made at time T_n. The point at time T_n is into the life of the swap, so that we have $T_n > T > t$. At the time of valuation, the option time to expiry is $T - t$ and there is $T_n - t$ until the nth payment. The value of this payment is given by:

$$C_t = MFe^{-r(T_n - t)}[rsN(d_1) - rXN(d_2)] \tag{6.20}$$

where

C_t	is the price of the call option on a single payment in the swap
r	is the risk-free instantaneous interest rate
$N(.)$	is the cumulative normal distribution
σ	is the interest rate volatility

and where

$$d_1 = \frac{\ln(rs_t/rX) + \frac{\sigma^2}{2}(T-t)}{\sigma\sqrt{T-t}}$$

$$d_2 = d_1 - \sigma\sqrt{T-t}.$$

The remaining life of the swaption $(T-t)$ governs the probability that it will expire in-the-money, determined using the lognormal distribution. On the other hand the interest payment itself is discounted (using $e^{-r(T_n-t)}$) over the period $T_n - t$ as it is not paid until time T_n.

Having valued a single interest payment, viewing the swap as a collection of interest payments, we value the call swaption as a collection of calls. Its value is given therefore by:

$$PV_{Swaption_t} = \sum_{n=1}^{N} MFe^{-r(T_n-t)}[rsN(d_1) - rXN(d_2)] \tag{6.21}$$

where t, T and n are as before.

If we substitute discrete spot rate discount factors instead of the continuous form given by equation (6.21) the equation becomes

$$PV_{Swaption_t} = MF[rsN(d_1) - rXN(d_2)] \sum_{n=1}^{N} Df_{t,T_n}. \tag{6.22}$$

Example 6.2: Swaption pricing[4]

We present a new term structure environment in this example to illustrate the basic concepts. This is shown in Table 6.5. We wish to price a forward-starting annual interest swap starting in two years for a term of three years. The swap has a notional of £10 million.

Date	Term (years)	Discount factor	Par yield	Zero-coupon rate	Forward rate
18/02/2001	0	1	5	5	5
18/02/2002	1	0.95238095	5.00	5	6.03015
18/02/2003	2	0.89821711	5.50	5.51382	7.10333
18/02/2004	3	0.83864539	6.00	6.04102	6.66173
18/02/2005	4	0.78613195	6.15	6.19602	6.71967
20/02/2006	5	0.73637858	6.25	6.30071	8.05230
19/02/2007	6	0.68165163	6.50	6.58946	8.70869
18/02/2008	7	0.62719194	6.75	6.88862	9.40246
18/02/2009	8	0.57315372	7.00	7.20016	10.18050
18/02/2010	9	0.52019523	7.25	7.52709	5.80396
18/02/2011	10	0.49165950	7.15	7.35361	6.16366

Table 6.5: Interest rate data for swaption valuation

[4] This example follows the approach described in Kolb (2000); although here we use discount factors in the calculation whereas Kolb uses *zero-coupon factors* which are $(1 + \text{spot rate})$.

The swap rate is given by

$$rs = \frac{\sum_{t=1}^{N} rf_{(t-1),t} \times Df_{0,t}}{\sum_{t=1}^{N} Df_{0,t}}$$

where rf is the forward rate.

Using the above equation, the numerator in this example is

$$(0.0666 \times 0.8386) + (0.0672 \times 0.7861) + (0.0805 \times 0.7634) \text{ or } 0.1701.$$

The denominator is

$$0.8386 + 0.7861 + 0.7634 \text{ or } 2.3881.$$

Therefore the forward-starting swap rate is 0.1701/2.3881 or 0.071228 (7.123%).

We now turn to the call swaption on this swap, the buyer of which acquires the right to enter into a three-year swap paying the fixed-rate swap rate of 7.00%. If the volatility of the forward swap rate is 0.20, the d_1 and d_2 terms are

$$d_1 = \frac{\ln\left(\frac{rs}{rX}\right) + \frac{\sigma^2}{2}(T-t)}{\sigma\sqrt{T-t}} = \frac{\ln\left(\frac{0.071228}{0.07}\right) + \left(\frac{0.2^2}{2} \times 2\right)}{0.2\sqrt{2}}$$

or 0.2029068

$$d_2 = d_1 - \sigma\sqrt{T-t} = 0.20290618 - 0.2(1.4142)$$

or -0.079934.

The cumulative normal values are

$N(d_1) = N(0.2029)$ which is 0.580397
$N(d_2) = N(-0.079934)$ which is 0.468145.[5]

From above we know that $\sum Df_{t,T_n}$ is 2.3881. So using equation (6.22) we calculate the value of the call swaption to be

$$\begin{aligned} PV_{Swaption_t} &= MF[rsN(d_1) - rXN(d_2)] \sum_{n=1}^{N} Df_{t,T_n} \\ &= 10{,}000{,}000 \times 1 \times [0.07228 \times 0.580397 - 0.07 \times 0.468145] \times 2.3881 \\ &= 219{,}250 \end{aligned}$$

or £219,250. Option premiums are frequently quoted as basis points of the notional amount, so in this case the premium is (219,250/10,000,000) or 219.25 basis points.

5 These values may be found from standard normal distribution tables or using the Microsoft Excel formula = NORMSDIST().

6.6 An overview of interest rate swap applications

In this section we review some of the principal uses of swaps as a hedging tool for bond instruments and also how to hedge a swap book.

6.6.1 *Corporate applications*

Swaps are part of the over-the-counter (OTC) market and so they can be tailored to suit the particular requirements of the user. It is common for swaps to be structured so that they match particular payment dates, payment frequencies and LIBOR margins, which may characterise the underlying exposure of the customer. As the market in interest rate swaps is so large, liquid and competitive, banks are willing to quote rates and structure swaps for virtually all customers, although it may be difficult for smaller customers to obtain competitive price quotes as notional values below £10 million or £5 million.

Swap applications can be viewed as being one of two main types, asset-linked swaps and liability-linked swaps. Asset-linked swaps are created when the swap is linked to an asset such as a bond in order to change the characteristics of the income stream for investors. Liability-linked swaps are traded when borrowers of funds wish to change the pattern of their cash flows. Of course, just as with repo transactions, the designation of a swap in such terms depends on from whose point of view one is looking at the swap. An asset-linked swap hedge is a liability-linked hedge for the counterparty, except in the case of swap market making banks who make two-way quotes in the instruments.

A straightforward application of an interest rate swap is when a borrower wishes to convert a floating-rate liability into a fixed-rate one, usually in order to remove the exposure to upward moves in interest rates. For instance a company may wish to fix its financing costs. Let us assume a company currently borrowing money at a floating rate, say six-month LIBOR + 100 basis points fears that interest rates may rise in the remaining three years of its loan. It enters into a three-year semi-annual interest rate swap with a bank, as the fixed-rate payer, paying say 6.75% against receiving six-month LIBOR. This fixes the company's borrowing costs for three years at 7.75% (7.99% effective annual rate). This is shown in Figure 6.5.

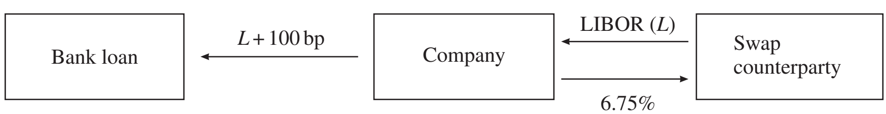

Figure 6.5: Changing liability from floating to fixed rate

Example 6.3: Liability-linked swap, fixed- to floating- to fixed-rate exposure

A corporate borrows for five years at a rate of 6¼% and shortly after enters into a swap paying floating rate, so that its net borrowing cost is LIBOR + 40 bps. After one year swap rates have fallen such that the company is quoted four-year swap rates as 4.90–84%. The company decides to switch back into fixed-rate liability in order to take advantage of the lower interest rate environment. It enters into a second swap paying fixed at 4.90% and receiving LIBOR. The net borrowing cost is now 5.30%. The arrangement is illustrated in Figure 6.6. The company has saved 95 basis points on its original borrowing cost, which is the difference between the two swap rates.

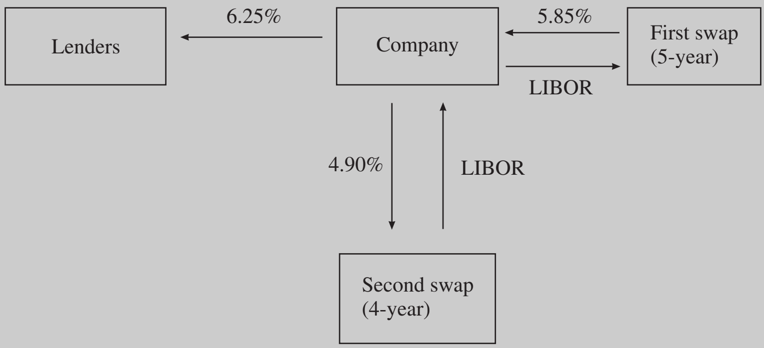

Figure 6.6

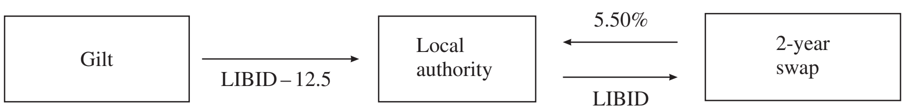

Figure 6.7: Transforming a floating-rate asset to fixed rate

Asset-linked swap structures might be required when for example, investors require a fixed interest security when floating-rate assets are available. Borrowers often issue FRNs, the holders of which may prefer to switch the income stream into fixed coupons. As an example, consider a local authority pension fund holding two-year floating-rate gilts. This is an asset of the highest quality, paying LIBID minus 12.5 bps. The pension fund wishes to swap the cash flows to create a fixed-interest asset. It obtains a quote for a tailor-made swap where the floating leg pays LIBID, the quote being 5.55–50%. By entering into this swap the pension fund has in place a structure that pays a fixed coupon of 5.375%. This is shown in Figure 6.7.

6.6.2 *Hedging bond instruments using interest rate swaps*

We illustrate here a generic approach to the hedging of bond positions using interest rate swaps. The bond trader has the option of using other bonds, bond futures or bond options, as well as swaps, when hedging the interest rate risk exposure of a bond position. However swaps are particularly efficient instruments to use because they display positive convexity characteristics, that is the increase in value of a swap for a fall in interest rates exceeds the loss in value with a similar magnitude rise in rates. This is exactly the price/yield profile of vanilla bonds.

The primary risk measure we require when hedging using a swap is its present value of a basis point or PVBP.[6] This measures the price sensitivity of the swap for a basis point

[6] This is also known as DVBP or dollar value of a basis point in the US market.

change in interest rates. The PVBP measure is used to calculate the hedge ratio when hedging a bond position. The PVBP can be given by

$$PVBP = \frac{\text{Change in swap value}}{\text{Rate change in basis points}} \tag{6.23}$$

which can be written as

$$PVBP = \frac{dS}{dr}. \tag{6.24}$$

Using the basic relationship for the value of a swap, which is viewed as the difference between the values of a fixed-coupon bond and an equivalent-maturity floating-rate bond (see Table 6.1) we can also write

$$PVBP = \frac{d\text{Fixed bond}}{dr} - \frac{d\text{Floating bond}}{dr} \tag{6.25}$$

which essentially states that the basis point value of the swap is the difference in the basis point values of the fixed-coupon and floating-rate bonds. The value is usually calculated for a notional £1 million of swap. The calculation is based on the duration and modified duration calculations used for bonds[7] and assumes that there is a parallel shift in the yield curve.

Figure 6.8 illustrates how equations(6.24) and (6.25) can be used to obtain the PVBP of a swap. Hypothetical five-year bonds are used in the example. The PVBP for a bond can be calculated using Bloomberg or the MDURATION function on Microsoft Excel. Using either of the two equations above we see that the PVBP of the swap is £425.00. This is shown below.

Calculating the PVBP using equation (6.24) we have:

$$PVBP_{swap} = \frac{dS}{dr} = \frac{4264 - (-4236)}{20} = 425 \tag{6.26}$$

Interest rate swap			
Term to maturity	5 years		
Fixed leg	6.50%		
Basis	Semi-annual, act/365		
Floating leg	6-month LIBOR		
Basis	Semi-annual, act/365		
Nominal amount	£1,000,000		
		Present value (£)	
	Rate change – 10 bps	0 bps	Rate change + 10 bps
Fixed-coupon bond	1,004,940	1,000,000	995,171
Floating-rate bond	1,000,640	1,000,000	999,371
Swap	4,264	0	4,236

Figure 6.8: PVBP for interest rate swap

[7] See Chapter 2.

while using equation (6.25) we obtain the same result using the bond values

$$\begin{aligned} PVBP_{swap} &= PVBP_{fixed} - PVBP_{floating} \\ &= \frac{1004940 - 995171}{20} - \frac{1000640 - 999371}{20} \\ &= 488.45 - 63.45 \\ &= 425.00. \end{aligned} \tag{6.27}$$

The swap basis point value is lower than that of the five-year fixed-coupon bond, that is £425 compared to £488.45. This is because of the impact of the floating-rate bond risk measure, which reduces the risk exposure of the swap as a whole by £63.45. As a rough rule of thumb, the PVBP of a swap is approximately equal to that of a fixed-rate bond that has a maturity similar to the period from the next coupon reset date of the swap through to the maturity date of the swap. This means that a 10-year semi-annual paying swap would have a PVBP close to that of a 9.5-year fixed-rate bond, and a 5.50-year swap would have a PVBP similar to that of a 5-year bond.

When using swaps as hedge tools, we bear in mind that over time the PVBP of swaps behaves differently to that of bonds. Immediately preceding an interest reset date the PVBP of a swap will be near-identical to that of the same-maturity fixed-rate bond, because the PVBP of a floating-rate bond at this time has essentially nil value. Immediately after the reset date the swap PVBP will be near-identical to that of a bond that matures at the next reset date. This means that at the point (and this point only) right after the reset the swap PVBP will decrease by the amount of the floating-rate PVBP. In between reset dates the swap PVBP is quite stable, as the effects of the fixed- and floating-rate PVBP changes cancel each other out. Contrast this with the fixed-rate PVBP, which decreases in value over time in stable fashion.[8] This feature is illustrated in Figure 6.9. A slight anomaly is that the PVBP of a swap actually increases by a small amount between reset dates; this is because the PVBP of a floating-rate bond decreases at a slightly faster rate than that of the fixed-rate bond during this time.

Hedging bond instruments with interest rate swaps is conceptually similar to hedging with another bond or with bond futures contracts. If one is holding a long position in a vanilla bond, the hedge requires a long position in the swap: remember that a long position in a swap is to be paying fixed (and receiving floating). This hedges the receipt of fixed from the bond position. The change in the value of the swap will match the change in value of the bond, only in the opposite direction.[9] The maturity of the swap should match as closely as possible that of the bond. As swaps are OTC contracts, it should be possible to match interest dates as well as maturity dates. If one is short the bond, the hedge is to be short the swap, so the receipt of fixed matches the pay-fixed liability of the short bond position.

The correct nominal amount of the swap to put on is established using the PVBP hedge ratio. This is given as

$$\text{Hedge ratio} = \frac{PVBP_{bond}}{PVBP_{swap}}. \tag{6.28}$$

8 This assumes no large-scale sudden yield movements.

9 The change will not be an exact mirror. It is very difficult to establish a precise hedge for a number of reasons, which include differences in day count bases, maturity mismatches and basis risk.

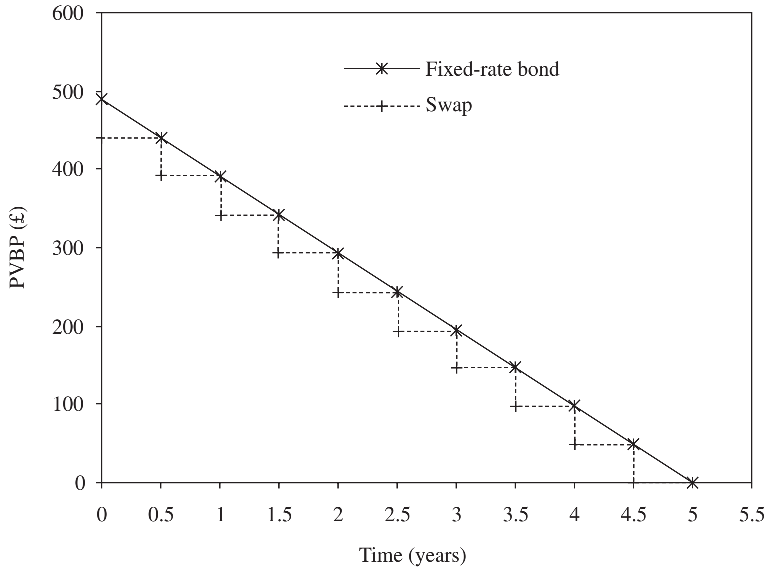

Figure 6.9: PVBP of 5-year swap

This technique is still used in the market but suffers from the assumption of parallel yield curve shifts and can therefore lead to significant hedging error at times. More advanced techniques are used by banks when hedging books using swaps, which space does not permit any discussion of. Some of these techniques are discussed in the bibliography.

Selected bibliography and references

The swaps market is extensively covered in the academic and practitioner literature. Das (1994) contains all one would ever wish to know and is written in a very user-friendly style. Its publication date means that it does not cover the very latest developments, but it is well worth buying. Be warned though, it is a weighty tome. A good general introduction is in Decovny (1998), while another accessible and reader-friendly discussion is given in Eales (1995). A superlative overview treatment is given in Chapter 14 of Jarrow and Turnbull (2000), and this is a market standard text. Technical aspects are covered in the journal references.

A good account of hedging with swaps can be found in Chapter 11 of Fabozzi (1998), which is authored by Shrikant Ramamurthy. For budding treasury traders amongst you Chapter 6 in Gup and Brooks (1993) is a readable introduction to swaps and to asset and liability management.

For the latest on swap trading volumes, and electronic trading on platforms such as *SwapsWire* see Dunbar (2000), Khan (2000) and Lindsay (2000).

Bicksler, J., Chen, A., "An Economic Analysis of Interest Rate Swaps", *Journal of Finance* 41(3), 1986, pp. 645–655

Brotherton-Ratcliffe, R., Iben, B., "Yield Curve Applications of Swap Products", in Schwartz, R., Smith, C. (eds), *Advanced Strategies in Financial Risk Management*, New York Institute of Finance 1993

Choudhry, M., *Bond Market Securities*, FT Prentice Hall 2001
Das, S., *Swaps and Financial Derivatives*, 2nd edition, IFR Publishing 1994
Decovny, S., *Swaps*, 2nd edition, FT Prentice Hall 1998
Dunbar, N., "Swaps volumes see euro wane", *Risk*, September 2000
Eales, B., *Financial Risk Management*, McGraw-Hill 1995, Chapter 3
Fabozzi, F. (editor), *Perspectives on Interest Rate Risk Management for Money Managers and Traders*, FJF Associates 1998
Gup, B., Brooks, R., *Interest Rate Risk Management*, Irwin 1993
Henna, P., *Interest-rate Risk Management using Futures and Swaps*, Probus 1991
International Swaps and Derivatives Association, *Code of Standard Working, Assumptions and Provisions for Swaps*, New York, 1991
Jarrow, R., Turnbull, S., *Derivative Securities*, 2nd edition, South-Western 2000
Khan, M., "Online platforms battle for business", *Risk*, September 2000
Kolb, R., *Futures, Options and Swaps*, 3rd edition, Blackwell 2000
Li, A., Raghavan, V. R., "LIBOR-In-Arrears Swaps", *Journal of Derivatives* 3, Spring 1996, pp. 44–48
Lindsay, R., "High Wire Act", *Risk*, August 2000
Marshall, J., Kapner, K., *Understanding Swap Finance*, South-Western 1990
Turnbull, S., "Swaps: A Zero Sum Game", *Financial Management* 16, Spring 1987, pp. 15–21

7 Credit Derivatives

Credit derivatives allow investors to manage the credit risk exposure of their portfolios or asset holdings, essentially by providing insurance against a deterioration in credit quality of the borrowing entity.[1] If there is a technical default by the borrower[2] or an actual default on the loan itself, and the bond is marked down in price, the losses suffered by the investor can be recouped in part or in full through the payout made by the credit derivative.

7.1 Credit risk

Credit risk is the risk that a borrowing entity will default on a loan, either through inability to maintain the interest servicing or because of bankruptcy or insolvency leading to inability to repay the principal itself. When technical or actual default occurs, bondholders suffer a loss as the value of their asset declines, and the potential greatest loss is that of the entire asset. The extent of credit risk fluctuates as the fortunes of borrowers change in line with their own economic circumstances and the macroeconomic business cycle. The magnitude of risk is described by a firm's credit rating. Ratings agencies undertake a formal analysis of the borrower, after which a rating is announced; the issues considered in the analysis include:

- the financial position of the firm itself, for example, its balance sheet position and anticipated cash flows and revenues;
- other firm-specific issues such as the quality of the management and succession planning;
- an assessment of the firm's ability to meet scheduled interest and principal payments, both in its domestic and foreign currencies;
- the outlook for the industry as a whole, and competition within it;
- general assessments for the domestic economy.

Another measure of credit risk is the credit risk premium, which is the difference between yields on the same-currency government benchmark bonds and corporate bonds. This premium is the compensation required by investors for holding bonds that are not default-free. The credit premium required will fluctuate as individual firms and sectors are perceived to offer improved or worsening credit risk, and as the general health of the economy improves or worsens. For example, Figure 7.1 illustrates the variability in credit spread premium in the sterling market, illustrated by the change in yields for 10-year bonds rated AAA, AA and A against the benchmark gilt yield during 1998–2000.

---

1 The simplest credit derivative works exactly like an insurance policy, with regular premiums paid by the protection-buyer to the protection-seller, and a payout in the event of a specified credit event.

2 A technical default is a delay in timely payment of the coupon, or non-payment of the coupon altogether.

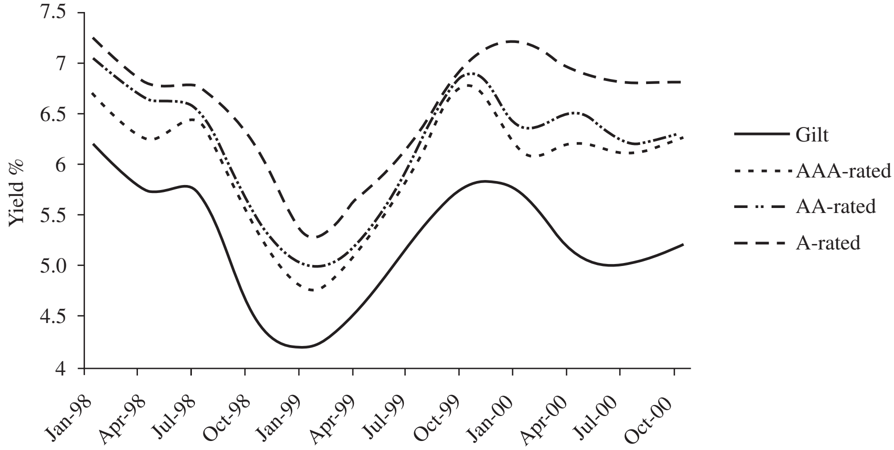

Figure 7.1: 10-year sterling bond yields.
(Source: Bank of England and Bloomberg)

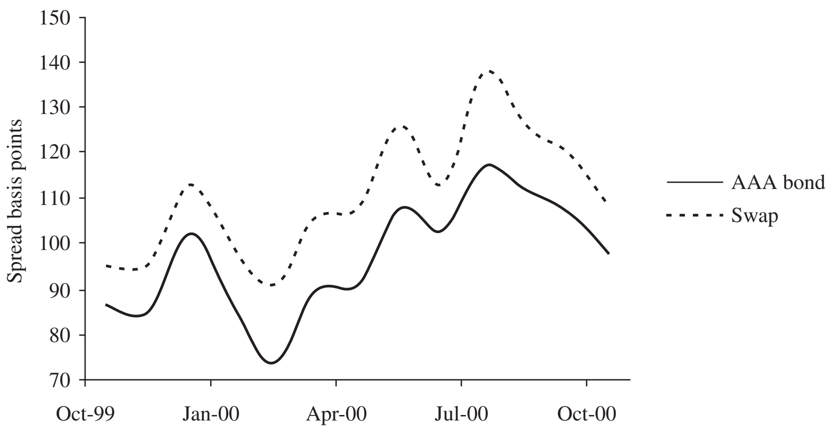

Figure 7.2: 10-year sterling credit spread.
(Source: Bank of England and Bloomberg)

7.2 Credit risk and credit derivatives

Credit derivatives are financial contracts designed to reduce or eliminate credit risk exposure by providing insurance against losses suffered due to credit events. A payout under a credit derivative is triggered by a credit event. As banks define default in different ways, the terms under which a credit derivative is executed usually include a specification of what constitutes a credit event.

The principle behind credit derivatives is straightforward. Investors desire exposure to non-default-free sovereign debt because of the higher returns this offers. However such exposure brings with it concomitant credit risk. This can be managed with credit derivatives. At the same time, the exposure itself can be taken on synthetically if for instance, there are compelling reasons why a cash market position cannot be established. The flexibility of credit derivatives provides users with a number of advantages and as they are over-the-counter (OTC) products they can be designed to meet specific user requirements.

We focus on credit derivatives as instruments that may be used to manage risk exposure inherent in a corporate or non-AAA sovereign bond portfolio. They may also be used to manage the credit risk of commercial loan books. The intense competition amongst commercial banks, combined with rapid disintermediation, has meant that banks have been forced to evaluate their lending policy, with a view to improving profitability and return on capital. The use of credit derivatives assists banks with restructuring their businesses, because they allow banks to repackage and parcel out credit risk, while retaining assets on-balance sheet (when required) and thus maintain client relationships. As the instruments isolate certain aspects of credit risk from the underlying loan or bond and transfer them to another entity, it becomes possible to separate the ownership and management of credit risk from the other features of ownership associated with the assets in question. This means that illiquid assets such as bank loans, and illiquid bonds, can have their credit risk exposures transferred; the bank owning the assets can protect against credit loss even if it cannot transfer the assets themselves.

The same principles carry over to the credit risk exposures of portfolio managers. For fixed-income portfolio managers some of the advantages of credit derivatives include the following:

- they can be tailor-made to meet the specific requirements of the entity buying the risk protection, as opposed to the liquidity or term of the underlying reference asset;
- in theory, they can be "sold short" without risk of a liquidity or delivery squeeze, as it is a specific credit risk that is being traded. In the cash market it is not possible to "sell short" a bank loan for example, but a credit derivative can be used to establish synthetically the economic effect of such a position;
- as they theoretically isolate credit risk from other factors such as client relationships and interest rate risk, credit derivatives introduce a formal pricing mechanism to price credit issues only. This means a market can develop in credit only, allowing more efficient pricing, and it becomes possible to model a term structure of credit rates;
- they are off-balance sheet instruments[3] and as such incorporate tremendous flexibility and leverage, exactly like other financial derivatives. For instance, bank loans are not particularly attractive investments for certain investors because of the administration required in managing and servicing a loan portfolio. However an exposure to bank loans and their associated return can be achieved by say, a total return swap while simultaneously avoiding the administrative costs of actually owning the assets. Hence credit derivatives allow investors access to specific credits while allowing banks access to further distribution for bank loan credit risk.

Thus credit derivatives can be an important instrument for bond portfolio managers as well as commercial banks, who wish to increase the liquidity of their portfolios, gain from the

---

[3] When credit derivatives are embedded in certain fixed-income products, such as structured notes and credit-linked notes, they are then off-balance sheet but part of an on-balance sheet structure.

relative value arising from credit pricing anomalies, and enhance portfolio returns. Some key applications are summarised below.

7.2.1 *Diversifying the credit portfolio*

A bank or portfolio manager may wish to take on credit exposure by providing credit protection on assets that it already owns, in return for a fee. This enhances income on their portfolio. They may sell credit derivatives to enable non-financial counterparties to gain credit exposures, if these clients do not wish to purchase the assets directly. In this respect the bank or asset manager performs a credit intermediation role.

7.2.2 *Reducing credit exposure*

A bank can reduce credit exposure either for an individual loan or a sectoral concentration, by buying a credit default swap. This may be desirable for assets in their portfolio that cannot be sold for client relationship or tax reasons. For fixed-income managers a particular asset or collection of assets may be viewed as favourable holdings in the long-term, but at risk from short-term downward price movement. In this instance a sale would not fit in with long-term objectives, however short-term credit protection can be obtained via credit swap.

7.2.3 *Acting as a credit derivatives market maker*

A financial entity may wish to set itself up as a market maker in credit derivatives. In this case it may or may not hold the reference assets directly, and depending on its appetite for risk and the liquidity of the market, it can offset derivative contracts as and when required.

7.3 Credit event

The occurrence of a specified credit event will trigger payment of the default payment by the seller of protection to the buyer of protection. Contracts specify physical or cash settlement. In physical settlement, the protection buyer transfers to the protection seller the deliverable obligation (usually the reference asset or assets), with the total principal outstanding equal to the nominal specified in the default swap contract. The protection seller simultaneously pays to the buyer 100% of the nominal. In cash settlement, the protection seller hands to the buyer the difference between the nominal amount of the default swap and the final value for the same nominal amount of the reference asset. This final value is usually determined by means of a poll of dealer banks.

The following may be specified as credit events in the legal documentation between counterparties:

- downgrade in S&P and/or Moody's credit rating below a specified minimum level;
- financial or debt restructuring, for example occasioned under administration or as required under US bankruptcy protection;
- bankruptcy or insolvency of the reference asset obligor;
- default on payment obligations such as bond coupon and continued non-payment after a specified time period.
- technical default, for example the non-payment of interest or coupon when it falls due;
- a change in credit spread payable by the obligor above a specified maximum level.

The 1999 ISDA credit default swap documentation specifies bankruptcy, failure to pay, obligation default, debt moratorium and restructuring to be credit events. Note that it does not specify a rating downgrade to be a credit event.

Credit derivative instruments

We now review some of the most commonly encountered credit derivative instruments.

Credit default swap

The most common credit derivative is the *credit default swap, credit swap* or *default swap*. This is a bilateral contract in which a periodic fixed fee or a one-off premium is paid to a *protection seller*, in return for which the seller will make a payment on the occurrence of a specified credit event. The fee is usually quoted as a basis point multiplier of the nominal value. The swap can refer to a single asset, known as the reference asset or underlying asset, or a basket of assets. The default payment can be paid in whatever way suits the protection buyer or both counterparties. For example it may be linked to the change in price of the reference asset or another specified asset, it may be fixed at a predetermined recovery rate, or it may be in the form of actual delivery of the reference asset at a specified price. The basic structure is illustrated in Figure 7.3.

However it is structured, the credit default swap enables one party to transfer its credit exposure to another party. Banks may use default swaps to trade sovereign and corporate credit spreads without trading the actual assets themselves; for example someone who has gone long a default swap (the protection buyer) will gain if the reference asset obligor suffers a rating downgrade or defaults, and can sell the default swap at a profit if he can find a buyer counterparty. This is because the cost of protection on the reference asset will have increased as a result of the credit event. The original buyer of the default swap need never have owned a bond issued by the reference asset obligor.

The maturity of the credit swap does not have to match the maturity of the reference asset and in most cases does not. On default the swap is terminated and default payment by the protection seller or guarantor is calculated and handed over. The guarantor may have the asset delivered to him and pay the nominal value, or may cash settle the swap contract.

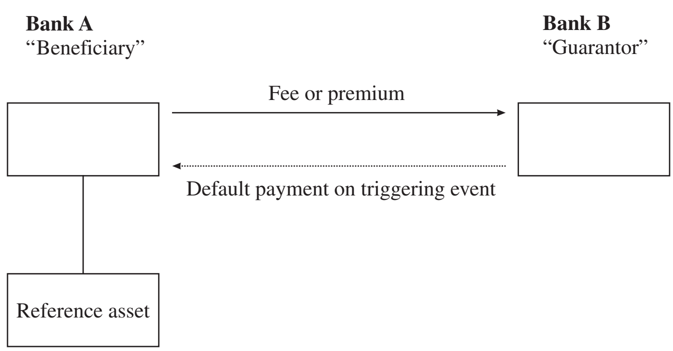

Figure 7.3: Credit default swap

Example 7.1: Credit default swap

XYZ plc credit spreads are currently trading at 120 bps over government for five-year maturities and 195 bps over for 10-year maturities. A portfolio manager hedges a $10 million holding of 10-year paper by purchasing the following credit default swap, written on the five-year bond. This hedge protects for the first five years of the holding, and in the event of XYZ's credit spread widening, will increase in value and may be sold on before expiry at profit. The 10-year bond holding also earns 75 bps over the shorter-term paper for the portfolio manager.

Term	5 years
Reference credit	XYZ plc five-year bond
Credit event	The business day following occurrence of specified credit event
Default payment	Nominal value of bond × (100 – price of bond after credit event)
Swap premium	3.35%

We assume now that midway into the life of the swap there is a technical default on the XYZ plc five-year bond, such that its price now stands at $28. Under the terms of the swap the protection buyer delivers the bond to the seller, who pays out $7.2 million to the buyer.

Credit options

Credit options are also bilateral OTC financial contracts. A credit option is a contract designed to meet specific hedging or speculative requirements of an entity, which may purchase or sell the option to meet its objectives. A credit call option gives the buyer the right without the obligation to purchase the underlying credit-sensitive asset, or a credit spread, at a specified price and specified time (or period of time). A credit put option gives the buyer the right without the obligation to sell the underlying credit-sensitive asset or credit spread. By purchasing credit options banks and other institutions can take a view on credit spread movements for the cost of the option premium only, without recourse to actual loans issued by an obligor. The writer of credit options seeks to earn premium income.

Credit option terms are similar to those used for conventional equity options. A *call* option written on a stock grants the purchaser the right but not the obligation to purchase a specified amount of the stock at a set price and time. A credit option can be used by bond investors to hedge against a decline in the price of specified bonds, in the event of a credit event such as a ratings downgrade. The investor would purchase an option whose payoff profile is a function of the credit quality of the bond, so that a loss on the bond position is offset by the payout from the option.

As with conventional options, there are both vanilla credit options and exotic credit options. The vanilla credit option[4] grants the purchaser the right but not the obligation to buy (or sell if a *put* option) an asset or credit spread at a specified price (the *strike* price) for a specified period of time up to the maturity of the option. A credit option allows a market

[4] Sometimes referred to as the *standard* credit option.

participant to take a view on credit only, and no other exposure such as interest rates. As an example consider an investor who believes that a particular credit spread, which can be that of a specific entity or the average for a sector (such as "all AA-rated sterling corporates"), will widen over the next six months. She can buy a six-month call option on the relevant credit spread, for which a one-off premium (the price of the option) is paid. If the credit spread indeed does widen beyond the strike during the six months, the option will be in-the-money and the investor will gain. If not, the investor's loss is limited to the premium paid.[5]

Exotic credit options are options that have one or more of their parameters changed from the vanilla norm; the same terms are used as in other option markets. Examples include the barrier credit option, which specifies a credit event that would *trigger* (activate) the option or inactivate it. A digital credit option would have a payout profile that would be fixed, irrespective of how much in-the-money it was on expiry, and a zero payout if out-of-the-money.

Credit-linked note

Credit-linked notes (CLNs) exist in a number of forms, but all of them contain a link between the return they pay and the credit-related performance of the underlying asset. A standard credit-linked note is a security, usually issued by an investment-graded entity, that has an interest payment and fixed maturity structure similar to a vanilla bond. The performance of the note however, including the maturity value, is linked to the performance of specified underlying assets as well as that of the issuing entity. The notes are often used by borrowers to hedge against credit risk, and by investors to enhance the yield received on their holdings. Essentially CLNs are hybrid instruments that combine a credit derivative with a vanilla bond. The CLN pays regular coupons, however the credit derivative element is usually set to allow the issuer to decrease the principal amount if a credit event occurs.

For example consider an issuer of credit cards that wants to fund its (credit card) loan portfolio via an issue of debt. In order to reduce the credit risk of the loans, it issues a two-year CLN. The principal amount of the bond is 100% as usual, and it pays a coupon of 7.50%, which is 200 basis points above the two-year benchmark. If however, the incidence of bad debt amongst credit card holders exceeds 10% then the terms state that note holders will only receive back £85 per £100 nominal. The credit card issuer has in effect purchased a credit option that lowers its liability in the event that it suffers from a specified credit event, which in this case is an above-expected incidence of bad debts. The credit card bank has issued the CLN to reduce its credit exposure, in the form of this particular type of credit insurance. If the incidence of bad debts is low, the note is redeemed at par. However if there is a high incidence of such debt, the bank will only have to repay a part of its loan liability.

Investors may wish to purchase the CLN because the coupon paid on it will be above what the credit card bank would pay on a vanilla bond it issued, and higher than other comparable investments in the market. In addition such notes are usually priced below par on issue. Assuming the notes are eventually redeemed at par, investors will also have realised a substantial capital gain.

Further detail on CLNs is given in Choudhry (2003).

---

5 Whether the option is an American or European one will determine whether it can be exercised before its expiry date or on its expiry date only.

The total return swap

A *total return swap* (TRS), sometimes known as a *total rate of return swap* or *TR swap*, is an agreement between two parties that exchanges the total return from a financial asset between them. This is designed to transfer the credit risk from one party to the other. It is one of the principal instruments used by banks and other financial instruments to manage their credit risk exposure, and as such is a credit derivative. One definition of a TRS is given in Francis *et al.* (1999), which states that a TRS is a swap agreement in which the *total return* of a bank loan or credit-sensitive security is exchanged for some other cash flow, usually tied to LIBOR or some other loan or credit-sensitive security.

In some versions of a TRS the actual underlying asset is actually sold to the counterparty, with a corresponding swap transaction agreed alongside; in other versions there is no physical change of ownership of the underlying asset. The TRS trade itself can be to any maturity term, that is, it need not match the maturity of the underlying security. In a TRS the total return from the underlying asset is paid over to the counterparty in return for a fixed or floating cash flow. This makes it slightly different to other credit derivatives, as the payments between counterparties to a TRS are connected to changes in the market value of the underlying asset, as well as changes resulting from the occurrence of a credit event.

Figure 7.4 illustrates a generic TR swap. The two counterparties are labelled as banks, but the party termed "Bank A" can be another financial institution, including cash-rich fixed-income portfolio managers such as insurance companies and hedge funds. In Figure 7.4 Bank A has contracted to pay the "total return" on a specified reference asset, while simultaneously receiving a LIBOR-based return from Bank B. The reference or underlying asset can be a bank loan such as a corporate loan or a sovereign or corporate bond. The total return payments from Bank A include the interest payments on the underlying loan as well as any appreciation in the market value of the asset. Bank B will pay the LIBOR-based return; it will also pay any difference if there is a depreciation in the price of the asset. The economic effect is as if Bank B owned the underlying asset, as such TR swaps are synthetic loans or securities. A significant feature is that Bank A will usually hold the underlying asset on its balance sheet, so that if this asset was originally on Bank B's balance sheet, this is a means by which the latter can have the asset removed from its balance sheet for the term of

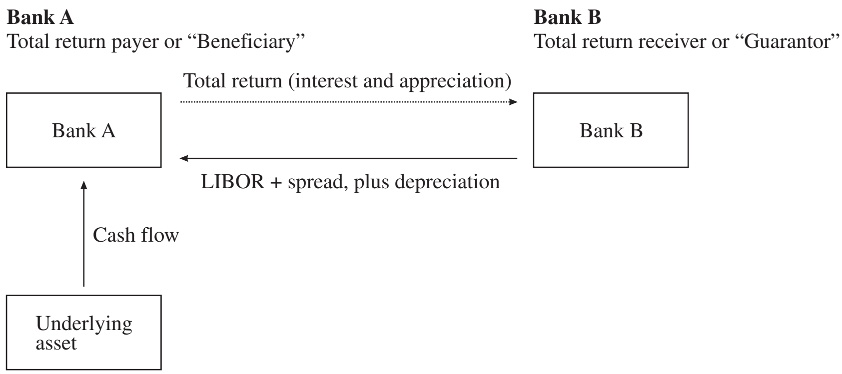

Figure 7.4: Total return swap

the TR swap.[6] If we assume Bank A has access to LIBOR funding, it will receive a spread on this from Bank B. Under the terms of the swap, Bank B will pay the difference between the initial market value and any depreciation, so it is sometimes termed the “guarantor” while Bank A is the “beneficiary”.

The total return on the underlying asset is the interest payments and any change in the market value if there is capital appreciation. The value of an appreciation may be cash settled, or alternatively there may be physical delivery of the reference asset on maturity of the swap, in return for a payment of the initial asset value by the total return “receiver”. The maturity of the TR swap need not be identical to that of the reference asset, and in fact it is rare for it to be so.

The swap element of the trade will usually pay on a quarterly or semi-annual basis, with the underlying asset being revalued or *marked-to-market* on the re-fixing dates. The asset price is usually obtained from an independent third party source such as Bloomberg or Reuters, or as the average of a range of market quotes. If the *obligor* of the reference asset defaults, the swap may be terminated immediately, with a net present value payment changing hands according to what this value is, or it may be continued with each party making appreciation or depreciation payments as appropriate. This second option is only available if there is a market for the asset, which is unlikely in the case of a bank loan. If the swap is terminated, each counterparty will be liable to the other for accrued interest plus any appreciation or depreciation of the asset. Commonly under the terms of the trade, the guarantor bank has the option to purchase the underlying asset from the beneficiary bank, and then to deal directly with the loan defaulter.

There are a number of reasons why portfolio managers may wish to enter into TR swap arrangements. One of these is to reduce or remove credit risk. Using TR swaps as a credit derivative instrument, a party can remove exposure to an asset without having to sell it. In a vanilla TR swap the total return payer retains rights to the reference asset, although in some cases servicing and voting rights may be transferred. The total return receiver gains an exposure to the reference asset without having to pay out the cash proceeds that would be required to purchase it. As the maturity of the swap rarely matches that of the asset, the swap receiver may gain from the positive funding or *carry* that derives from being able to roll over short-term funding of a longer-term asset.[7] The total return payer on the other hand benefits from protection against market and credit risk for a specified period of time, without having to liquidate the asset itself. On maturity of the swap the total return payer may reinvest the asset if it continues to own it, or it may sell the asset in the open market. Thus the instrument may be considered a synthetic repo. A TR swap agreement entered into as a credit derivative is a means by which banks can take on unfunded off-balance sheet credit exposure. Higher-rated banks that have access to LIBID funding can benefit by funding on-balance sheet assets that are credit protected through a credit derivative such as a TR swap, assuming the net spread of asset income over credit protection premium is positive.

---

[6] Although it is common for the receiver of the LIBOR-based payments to have the reference asset on its balance sheet, this is not always the case.

[7] This assumes a positively-sloping yield curve.

A TR swap conducted as a synthetic repo is usually undertaken to effect the temporary removal of assets from the balance sheet. This may be desired for a number of reasons, for example if the institution is due to be analysed by credit rating agencies or if the annual external audit is due shortly. Another reason a bank may wish to temporarily remove lower credit-quality assets from its balance sheet is if it is in danger of breaching capital limits in between the quarterly return periods. In this case, as the return period approaches, lower-quality assets may be removed from the balance sheet by means of a TR swap, which is set to mature after the return period has passed.

Banks have employed a number of methods to price credit derivatives and TR swaps. Space does not permit an in-depth discussion of these techniques here. Essentially the pricing of credit derivatives is linked to that of other instruments; however the main difference between credit derivatives and other off-balance sheet products such as equity, currency or bond derivatives is that the latter can be priced and hedged with reference to the underlying asset, which can be problematic when applied to credit derivatives. Credit products pricing uses statistical data on likelihood of default, probability of payout, level of risk tolerance and a pricing model. The basic ingredients of a TR swap are that one party "funds" an underlying asset and transfers the total return of the asset to another party, in return for a (usually) floating return that is a spread to LIBOR. This spread is a function of:

- the credit rating of the swap counterparty;
- the amount and value of the reference asset;
- the credit quality of the reference asset;
- the funding costs of the beneficiary bank;
- any required profit margin;
- the capital charge associated with the TR swap.

The TR swap counterparties must consider a number of risk factors associated with the transaction, which include:

- the probability that the TR beneficiary may default while the reference asset has declined in value;
- the reference asset obligor defaults, followed by default of the TR swap receiver before payment of the depreciation has been made to the payer or "provider".

The first risk measure is a function of the probability of default by the TR swap receiver and the market volatility of the reference asset, while the second risk is related to the joint probability of default of both factors as well as the recovery probability of the asset.

Applications for portfolio managers

Applications overview

Credit derivatives have allowed market participants to separate and disaggregate credit risk, and thence to trade this risk in a secondary market (see for example Das, (2000), Chapters 2–4). Initially portfolio managers used them to reduce credit exposure; subsequently they have been used in the management of portfolios, to enhance portfolio yields and in the structuring of synthetic collateralised debt obligations. We summarise portfolio managers' main uses of credit derivatives below.

Enhancing portfolio returns

Asset managers can derive premium income by trading credit exposures in the form of derivatives issued with synthetic structured notes. The multi-tranching aspect of structured products enables specific credit exposures (credit spreads and outright default), and their expectations, to be sold to specific areas of demand. By using structured notes such as credit-linked notes, tied to the assets in the reference pool of the portfolio manager, the trading of credit exposures is crystallised as added yield on the asset manager's fixed income portfolio. In this way the portfolio manager has enabled other market participants to gain an exposure to the credit risk of a pool of assets but not to any other aspects of the portfolio, and without the need to hold the assets themselves.

Reducing credit exposure

Consider a portfolio manager that holds a large portfolio of bonds issued by a particular sector (say, utilities) and believes that spreads in this sector will widen in the short term. Previously, in order to reduce its credit exposure it would have to sell bonds, however this may crystallise a mark-to-market loss and may conflict with its long-term investment strategy. An alternative approach would be to enter into a credit default swap, purchasing protection for the short term; if spreads do widen these swaps will increase in value and may be sold at a profit in the secondary market. Alternatively the portfolio manager may enter into total return swaps on the desired credits. It pays the counterparty the total return on the reference assets, in return for LIBOR. This transfers the credit exposure of the bonds to the counterparty for the term of the swap, in return for the credit exposure of the counterparty.

Consider now the case of a portfolio manager wishing to mitigate credit risk from a growing portfolio (say, one that has just been launched). Figure 7.5 shows an example of an unhedged credit exposure to a hypothetical credit-risky portfolio. It illustrates the manager's expectation of credit risk building up to $250 million as the portfolio is ramped up, and then reducing to a more stable level as the credits become more established. A three-year credit default swap entered into shortly after provides protection on half of the notional exposure, shown as the broken line. The net exposure to credit events has been reduced by a significant margin.

Credit switches and zero-cost credit exposure

Protection buyers utilising credit default swaps must pay premium in return for laying off their credit risk exposure. An alternative approach for an asset manager involves the use of credit switches for specific sectors of the portfolio. In a credit switch the portfolio manager purchases credit protection on one reference asset or pool of assets, and simultaneously sells protection on another asset or pool of assets.[8] So for example, the portfolio manager would purchase protection for a particular fund and sell protection on another. Typically the entire transaction would be undertaken with one investment bank, which would price the structure so that the net cash flows would be zero. This has the effect of synthetically diversifying the credit exposure of the portfolio manager, enabling it to gain and/or reduce exposure to sectors it desires.

8 A pool of assets would be concentrated on one sector, such as utility company bonds.

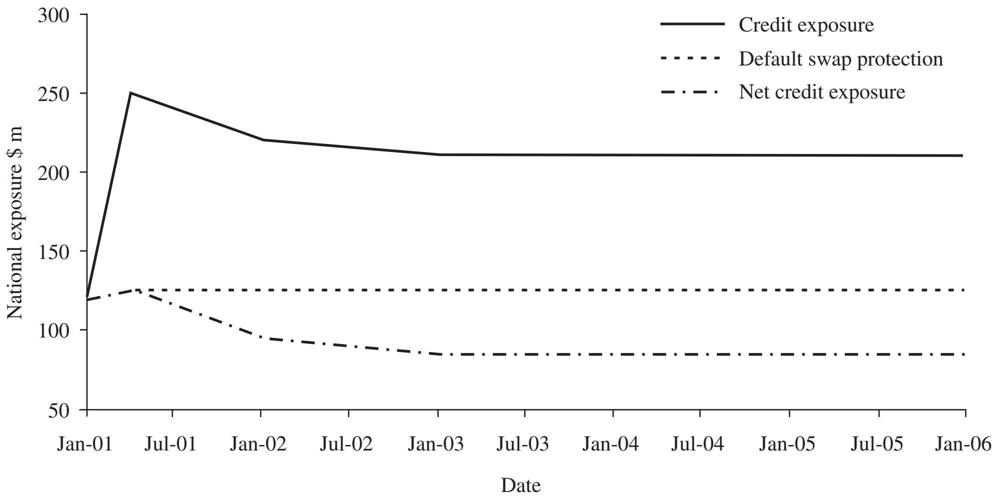

Figure 7.5: Reducing credit exposure

Exposure to market sectors

Investors can use credit derivatives to gain exposure to sectors for which they do not wish a cash market exposure. This can be achieved with an *index* swap, which is similar to a TR swap, with one counterparty paying a total return that is linked to a external reference index. The other party pays a LIBOR-linked coupon or the total return of another index. Indices that are used might include the government bond index, a high-yield index or a technology stocks index. Assume that an investor believes that the bank loan market will outperform the mortgage-backed bond sector; to reflect this view the investor enters into an index swap in which he pays the total return of the mortgage index and receives the total return of the bank loan index.

Another possibility is synthetic exposure to foreign currency and money markets. Again we assume that an investor has a particular view on an emerging market currency. If he wishes he can purchase a short-term (say one-year) domestic coupon-bearing note, whose principal redemption is linked to a currency factor. This factor is based on the ratio of the spot value of the foreign currency on issue of the note to the value on maturity. Such currency-linked notes can also be structured so that they provide an exposure to sovereign credit risk. The downside of currency-linked notes is that if the exchange rate goes the other way, the note will have a zero return, in effect a negative return once the investor's funding costs have been taken into account.

Credit spreads

Credit derivatives can be used to trade credit spreads. Assume that an investor has negative views on a certain emerging market government bond credit spread relative to UK gilts. The simplest way to reflect this view would be to go long a credit default swap on the sovereign, paying X basis points. Assuming that the investor's view is correct and the sovereign bonds decrease in price as their credit spread widens, the premium payable on the credit swap will increase. The investor's swap can then be sold into the market at this higher premium.

Application of total return swaps

Capital structure arbitrage A capital structure arbitrage describes an arrangement whereby investors exploit mispricing between the yields received on two different loans by the same issuer. Assume that the reference entity has both a commercial bank loan and a subordinated bond issue outstanding, but that the former pays LIBOR plus 330 basis points while the latter pays LIBOR plus 230 basis points. An investor enters into a total return swap in which it effectively is purchasing the bank loan and selling short the bond. The nominal amounts will be at a ratio, for argument's sake let us say 2:1, as the bonds will be more price-sensitive to changes in credit status than the loans.

The trade is illustrated in Figure 7.6. The investor receives the "total return" on the bank loan, while simultaneously paying the return on the bond in addition to LIBOR plus 30 basis points, which is the price of the TR swap. The swap generates a net spread of 175 basis points, given by $[(100\,\text{bps} \times \frac{1}{2}) + (250\,\text{bps} \times \frac{1}{2})]$.

Synthetic repo A portfolio manager believes that a particular bond that it does not hold is about to decline in price. To reflect this view the portfolio manager may do one of the following:

- *Sell the bond in the market and cover the resulting short position in repo*: The cash flow out is the coupon on the bond, with capital gain if the bond falls in price. Assume that the repo rate is floating, say LIBOR plus a spread. The manager must be aware of the funding costs of the trade, so that unless the bond can be covered in repo at *general collateral* rates,[9] the funding will be at a loss. The yield on the bond must also be lower than the LIBOR plus spread received in the repo.
- *As an alternative, enter into a TR swap*: The portfolio manager pays the total return on the bond and receives LIBOR plus a spread. If the bond yield exceeds the LIBOR spread, the funding will be negative, however the trade will gain if the trader's view is proved

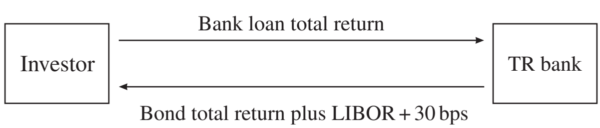

Figure 7.6: Total return swap in capital structure arbitrage

[9] That is, the bond cannot be *special*. A bond is special when the repo rate payable on it is significantly (say, 20–30 basis points or more) below the *general collateral* repo rate, so that covering a short position in the bond entails paying a substantial funding premium.

correct and the bond falls in price by a sufficient amount. If the breakeven funding cost (which the bond must exceed as it falls in value) is lower in the TR swap, this method will be used rather than the repo approach. This is more likely if the bond is special.

Overview of TR swap applications Total return swaps are increasingly used as synthetic repo instruments, most commonly by investors that wish to purchase the credit exposure of an asset without purchasing the asset itself. This is conceptually similar to what happened when interest rate swaps were introduced, which enabled banks and other financial institutions to trade interest rate risk without borrowing or lending cash funds.

Under a TR swap an asset such as a bond position may be removed from the balance sheet. In order to avoid adverse impact on regular internal and external capital and credit exposure reporting a bank may use TR swaps to reduce the amount of lower-quality assets on the balance sheet. This can be done by entering into a short-term TR swap with say, a two-week term that straddles the reporting date. Bonds are removed from the balance sheet if they are part of a sale plus TR swap transaction. This is because legally the bank selling the asset is not required to repurchase bonds from the swap counterparty, nor is the total return payer obliged to sell the bonds back to the counterparty (or indeed sell the bonds at all on maturity of the TR swap).

Pricing

There are a number of approaches to the pricing of credit derivatives. A good overview of the main methods is contained in Das (2000). Here we introduce asset swap pricing.

Asset swap pricing

Credit derivatives are commonly valued using the asset swap pricing technique. In addition to its use by dealers, risk management departments who wish to independently price such swaps also adopt this technique. The asset swap market is a reasonably reliable indicator of the returns required for individual credit exposures, and provides a mark-to-market framework for reference assets as well as a hedging mechanism.

A par asset swap typically combines the sale of an asset such as a fixed-rate corporate bond to a counterparty, at par and with no interest accrued, with an interest rate swap. The coupon on the bond is paid in return for LIBOR, plus a spread if necessary. This spread is the asset swap spread and is the price of the asset swap. In effect the asset swap allows market participants that pay LIBOR-based funding to receive the asset swap spread. This spread is a function of the credit risk of the underlying bond asset, which is why it in effect becomes the cornerstone of the price payable on a credit default swap written on that reference asset.

The generic pricing is given by equation (7.1):

$$Y_a = Y_b - ir \tag{7.1}$$

where

Y_a is the asset swap spread;
Y_b is the asset spread over the benchmark;
ir is the interest rate swap spread.

The asset spread over the benchmark is simply the bond (asset) redemption yield over that of the government benchmark. The interest rate swap spread reflects the cost involved in converting fixed-coupon benchmark bonds into a floating-rate coupon during the life of the asset (or default swap), and is based on the swap rate for that maturity.

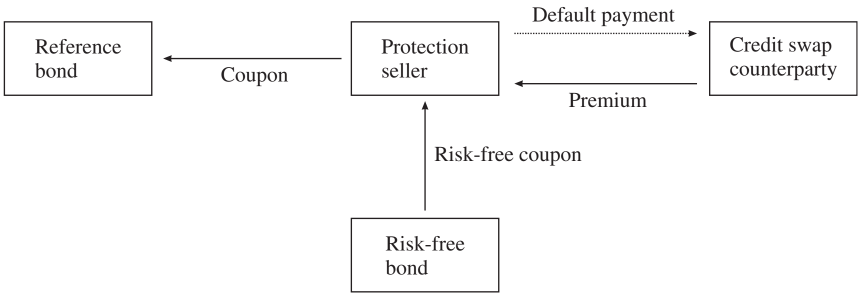

Figure 7.7: Credit default swap and asset swap hedge

Asset swap pricing example

XYZ plc is a Baa2-rated corporate. The seven-year asset swap for this entity is currently trading at 93 basis points; the underlying seven-year bond is hedged by an interest rate swap with an Aa2-rated bank. The risk-free rate for floating-rate bonds is LIBID minus 12.5 basis points (assume the bid–offer spread is 6 basis points). This suggests that the credit spread for XYZ plc is 111.5 basis points. The credit spread is the return required by an investor for holding the credit of XYZ plc. The protection seller is conceptually long the asset, and so would short the asset as a hedge of its position. This is illustrated in Figure 7.7. The price charged for the default swap is the price of shorting the asset, which works out as 111.5 basis points each year.

Therefore we can price a credit default written on XYZ plc as the present value of 111.5 basis points for seven years, discounted at the interest rate swap rate of 5.875%. This computes to a credit swap price of 6.25%.

Reference	XYZ plc
Term	7 years
Interest rate swap rate	5.875%
Asset swap	LIBOR plus 93 bps
Default swap pricing:	
Benchmark rate	LIBID minus 12.5 bps
Margin	6 bps
Credit default swap	111.5 bps
Default swap price	6.252%

Modelling credit spreads

Practitioners increasingly model credit risk as they do interest rates and use spread models to price associated derivatives. For example, the multi-factor Heath–Jarrow–Morton approach has been used, modelling interest rate risk, default risk and recovery risk.[10] However, spread models do not split the spread elements into default risk and recovery risk, but

[10] This is the rate of recovery on a defaulted loan. Bonds rarely descend to 0 in price in the event of default; there is always assumed to be some residual recovery value still available.

model the spread as a whole. For example Das and Tufano (1996) noted that equation (7.2) may be used to model credit spreads:

$$ds = k(\theta - s)dt + \sigma\sqrt{s}dZ \tag{7.2}$$

where

- s is the credit spread;
- k is the mean reversion rate;
- θ is the mean of the spread;
- σ is the volatility.

The standard Brownian motion or Weiner process is indicated by dZ. Under this approach it would be possible to price credit options in the same way as other option products, modelling the credit spread rather than say, the interest rate. An excellent introduction to modelling methods for pricing credit derivatives is given in Francis *et al.* (1999).

7.4 An introduction to collateralised debt obligations (CDOs)

Credit derivatives are the principal instrument employed in the structuring of what are known as synthetic collateralised debt obligations (CDOs), and are used to transfer credit risk from a portfolio of assets from one party to the other. Either or none of the parties may or may not be the beneficial owner of the assets, which can be bonds, loans or a mix of both. In this section we introduce CDOs, which are part of the family of debt market structured financial products, before considering briefly the motivation behind the structuring of synthetic CDOs.

A CDO is essentially a securitisation in which a special purpose vehicle or SPV issues bonds or notes against an investment in a diversified pool of assets. These assets can be bonds, loans such as commercial bank loans or a mixture of both bonds and loans. Where assets are bonds, these are usually high-yield bonds that provide a spread of interest over the interest liability of the issued notes; where the assets are loans, the CDO acts as a mechanism by which illiquid loans can be pooled into a marketable security or securities. The third type of CDO is known as a synthetic CDO and refers to a structure in which credit derivatives are used to construct the underlying pool of assets.

The investments are funded through the issue of the notes, and interest and principal payments on these notes are linked to the performance of the underlying assets. These underlying assets act as the *collateral* for the issues' notes, hence the name. The key difference between CDOs and asset-backed securities and multi-asset repackaged securities is that the collateral pool is more actively managed by a portfolio or collateral manager. Generally CDO feature a multi-tranche structure, with a number of issued securities, most or all of which are rated by a ratings agency. The priority of payment of the issued securities reflects the credit rating for each note, with the most senior note being the highest rated. The term *waterfall* is used to refer to the order of payments; sufficient underlying cash flows must be generated by the issuing vehicle in order to meet the fees of third-party servicers and all the note issue liabilities. In Europe issued securities may pay a fixed or floating coupon, usually on a semi-annual, quarterly or monthly basis,[11] with senior notes issued from AAA to A and junior and mezzanine notes rated BBB to B. There may be unrated subordinated and equity pieces issued. Investors in the subordinated notes receive coupon after payment of servicing fees and the coupon on senior notes. The equity and subordinated note are the first loss pieces and, as they carry the highest risk, have a higher expected return compared to that of the underlying collateral.

There are two types of CDO, *collateralised bond obligations* (CBOs) and *collateralised loan obligations* (CLOs). As the names suggest, the primary difference between each type is the nature of the underlying assets; a CBO will be collateralised by a portfolio of bonds while a CLO will represent an underlying pool of bank loans. CDOs have also been collateralised by credit derivatives and credit-linked notes. Following this distinction, CDOs can be broken into two main types, *balance sheet* CDOs and *arbitrage* CDOs. Balance sheet CDOs are most akin to a traditional securitisation; they are created to remove assets from the balance sheet of the originating bank or financial institution, usually to reduce capital requirements, increase return on capital or free up lending lines. An arbitrage CDO is created when the originator, who may be a bank or fund manager for instance, wishes to exploit the yield differential between the underlying assets and the overlying notes. This may be achieved by active management of the underlying portfolio, which might consist of high-yielding or emerging market bonds. Arbitrage CDOs are categorised further into *cash flow* and *market value* CDOs. Almost invariably balance sheet CDOs are cash flow transactions. Put simply a cash flow CDO is one in which the underlying collateral generates sufficient cash flow to pay the principal and interest on the issued notes, as well as the servicing fees of third-party agents. In a market value CDO, the collateral manager actively runs the portfolio and, by means of this trading activity, generates sufficient returns to pay the CDO obligations.

Banks and financial institutions use CDOs to diversify their sources of funding, to manage portfolio risk and to obtain regulatory capital relief. Investors are attracted to the senior notes in a transaction because these allow them to earn relatively high yields compared to other asset-backed bonds of a similar credit rating. Other advantages include:

- exposure to a diversified range of credits;
- access to the fund management and credit analysis skills of the portfolio manager;
- generally, a lower level of uncertainty and risk exposure compared to a single bond of similar rating.

A good overview introduction to CDOs is given in Fabozzi and Goodman (2000). We show a "family tree" of CDOs in Figure 7.8. The most common CDOs are balance sheet deals; however a later development, the *synthetic* CDO, now accounts for a large number of transactions.

[11] Hence proving once and for all that Eurobonds, defined as international securities issued by a syndicate of banks and clearing in Euroclear and Clearstream, may pay coupon on frequencies other than an annual basis!

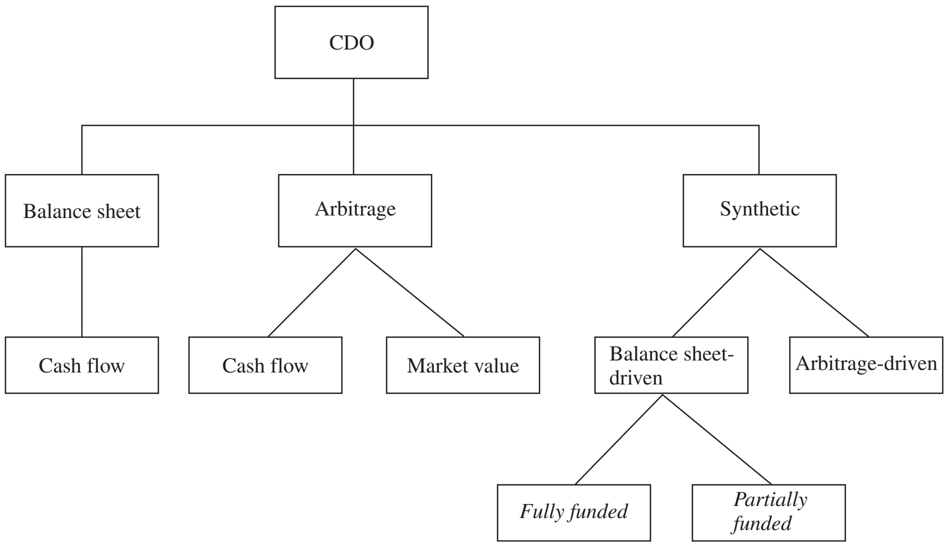

Figure 7.8: The CDO family

7.4.1 *Cash flow CDO*

These are intuitively the easiest to understand and are similar to other asset-backed securitisation involving an SPV. In essence assets such as bonds or loans are pooled together and the cash flows from these assets used to back the obligations of the issued notes. As the underlying assets are sold to the SPV, they are removed from the originator's balance sheet; hence the credit risk associated with these assets is transferred to the holders of the issued notes. The originator also obtains funding by issuing the notes. The structure is illustrated, admittedly in simple fashion, in Figure 7.9.

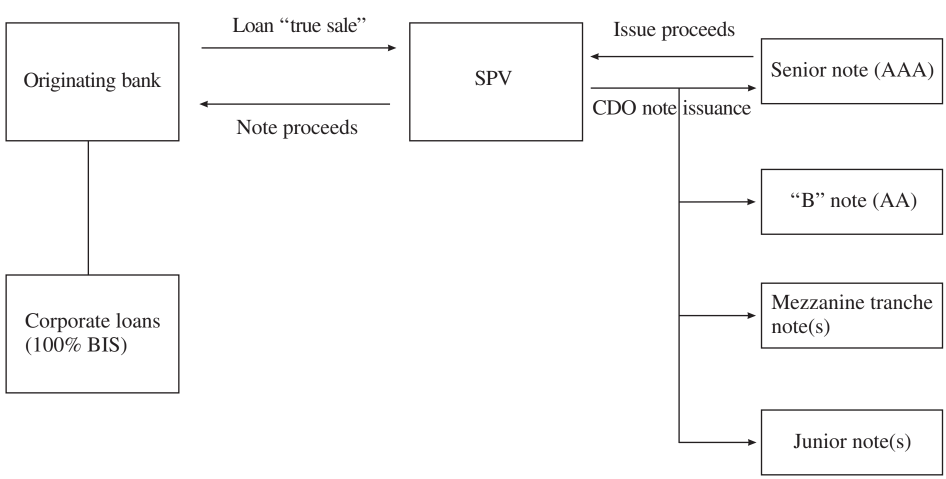

Figure 7.9: Cash flow CDO

7.4.2 *Synthetic CDO*

A synthetic CDO is so-called because the transfer of credit risk is achieved "synthetically" through the sale of a credit derivative, rather than a "true sale" to the SPV. Thus in a synthetic CDO the credit risk of the underlying loans or bonds is transferred to the SPV using credit default swaps and/or total return swaps (TRS). However the assets themselves are not legally transferred to the SPV, and they remain on the originator's balance sheet. Using a synthetic CDO, the originator can obtain regulatory capital relief and manage the credit risk on its balance sheet, but will not be receiving any funding. A typical structure is shown in Figure 7.10.[12]

Synthetic CDOs are popular in the European market because they are simpler and can be brought to the market quickly; in practice, in certain countries the legal infrastructure has not been sufficiently developed to enable true sale securitisation to be undertaken. In addition when the underlying asset pool is composed of bonds from different countries, a cash-funded CDO may present too many administrative difficulties. A synthetic CDO removes such issues by using credit derivatives, and in theory can be brought to market more quickly than a cash flow arrangement (although in practice this is not always the case). Synthetic CDOs have been issued as arbitrage CDOs or balance sheet CDOs. We briefly describe each type.

A synthetic arbitrage CDO is originated generally by collateral managers who wish to exploit the difference in yield between that obtained on the underlying assets and that payable on the CDO, both in note interest and servicing fees. The generic structure is as follows: the specially-created SPV enters into a TRS with the originating bank or financial institution, referencing the bank's underlying portfolio (the *reference portfolio*). The portfolio is actively managed and is funded on the balance sheet by the originating bank. The SPV receives the "total return" from the reference portfolio, and in return it pays LIBOR plus a spread to the originating bank. The SPV also issues notes that are sold into the market to

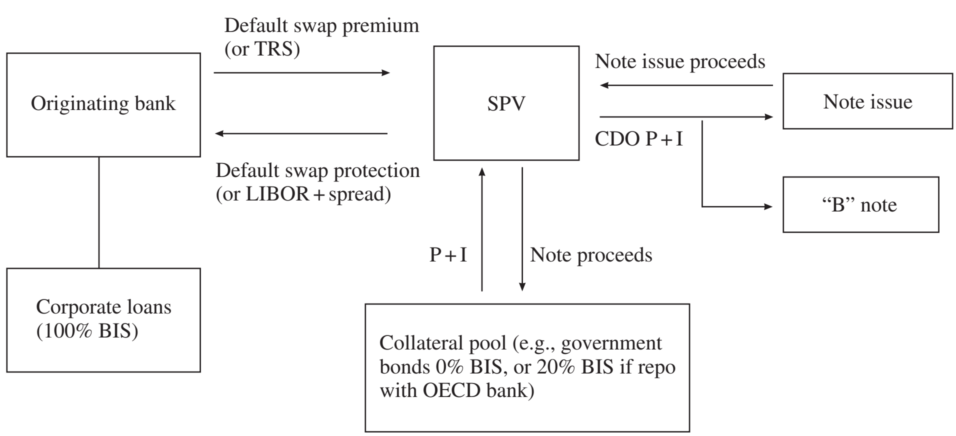

12 "P + I" is a common abbreviation for note (or loan) principal and interest.

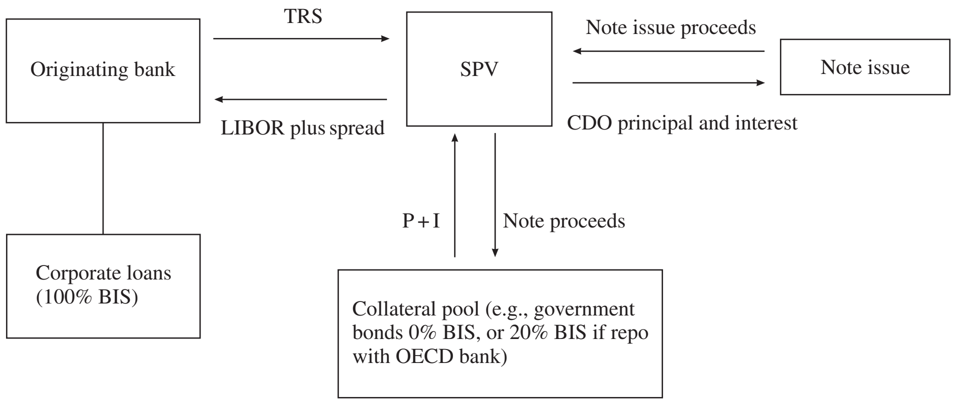

Figure 7.10: Synthetic CDO structure

CDO investors, and these notes can be rated as high as AAA as they are backed by high-quality collateral, which is purchased using the note proceeds. A typical structure is shown in Figure 7.11.

Balance sheet synthetic CDOs are employed by banks that wish to manage regulatory capital. As before, the underlying assets are bonds, loans and credit facilities originated by the issuing bank. In a balance sheet CDP the SPV enters into a credit default swap agreement with the originator, again with the specific collateral pool designated as the reference portfolio. The SPV receives the premium payable on the default swap, and thereby provides credit protection on the reference portfolio. There are two types of CDO within this structure. In a *partially funded* CDO, only the highest credit risk segment of the portfolio is transferred. The cash flow that would be needed to service the synthetic CDO overlying liability is received from the AAA-rated collateral that is purchased by the SPV with the proceeds of the note issue. An originating bank obtains maximum regulatory capital relief by means of a partially funded structure, through a combination of the synthetic CDO and what is known as a *super senior swap* arrangement with an OECD banking counterparty. A super senior swap provides additional protection to that part of the portfolio, the senior segment, that is already protected by the funded portion of the transaction.

A generic partially funded transaction is shown in Figure 7.12. It shows an arrangement whereby the issuer enters into two credit default swaps; the first with an SPV that provides protection for losses up to a specified amount of the reference pool,[13] while the second swap is set up with the OECD bank or, occasionally, an insurance company.[14]

A *fully funded* CDO is a structure where the credit risk of the entire portfolio is transferred to the SPV via a credit default swap. In a fully funded (or just "funded") synthetic CDO the issuer enters into the credit default swap with the SPV, which itself issues notes to the value of the assets on which the risk has been transferred. The proceeds from the notes are invested in risk-free government or agency debt such as gilts, bunds or Pfandbriefe,

13 In practice, to date this portion has been between 5% and 15% of the reference pool.

14 An "OECD" bank, thus guaranteeing a 20% risk weighting for capital ratio purposes, under Basle I rules.

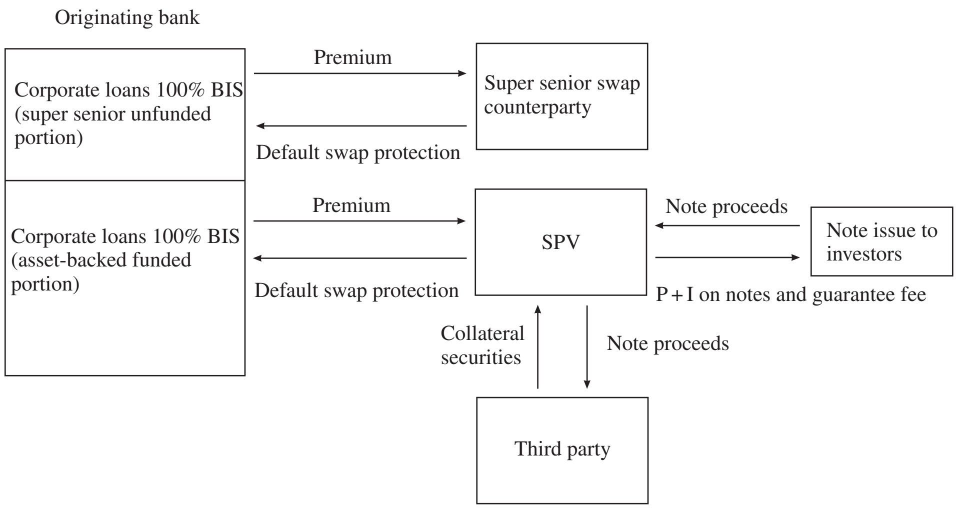

Figure 7.11: Synthetic arbitrage CDO structure

or in senior unsecured bank debt. Should there be a default on one or more of the underlying assets, the required amount of the collateral is sold and the proceeds from the sale paid to the issuer to recompense for the losses. The premium paid on the credit default swap must be sufficiently high to ensure that it covers the difference in yield between that on the collateral and that on the notes issued by the SPV. The generic structure is illustrated in Figure 7.13.

Fully funded CDOs are relatively uncommon. One of the advantages of the partially funded arrangement is that the issuer will pay a lower premium compared to a fully funded synthetic CDO, because it is not required to pay the difference between the yield on the collateral and the coupon on the note issue (the unfunded part of the transaction). The downside is that the issuer will receive a reduction in risk weighting for capital purposes to 20% for the risk transferred via the super senior default swap.

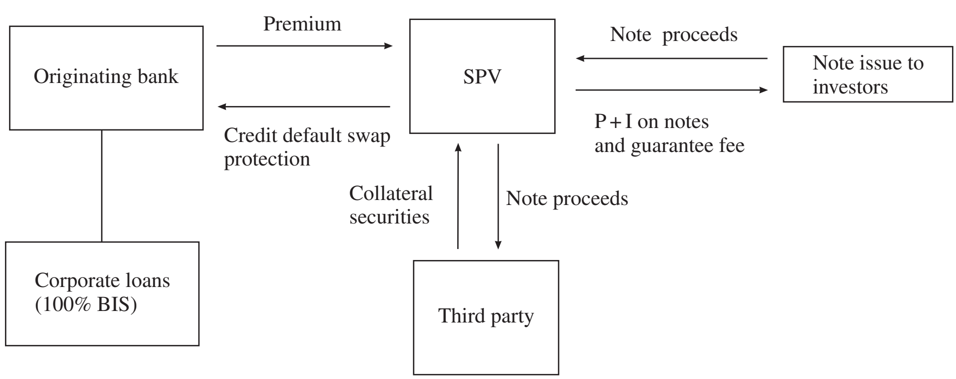

Figure 7.12: Partially funded synthetic CDO structure

Within the European market synthetic balance sheet CDOs are the most common structure. The reasons that banks originate them are two-fold:

- *capital relief*: banks can obtain regulatory capital relief by transferring lower-yield corporate credit risk such as corporate bank loans off their balance sheet. Under Basle I rules all corporate debt carries an identical 100% risk weighting; therefore with banks having to assign 8% of capital for such loans, higher-rated (and hence lower-yielding) corporate assets will require the same amount of capital but will generate a lower return on that capital. A bank may wish to transfer such higher-rated, lower-yield assets from its balance sheet, and this can be achieved using a CDO transaction. The capital requirements for a synthetic CDO are lower than for corporate assets; for example the funded segment of the deal will be supported by high-quality collateral such as government bonds, and through a repo arrangement with an OECD bank which would carry a 20% risk weighting, as does the super senior element;
- *transfer of credit risk*: the cost of servicing a fully funded CDO, and the premium payable on the associated credit default swap, can be prohibitive. With a partially funded structure, the issue amount is typically a relatively small share of the asset portfolio. This lowers substantially the default swap premium. Also, as the CDO investors suffer the first loss element of the portfolio, the super senior default swap can be entered into at a considerably lower cost than that on a fully funded CDO.

In essence synthetic CDOs represent relative ease of transaction and legal documentation requirements, which has been behind their popularity in the European market.

Selected bibliography and references

Choudhry, M., *Credit Derivatives and Synthetic Securitisation*, Wiley 2003
Das, S., *Credit Derivatives and Credit Linked Notes*, 2nd edition, Wiley 2000
Das, S., Tufano, P., "Pricing credit sensitive debt when interest rates, credit ratings and credit spreads are stochastic", *Journal of Financial Engineering*, 1996, pp. 161–198
Fabozzi, F., Goodman, L., *Investing in Collateralised Debt Obligation*, FJF Associates 2000
Francis, J., Frost, J., Whittakes, J., *The Handbook of Credit Derivatives*, McGraw-Hill 1999
Kasapi, A., *Mastering Credit Derivatives*, FT Prentice Hall 1999

Appendix 7.1: Credit ratings

FitchIBCA	Moody's	S&P	Summary description
Investment grade			
AAA	Aaa	AAA	Gilt edged, prime, maximum safety, lowest risk, and when sovereign borrower considered "default-free"
AA+	Aa1	AA+	
AA	Aa2	AA	High grade; high credit quality
AA−	Aa3	AA−	
A+	A1	A+	
A	A2	A	Upper-medium grade
A−	A3	A−	
BBB+	Baa1	BBB+	
BBB	Baa2	BBB	Lower-medium grade
BBB−	Baa3	BBB−	
Speculative grade			
BB+	Ba1	BB+	
BB	Ba2	BB	Low grade; speculative
BB−	Ba3	BB−	
B+	B1		
B	B	B	Highly speculative
B−	B3		
Predominantly speculative, substantial risk or in default			
CCC+		CCC+	
CCC	Caa	CCC	Substantial risk, in poor standing
CC	Ca	CC	May be in default, very speculative
C	C	C	Extremely speculative
		CI	Income bonds – no interest being paid
DDD			
DD			Default
D		D	

8 Equity Futures Contracts

8.1 Exchange-traded equity index and universal stock futures

Futures contracts represent an agreement between two parties to undertake a transaction at some agreed future date at a price agreed now. These contracts are exchange-based instruments and, as such, contracts are tightly specified so that market participants know exactly what they are agreeing to buy or sell. The contracts have standard delivery dates and the determination of a contract's value at a point in time is normally defined as an exchange specified multiplier times the current market quote for the futures contract. Table 8.1 illustrates a typical contract specification from LIFFE. Note that the multiplier that determines the price of a contract is defined as £10. The delivery months are clearly laid out as the cycle March, June, September and December. The delivery and last trading day (LTD) are specified to be the first business day after the LTD and the third Friday in the delivery month, respectively. Of course, the contract makes provision for the eventuality that the last trading and delivery days fall on public holidays. These contracts also define the smallest amount by which quotes can move, which is called a tick, and attach a value to those minimum price movements. In the case of the LIFFE FTSE 100 quote the tick size is 0.5 of an index point, in other words 0.5(£10.00) = £5.00.

Contract size	Valued at £10 per index point (e.g. value £65,000 at 6500.0)
Delivery months	March, June, September, December (nearest three available for trading)
Last trading day	10:30:30 (London time) third Friday in delivery month[a]
Delivery day	First business day after the expiry date
Quotation	Index points (e.g. 6500.0)
Minimum price movement (tick size and value)	0.5 (£5.00)
LIFFE CONNECT™ trading hours	08:00–17:30 (London time)
Daily settlement time	16:30 (London time)
Contract standard	Cash settlement based on the exchange delivery settlement price

Exchange delivery settlement price (EDSP): The EDSP is based on the average values of the FTSE 100 index every 15 seconds between (and including) 10:10 and 10:30 (London time) on the last trading day. Of the 81 measured values, the highest 12 and lowest 12 will be discarded and the remaining 57 will be averaged to calculate the EDSP. Where necessary, the calculation will be rounded to the nearest half index point.

Table 8.1: FTSE 100 index future

[a] In the event of the third Friday not being a business day, the last trading day shall normally be the last business day preceding the third Friday.

Source: LIFFE

The contract illustrated above is the index futures contract in respect of the FTSE 100 equity index. With the advent of Euroland several other indexes have appeared. These indexes are proving popular with equity fund managers wishing to take advantage of European-wide diversified portfolios. One such index is LIFFE's Eurotop 100, another is the MSCI index; there are many others. To get an appreciation of the number of indexes available the reader is encouraged to use LIFFE's web page – www.liffe.com – where the full range of contract specifications for equity (and other instruments) is freely available. For global coverage the site www.site-by-site.com provides convenient access to the world's major exchanges.

A typical European index futures contract – the Eurotop 100 – is presented in Table 8.2. Note that, once again, the multiplier is defined in the contract in order to determine the contract's value. In this case the multiplier is defined in euros and not sterling. The minimum price movement allowed here is 0.5 of an index point so, since the Eurotop multiplier is €20 the smallest price movement allowed for trading purposes is 0.5 (€20.00) = €10.00.

An innovation as far as the exchanges are concerned is the introduction of universal stock futures (USFs) also known as single stock futures (SSFs). These derivatives are a response to the popularity of an over-the-counter derivative "contracts for differences" (CFDs), about which more will be said later. LIFFE offers universal stock futures on US, European and UK equities denominated in euros, US dollars, sterling, Swiss francs, Swedish krona, Norwegian krone and Danish krone. As an indication of their rapid acceptance,

Contract size	Valued at €20 per index point (e.g. value €20,000 at 1000.0)
Delivery months	March, June, September, December (nearest three available for trading)
Delivery day	First business day after the last trading day
Last trading day	13:00:30 (CET), 12:00:30 (London time) third Friday in delivery month[a]
Quotation	Index points (e.g. 1000.0)
Minimum price movement (tick size and value)	0.5 (€10)
LIFFE CONNECT™ trading hours	10:00–18:00 (CET), 08:00–17:00 (London time)
Daily settlement time	17:30 (CET), 16:30 (London time)
Contract standard	Cash settlement based on the exchange delivery settlement price

Exchange delivery settlement price (EDSP): The EDSP is based on the average values of the index[b] every 15 seconds between (and including) 11:40 and 12:00 (London time) on the last trading day. Of the 81 measured values, the highest 12 and lowest 12 will be discarded and the remaining 57 will be averaged to calculate the EDSP. Where necessary, the calculation will be rounded to the nearest half index point.

Table 8.2: FTSE Eurotop 100 index futures

[a] In the event of the third Friday not being a business day, the last trading day shall normally be the last business day preceding the third Friday.

[b] In this instance, "the index" refers to the cash value of the FTSE Eurotop 100 index so as to reflect the underlying index value of the futures contract.

Source: LIFFE

Unit of trading	One future normally represents 100 shares, except for Italian futures where one future represents 1000 shares
Delivery months	Nearest two of March, June, September and December, plus nearest two serial months such that the nearest three calendar months are always available for trading
Quotation	Denmark: Danish kroner (DKK) per share Norway: Norwegian kroner (NOK) per share Sweden: Swedish krona (SEK) per share Switzerland: Swiss francs (CHF) per share All others: euro per share
Minimum price movement (tick size)	€0.01 Denmark DKK 0.5 Norway NOK 0.5 Sweden SEK 0.01 Switzerland CHF 0.1
Tick value	€1.00 Denmark DKK 50 Norway NOK 50 Italy €10 Sweden SEK 1.00 Switzerland CHF 10
Last trading day	refer to www.liffe.com
Settlement day	refer to www.liffe.com
Trading hours	refer to www.liffe.com

Table 8.3: Universal stock futures on continental European shares
Source: LIFFE

12 contracts were introduced in March 2001, by May 2002 the number of contracts available had risen to 122. Interest in these futures is set to grow even more when they are introduced on to the American markets in the near future.

Table 8.3 illustrates that, in essence, these contracts have the same form of specification as their index colleagues and, although universal stock futures are based on a quantity of shares in the company concerned, they are cash settled in the appropriate currency determined which is dependent on the share's domestic currency quotation. In other words there is no requirement to deliver the underlying shares and incur commission charges that such a transaction would involve. For the US and European stocks the number of shares involved is normally 100, for the UK it is normally 1000.

8.1.1 *Uses of futures*

1. Speculation

Perhaps the greatest use to which futures are put is that of speculation. Given the ease with which these contracts can be bought and sold, long or short positions can be opened quickly, efficiently and, in terms of institutional and transactions cost, cheaply. Basic speculative strategies involving futures would be to buy the futures if a rise in the price of the underlying is expected, and to sell the futures if the price of the underlying is expected to fall. These types of speculative positions will be opened and closed quickly, often within minutes but certainly on an intra-day basis. These speculative forays into the market are an essential feature of "fair" price determination.

Example 8.1 gives an illustration of a speculative strategy.

Example 8.1

The Eurotop 100 index is expected to rise on the easing of interest rates by the European Central Bank (ECB). At the moment the underlying index stands at 3000 and its associated futures contract stands at 3006.

Strategy: Buy one Eurotop 100 futures contract at the market quote of 3006.

Outcome: Five minutes later ECB announces that the interest rate is to be cut by one half of 1%. This reduction is 25 basis points more than the market expected. The initial market response is to push the futures contract to a quote of 3056. Clearly a gain of 50 index points. Since each point is multiplied by €20 the gain to the speculator, who sells the contract, is €1000.

As a second example consider the case where an investor is bearish about the general market direction but bullish about the performance of a particular stock in the index. USFs could be used to capitalise on this view. A short position in the index futures contract could be combined with a long USF position in the individual stock. If the investor's view proves correct he/she will gain on both legs of the transaction.

2. Arbitrage

Arbitrage is the making of a *risk-free* profit, in this case by taking advantage of price differentials between an instrument traded in different marketplaces. There are in practice several ways in which this might be achieved. For example, if a futures contract on an instrument is traded in Country A and a futures contract is traded on the same instrument in Country B, on occasions there may be small mispricing occurrences. An arbitrage strategy would be to buy at the cheap quote in one country and simultaneously sell at the high quote in the other country. The position would be unwound when the prices quoted on the markets in the different countries converge.

3. Hedging

With many equity funds now being managed on an index-tracking basis, the availability of futures contracts which match the underlying security benchmark indexes provides an easy-to-use, effective hedging vehicle. However, even non-index-tracking equity portfolios can benefit from the existence of equity futures contracts, as will be demonstrated later in this chapter.

Example 8.2

The first step in hedging a diversified equity portfolio is to identify the position's risk profile. Diagrammatically this is illustrated in Figure 8.1. Note that as the underlying security rises there is a positive payout while, naturally, if the security's price falls there is a loss.

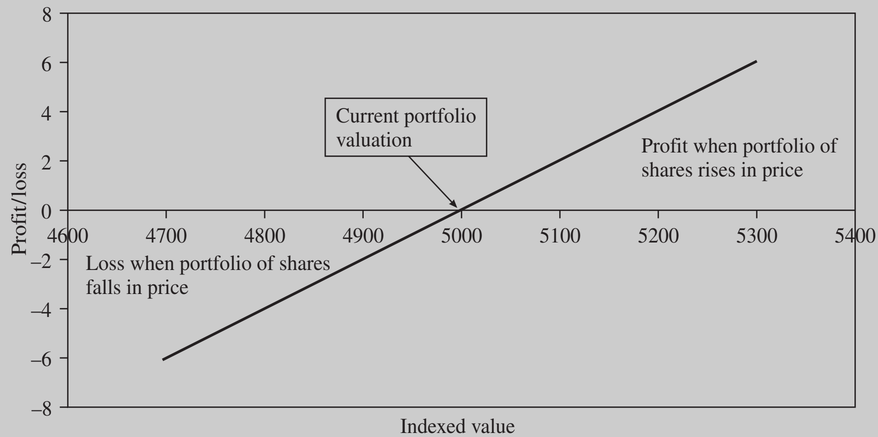

Figure 8.1: Diversified equity portfolio: long position

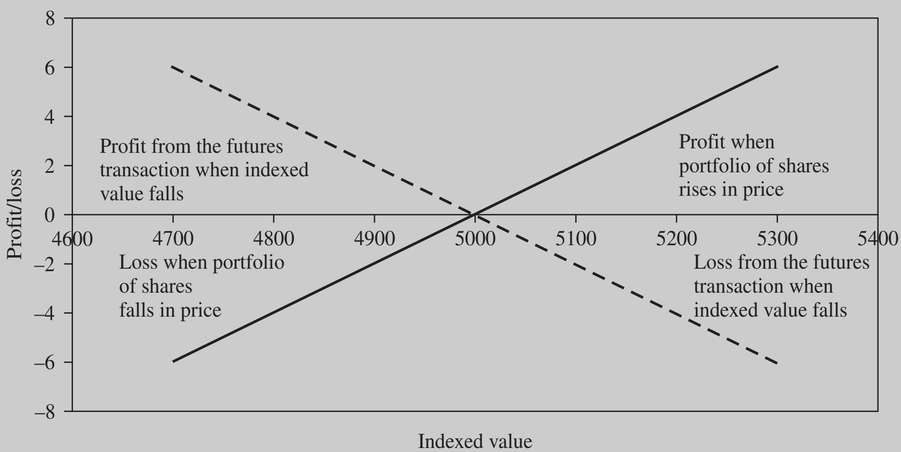

Figure 8.2: Diversified equity portfolio: long portfolio + short futures position

To counteract the risk of loss an instrument with the opposite payout profile needs to be used. This is achieved by selling (or shorting) an appropriate number of futures contracts. With each fall in the price of the underlying security, put simplistically, the potential profit from the short futures position rises. Figure 8.2 demonstrates the payout profiles of the cash and futures instruments. Combining the losses and gains from the two positions creates a flat profile: the gain from one instrument offsets the loss from the other.

In order to mount a hedge in practice, a number of values need to be used, as given in Table 8.4.

In addition, to provide an illustration for this example, the assumption is made that the cash market index will fall to 5900 by the maturity of the futures contract.

Current value of the cash market portfolio	£6,000,000
The current spot market quote	6256
The current futures quote	6268
The relationship between the portfolio and the appropriate futures contract (beta)	1 (index tracking)
The number of futures contracts needed to hedge the portfolio	Must be calculated
The time frame over which the hedge is to be in place	35 days
The exchange-based contract specification	LIFFE FTSE 100
Maturity date of the futures contract	35 days
Current interest rate	6% p.a.
Dividend yield	4% p.a.

Table 8.4

In this example it should be noted that the hedge horizon set at 35 days in the future exactly matches the expiry date of the futures contract used to hedge the position. Moreover there is a portfolio beta (β_p) of unity, in other words the unrealistic assumption of zero tracking error is being made.

To assess the effectiveness of this futures hedge it is useful to summarise the positions as indicated in Table 8.5.

Part of the information required to assess the effectiveness of the hedge appears as text in row two: the number of JUN contracts required to fully hedge the portfolio. This figure can be found using the following formula:

$$\text{No. contracts} = \frac{VP}{(IVP)(FQ)}\beta_p \text{ or } \frac{VP}{(IVP)(CI)}\beta_p \tag{8.1}$$

where

- VP is the current value of the portfolio;
- IVP is the value of an index point (also known as the multiplier);
- FQ is the current futures quote;
- CI is the current value of the cash index;
- β_p is a measure of the relationship between the portfolio to be hedged and the hedge instrument.

Date	**Cash market**	**Futures market**
xx May	£6,000,000	(£10)(6268)(No. of JUN contracts)
Horizon	To be estimated	(£10)(5900)(No. of JUN contracts)
Result		

Table 8.5

Digression

One question which could quite validly be posed is: "Why does the portfolio β qualify as a contract adjustment parameter?"

Minimum variance property of the hedge parameter *h*

The total value of a futures augmented portfolio at time t can be expressed as a linear relationship between the underlying securities held in the portfolio (S_t) – the number of futures contracts used to hedge the position multiplied by the value of one futures contract. This can be expressed as:

$$VP_t = S_t - hF_t. \tag{8.2}$$

Any change in the value of the portfolio can be measured by taking the difference between today's values and yesterday's values.

$$(VP_t - VP_{t-1}) = (S_t - S_{t-1}) - h(F_t - F_{t-1}) \tag{8.3}$$

or

$$\Delta VP_t = \Delta S_t - h\Delta F_t. \tag{8.4}$$

One approach that could be used to identify a suitable hedge parameter h would be to select that hedge parameter which minimises the variance of the changes in the value of the portfolio. Using the expectation operator this can be expressed as:

$$E[\Delta VP_t - E(\Delta VP_t)]^2 = E[(\Delta S_t - E(\Delta S_t)) - h(\Delta F_t - E(\Delta F_t))]^2$$

which, on expansion, can be rewritten as:

$$\sigma^2_{\Delta VP} = \sigma^2_{\Delta S} - 2h\sigma_{\Delta SF} + h^2\sigma^2_{\Delta F} \tag{8.5}$$

where

$\sigma^2_{\Delta VP}$	represents the variance of changes in portfolio value;
$\sigma^2_{\Delta F}$	represents the variance of changes in the futures value;
$\sigma^2_{\Delta S}$	represents the variance of changes in cash index value;
$\sigma_{\Delta SF}$	represents the co-variance of changes in cash index and futures values.

Finding the partial derivative of σ^2 with respect to h and setting it equal to zero will deliver an optimal value.

$$\frac{\partial\sigma^2_{\Delta VP}}{\partial h} = -2\sigma_{\Delta SF} + 2h\sigma^2_{\Delta F} = 0. \tag{8.6}$$

Solving for h yields

$$h = \frac{\sigma_{\Delta SF}}{\sigma^2_{\Delta F}} = \beta_p \tag{8.7}$$

where the beta coefficient is calculated from the relationship

$$\Delta S_t = \alpha + \beta\Delta F_t + \varepsilon_t. \tag{8.8}$$

So there is a natural hedge parameter that relates back to the relationship that exists between the cash market and futures market instruments.

That the h parameter actually refers to a minimum can be verified by considering the second derivative with respect to h.

$$\frac{\partial^2 \sigma^2_{\Delta VP}}{\partial^2 h} = 2\sigma^2_{\Delta F} > 0$$

and since the second derivative is strictly greater than zero the conclusion can be drawn that β_p represents a minimum value.

Another interesting relationship that comes from this analysis is the correlation coefficient.

By manipulating the equation for β_p it becomes evident that the equation can be restated in terms of the coefficient of correlation (ρ).

$$\rho = \beta \frac{\sigma_{\Delta S}}{\sigma_{\Delta F}}. \tag{8.9}$$

Based on the data provided the number of contracts required is found to be:

$$\text{No. contracts} = \frac{£6{,}000{,}000}{(£10)(6268)} 1 = 95.72 \text{ or } \frac{£6{,}000{,}000}{(£10)(6256)} 1 = 95.91.$$

Adopting either approach yields a figure of 96 contracts to put on a full hedge. By applying this figure to the appropriate cells in Table 8.5, Table 8.6 is obtained.

The next cell to require attention is the cash market at the horizon cell. In order to complete this cell the effect of a fall in the index to 5900 on the cash market portfolio needs to be calculated. An approximate value can be obtained by subtracting the product of the starting value of the portfolio and the beta adjusted relative decline in the index from the starting value of the portfolio. This is presented in equation (8.10).

$$VP_t = VP_{t-1} + VP_{t-1}\left(\frac{Index_t - Index_{t-1}}{Index_t}\right)\beta_p \tag{8.10}$$

where $Index_{t-i}$ represents the cash market index level at time $t - i$, for $i = 1, 2$.

All other variables are as defined in equation (8.1).

Inserting the data provided into equation (8.1) yields:

$$VP_t = £6{,}000{,}000 + £6{,}000{,}000\left(\frac{5900 - 6256}{6256}\right) \times 1 = £5{,}658{,}568$$

so that Table 8.6 can be updated and the outcome of the hedge obtained.

It is evident from Table 8.7 that the value of the cash market portfolio has fallen by £341,432. However, the futures market transaction has generated a profit of £353,280 resulting in an overall profit of £11,848. Of course this calculated profit is an illusion. In

Date	**Cash market**	**Futures market**
xx May	£6,000,000	£6,017,280
Horizon	To be estimated	£5,664,000
Result		

Table 8.6

Date	Cash market	Futures market	Overall result
xx May	£6,000,000	£6,017,280	
Horizon	£5,658,568	£5,664,000	
Result	−£341,432	£353,280	£11,848

Table 8.7

reality a number of factors would mitigate against this profit being obtained or indeed the reported outcome actually being observed. For example, no account has been taken of the margin account which would be running in the background and would face mark-to-market on a daily basis. No account has been taken of the possible need to borrow the funds necessary to enter a futures transaction or alternatively the opportunity cost of earmarking own funds for the futures operation. On the cash market side of the balance sheet no account has been taken of dividends that might have been received or rights issues, share splits, etc. that might have occurred over the period of the hedge. Notwithstanding these omissions the arithmetic demonstrates that the hedge is effective and the portfolio is hedged through the shorting of the required number of futures contracts.

4. Engineering
Futures contracts can be used to alter the payout profile of a portfolio by buying or selling futures contracts. For example, a portfolio's beta can be adjusted to reflect bullish or bearish sentiment. In the case of a long equity, portfolio buying futures contracts will increase market exposure and would be an appropriate strategy when expectations are that the market is set to rise. Selling futures might simply hedge the portfolio to yield a flat (no gain, no loss) position or they may be used more aggressively to take advantage of and make a gain should the anticipated fall in price materialise.

8.2 Operational characteristics of equity futures contracts

8.2.1 *Margin*

There are a number of advantages to using equity index or universal stock futures contracts. They are, for example, highly liquid instruments, which implies that positions can be opened and closed very quickly. Short positions can be opened just as easily as long positions even though the short seller may not hold the underlying instrument in their portfolio. In addition they normally enjoy low bid–ask spreads as a result of their popularity and high trading volumes. Perhaps, though, the feature that really sets them apart from the underlying security is that they are bought and sold on *margin*.

One form of margin is termed *initial margin*, in effect, this is a deposit that represents a sign of good faith when entering into a contract. In legal terms this deposit is regarded as *consideration*. In the absence of consideration the contract would be deemed to be invalid in English law. Initial margin can take several forms. It could, for example, be a cash payment, certificates of deposit, a bank guarantee, government bonds, and in some cases even equity might be accepted as collateral as it is deposited with a clearing house which takes on the management of the position. Whatever form this initial margin takes, an attractive feature is that the amount involved is only a small percentage of the value of the underlying security – often the initial margin is no more than 5% of the underlying security's full market value.

A second, and crucially important, type of margin is known as *variation* margin. This margin measures the profit or loss accruing to open futures positions each day in a process known as marking-to-market. Variation margin minimises the risk of counterparty default by requiring one of the counterparties to add money to their margin account in order to maintain their open position. This type of margin, unlike initial margin, can only be paid in cash. Example 8.3 will help to illustrate how the margin account works.

Example 8.3

A fund manager takes a bullish view of the stock market and decides to use five equity index futures contracts to take advantage of the anticipated rise in the market. The current quote for the nearby futures contract is 5000, the futures contract specifies a £10 per index point multiplier and initial margin is set by the appropriate clearing house at £12,500. Table 8.8 summarises the cash flows involved in the operation of the margin account.

The fund manager decides at the end of Day 4 that the position is no longer worth holding because the anticipated market rise has fizzled out. The position is closed out at 4975. The resulting overall loss on the position is:

Buy at 5000, sell at 4975 = 25 index points lost.

Each index point is valued at £10 so the 25 index points lost converts to a monetary value of £250. There are 5 contracts involved in the transaction, which means that the total loss is £1250.

Notice that the margin account has ensured that on Day 5 the only amount that needs to be paid is the loss experienced on Day 4. The variation margin payments have settled to profits and loss on a daily basis at the commencement of business each day. The net result will still be a loss of £1250. This overall loss will be the sum of the daily variation margin cash flows: (£1000 + £800 − £2550 − £500) = −£1250.

Time	**Action**
Day 0 – p.m.	At close of business **buy** 5 equity index contracts to **open** @ 5000
Day 1 – a.m.	**a.m.**: At start of business deposit £12,500 initial margin – **no variation margin**
Day 1 – p.m.	**p.m.**: Futures closes at 5020
Day 2 – a.m.	**a.m.**: At start of business
	Receive: (5)(£10)(5020 − 5000) = £1000 **variation margin**
Day 2 – p.m.	**p.m.**: Futures closes at 5036
Day 3 – a.m.	**a.m.**: At start of business
	Receive: (5)(£10)(5036 − 5020) = £800 **variation margin**
Day 3 – p.m.	**p.m.**: Futures closes at 4985
Day 4 – a.m.	**a.m.**: At start of business
	Pay: (5)(£10)(4985 − 5036) = £2550 **variation margin**
Day 4 – p.m.	**p.m.**: **sell** 5 equity index contracts to **close** at 4975
Day 5 – a.m.	**a.m.**: At start of business
	Pay: (5)(£10)(4975 − 4985) = £500 **variation margin**

Table 8.8

Other types of margin may also exist. Buffer margin may be called by the clearing house in advance of major events which are expected to impact heavily on financial markets. In some circumstances intra-day margin may be called when increased volatility leads to large market swings occurring.

8.2.2 *Beta and basis*

In the current context basis can be defined as the difference between the cash or spot market instrument and the futures market instrument:

(a) Basis = Futures price – Spot price

but could also appear as:

(b) Basis = Spot price – Futures price.

In either (a) or (b) basis can be positive, negative or zero. If (a) is adopted as the working approach to calculating then basis will be positive when the futures price is greater than the spot price. This situation is termed *contango*. The currently quoted futures price is greater than the currently quoted spot and will tend to fall over time to equal the spot price on delivery day. Still using (a), should the futures price be less than the spot price, basis will be negative. This situation is termed *backwardation* and describes the case where the futures price will tend to rise over time to equal the spot on delivery day – a feature of some instruments on commodity markets in London in recent years, for example.

The price differences that exist between the two instruments at any time occur because, although clearly related, the instruments are not perfectly correlated. In fact only on delivery day, when the basis is zero, will the prices on the two markets converge. At any other time before maturity accurately predicting basis is virtually impossible. There are a number of reasons why this is the case:

1. Prices in both markets will respond to news entering those markets, but the impact of this news may be different in each case. For example, an unexpected announcement by the Federal Reserve of a cut in the interest rate may bring about a 20 point rise in the spot market as an initial reaction as the prices of equity are marked up. The futures market response may be an immediate 50 point gain.
2. The operational rules in each market can be substantially different. One particular example of this would be where an exchange imposes a maximum allowable price movement on a futures contract on any day. Where such a limit exists and the specified limit is reached, the exchange can close business in that contract. This, theoretically, should help to re-impose the semblance of an orderly market in that instrument. However, closing down the futures market, effectively freezing prices and making contracts temporarily illiquid, does not stop business on the spot market where transactions in the underlying security continue. Far from giving the market a chance to settle down, this course of action can allow a widening of the basis to creep in between the two instruments.
3. The technical construction of indexes, the updating of those indexes to reflect price movements, and the reporting of those indexes. On the technical side a futures index

quote will include some premium. This typically will include a cost of carry which will be the amount of interest that must be paid to finance the purchase of the underlying security (plus storage costs in the case of commodities) less dividends, interest or other income received from the underlying security during the life of the futures contract.

Basis can be interpreted in different ways and to an extent the interpretation depends on the context in which the word is being used. Later chapters will use basis to describe differences between the prices of cash instruments in different market sectors.

As mentioned above, on delivery day the basis between the spot and futures instruments will be zero. The price at which the futures contract closes is known as the exchange delivery settlement price (EDSP) and the way in which the EDSP is calculated will be clearly defined in the futures contract.

Of course, in general the β of a portfolio will not be equal to unity even if the portfolio in question has been carefully constructed to track a particular index. The data in Table 8.9 reports the closing prices for the FTSE 100 and the futures contract on that index. Note that in the series the 14 May observation shows the cash market 145 below the futures market at close. This figure is an outlier. It reflects a trade for the shares in one company where two extra zeros where inputted in error. The trade mistake was quickly spotted and corrected in the futures market but took time to readjust in the spot market resulting in a 145 difference being recorded at close of business. This resulting increase in basis is illustrated in Figure 8.3.

Date	**FTSE 100**	**Futures**	**Basis (FQ-FTSE 100)**
23 Apr 01	5871.3	5882.5	11.2
24 Apr 01	5840.3	5870.0	29.7
25 Apr 01	5827.5	5855.0	27.5
26 Apr 01	5868.3	5892.0	23.7
27 Apr 01	5951.4	5974.0	22.6
30 Apr 01	5966.9	5996.0	29.1
01 May 01	5928.0	5951.5	23.5
02 May 01	5904.2	5925.0	20.8
03 May 01	5765.8	5779.5	13.7
04 May 01	5870.3	5902.0	31.7
08 May 01	5886.4	5914.5	28.1
09 May 01	5893.7	5913.5	19.8
10 May 01	5964.0	5973.5	9.5
11 May 01	5896.8	5916.0	19.2
14 May 01	5690.5	5835.5	145.0
15 May 01	5842.9	5851.0	8.1
16 May 01	5884.0	5900.0	16.0
17 May 01	5904.5	5920.0	15.5
18 May 01	5915.0	5941.5	26.5
21 May 01	5941.6	5959.0	17.4
22 May 01	5976.6	5984.0	7.4
23 May 01	5897.4	5911.5	14.1
24 May 01	5915.9	5929.5	13.6
25 May 01	5878.3	5888.5	10.2

Table 8.9

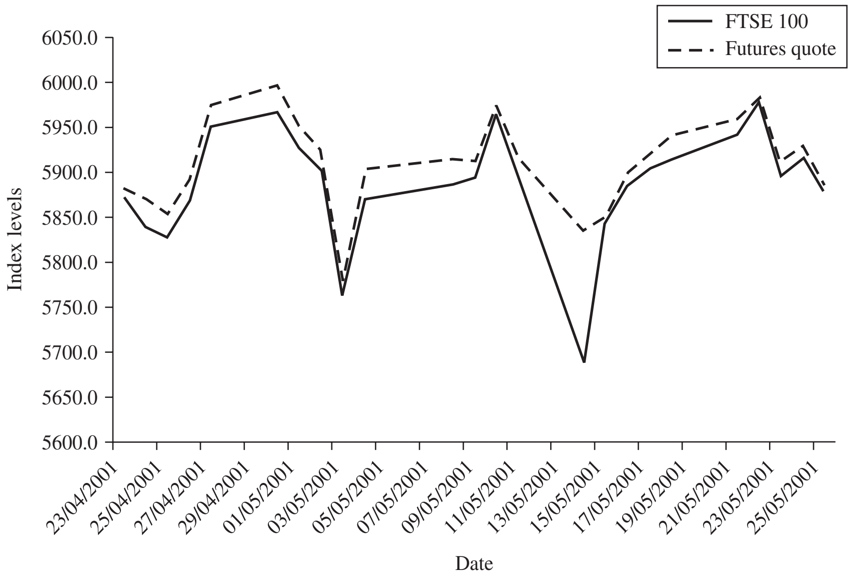

Figure 8.3: Futures and cash market quotes

The parameter β can be calculated in several ways using linear regression. It could, for example, be calculated using price increments, that is, the difference between the price today and where it was yesterday (or some previous time period). This indeed is the approach used to justify the choice of beta as an optimal hedge parameter earlier in this chapter. Alternatively β could be found using relative price changes or using natural log price increments, which are referred to as natural log returns. This then would not fit with the algebra presented earlier but the effect on hedge is unlikely to be dramatic when using this sub-optimal hedge parameter.

Whatever approach is adopted it must be emphasised that the estimate of β in the spot versus futures linear relationship is clearly going to vary over time. The estimate obtained

Regression statistics					
R square	0.886027				
Adjusted R square	0.883793				
Standard error	31.01521				
Observations	53				
ANOVA	*df*	*SS*	*MS*	*F*	Significance *F*
Regression	1	381387.2	381387.2	396.4757	1.04634E − 25
Residual	51	49059.12	961.9435		
Total	52	430446.4			
	Coefficient	Standard err	*t*-stat	*p*-value	
Intercept	0.749725	4.263123	0.175863	0.861098	
X variable 1	1.051984	0.052832	19.9117	1.05E-25	

Table 8.10: Summary output

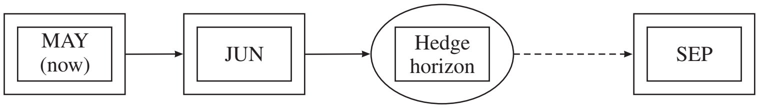

Figure 8.4: Non-matching maturity dates

will depend and vary upon the number of observations used to obtain an estimate and it will also vary depending on whether or not the data contains outliers that are included or omitted in the regression analysis.

Basing a regression here on the natural log returns for a total of 52 observations (9 March–25 May 2001) the estimated β-hat is 1.05. If the 14 May (145 basis) is omitted the portfolio β-hat is 1.02. In a statistical sense, as evidenced by *p*-values and *t*-statistics, both estimates of β-hat are highly significant, and the relationship between the spot and future log returns is very good – the coefficient of correlation is 89% if all the data are used and 96% if the outlier is omitted.

8.2.3 *More advanced hedges*

On many occasions, of course, it is pure coincidence when the hedge horizon date matches the maturity date of an exchange-traded futures contract. Consider the case where the horizon required by a fund manager is 35 days after the maturity date of the nearby contract (JUN). The maturity date of the next contract out, which might provide an alternative hedge vehicle, lies far beyond the required horizon date.

When this situation arises and the hedge horizon does not match the maturity date of the contract exactly, the method of hedging using futures contracts needs to be adjusted. There are several ways in which this can be achieved.

In the example that will be developed the nearby contract (JUN) clearly expires before the hedge horizon date. One way to hedge the long equity position would be to sell the appropriate number of nearby futures contracts today and when they expire roll over the position into the SEP contract. Alternatively an appropriate number of contracts in the longer-dated contract (SEP) could be sold today thus avoiding the problem of rollover. Both of these approaches are known as *stack hedges* – all the futures contracts required to hedge the position are stacked into one maturity date.

Example 8.4: Stack hedge

A fund manager is running an index-tracking portfolio designed to follow the movements in the Eurotop 100 index. The current value of the portfolio is €9,848,485, the estimated beta is 0.99, the hedge horizon date is 35 days after the maturity date of the nearby Eurotop 100 futures contract. The fund manager wishes to lock-in to today's value of the index in order to report good annual growth figures to his clients.

The Eurotop 100 cash market index currently stands at 3250 and the value of an index point, drawn from the LIFFE contract, is €20. According to the arithmetic suggested in equation (8.1) this implies that exactly 150 futures contracts will need to be shorted. The fund manager decides to avoid the problem of rollover and short the SEP contract. Given the current market quote for the SEP futures contract is 3298.

Variable	Current value
Value of portfolio	€9,848,485
Portfolio beta	0.99
Eurotop 100 multiplier	€20
Cash market index now	3250
JUN futures quote	3265
SEP futures quote	3298
Days to hedge horizon	70
Days to maturity of JUN contract from now	35
Days to maturity of SEP contract from now	125

Table 8.11: Background data

Moving forward in time to the hedge horizon date, if the assumption is made that the spot market now stands at 3050 and the associated futures is quoted at 3062, the outcome of the hedge will be as follows:

Date	Cash market	Futures market	Net outcome
xx May	€9,848,485	€9,894,000	
Horizon	€9,248,485	€9,186,000	
Result	€600,000	€708,000	€108,000

Table 8.12

Clearly the hedge has been quite effective, returning a net gain of €108,000. However, it must be recognised that the result is highly dependent upon the basis that exists both at the time that the hedge is mounted and when the hedge is removed. A second case will serve to clarify this assertion.

Suppose, for example, that when the hedge is initiated the basis is in backwardation, the futures quote at 3248 is lower than the cash market index. One justification for this would be that there is an expectation that the market is set to fall.

Suppose further that at the hedge horizon bullish sentiment reigns in the market and the futures is in contango – the futures quote at 3088 is higher than the cash market. Under these conditions the hedge outcome might have resulted in the following outcome:

Date	Cash market	Futures market	Net outcome
xx May	€9,848,485	€9,744,000	
Horizon	€9,248,485	€9,264,000	
Result	€600,000	€480,000	−€120,000

Table 8.13

In this case the futures market transaction fails to deliver the required protection. The first leg of the futures transaction in May is undertaken at a disadvantageous basis and the second leg is carried out at a disadvantageous basis, too. The overall result is a loss of €120,000. How this problem might be handled will be looked at in more detail in Section 8.2.4 which deals with hedging basis risk.

Example 8.5: Strip hedge

In this case the fund manager decides that a mixture of contracts from the JUN and the SEP maturity dates represents an appropriate strategy. In order to determine an appropriate split between the contracts the number of days involved in the hedge period is calculated as are the number of days until the maturity date of the far (SEP) contract. These values are, say, 70 days and 125 days, respectively. One contract split would then be given by the ratios:

Ratio	Number of contracts
$70/125 = 0.56$	$(150)(0.56) = 84$
$55/125 = 0.44$	$(150)(0.44) = 66$

Table 8.14

The required action today under this strategy is to short 84 JUN contracts and in addition 66 SEP contracts. In analysing the effectiveness of the hedge the process will need to be looked at in two stages. It should be noted that this split assumes that basis is well-behaved and declines at a constant rate over time until the contract reaches maturity. The suggested way of splitting the contracts between the two maturity dates is one of several possibilities, all of which will behave in a similar fashion to the one demonstrated below.

To illustrate the processes involved the strategy will be broken down and explained in two stages.

Stage 1: Until JUN contract maturity

Assume the cash market index stands at 3112 when the JUN contract matures.

At the end of this stage the hedge will be showing a loss of €156,960. This loss though ignores the profit that will have accrued to the 66 SEP contracts shorted when the hedge was initiated. Notwithstanding this failure to hedge the equity portfolio's loss completely using the 84 JUN contracts, there will be a profit from the JUN contracts that can be reinvested in the now cheaper stocks or placed on an interest-bearing deposit account for the remaining duration of the hedge. In this example it is assumed that the profits from stage 1 of the hedge are placed or deposited at the current rate of interest – say 3.50% p.a. (Note that like transactions and institutional costs and the cost associated with running a margin account this interest income will be ignored when calculating the outcome of the hedge.)

The revised value of the portfolio will be €9,434,485. Since there are only 66 SEP contracts hedging the position, more contracts will need to be shorted to re-instigate a "full" hedge. Basing the calculations on the level of the current cash market index and

Date	Cash market	Futures market	Net outcome
xx May	€9,848,485	€5,485,200	
Maturity contract JUN	€9,434,485	€5,228,160	
Result	€414,000	€257,040	−€156,960

Table 8.15

the 0.56 ratio used to determine the number of JUN contracts, when first deciding the nature of the contract split, a further 84 SEP contracts must be shorted at the prevailing futures quote of 3138.[1]

Stage 2: Until hedge horizon

The first set of SEP contracts (66 shorted at hedging initiation) is shown in Table 8.16. The second set of SEP contracts (84 shorted at rollover) is given in Table 8.17.

At this point the overall effectiveness of the hedge can be assessed. The profit from the JUN futures transaction is €257,040 + interest [€257,040(0.035)(35/365)] + profit from 66 SEP contracts shorted at initiation + profit from 84 SEP contracts shorted at rollover.

This more than matches the €600,000 lost on the cash market portfolio. The profits from the futures market can be used to "top-up" the portfolio, whilst retaining its index-tracking nature, at the prevailing lower share prices.

In this strip hedge case no account has been taken of basis risk – as illustrated in the stack hedge case. The problem of basis risk has not been eliminated and remains as a threat to the overall effectiveness of the hedge. One way in which this type of risk might be handled is demonstrated below.

Date	Cash market	Futures market	Net outcome
xx May	€9,848,485	€4,353,360	
Horizon	€9,248,485	€4,041,840	
Result	€600,000	€311,520	–€288,480

Table 8.16

Date	Cash market	Futures market	Net outcome
From rollover		€5,271,840	
Horizon		€5,144,160	
Result		€127,680	

Table 8.17

Profit generated by:	
84 short JUN contracts at initiation of hedge	€257,903
84 short SEP contracts from rollover	€127,680
66 short SEP contracts at initiation of hedge	€311,520
Total	€697,103

Table 8.18

[1] For more on calculating theoretical futures prices refer to Section 8.2.5.

8.2.4 *Basis risk hedging*

Example 8.6

Adopting the same example used to analyse the effectiveness of stack and strip hedges the number of contract maturities involved increases to three: JUN, SEP, as previously, but the DEC contract is now introduced, too.

Using the stack hedge the fund manager shorted 150 SEP contracts but since the hedge is to be removed toward the end of July the position will face basis risk. To counteract this risk, at least in part, the fund manager could set up a calendar spread by going short a further 150 SEP contracts and simultaneously going long the same quantity of DEC contracts.

To consider the outcome of this strategy take the results experienced in the case of the stack hedge presented in Table 8.13 where €120,000 were lost by disadvantageous basis. To recap the situation: when the hedge was initiated the future was in backwardation (futures quote 3248, spot market 3250) and when the hedge was removed the future was in contango (futures quote 3088, spot market 3050). Table 8.19 presents the data for this example as Case 1.

Case 1

If the assumption is now made that when the hedge is initiated it is possible to buy the DEC futures at a quote of 3320 and that when the hedge is removed at the horizon the DEC contracts can be sold to close the position at 3170, the result will generate the following payout scheme:

Gain on 300 short SEP contracts: €960,000.

This is found by multiplying the total number of short contracts by the contract multiplier and by the SEP opening futures quote (NOW) and then subtracting the total number of short contracts by the contract multiplier and by the SEP closing futures quote (HORIZON). This is shown in equation (8.11) where the Case 1 data is applied.

$$\begin{aligned}&(300)(20)(3248) - (300)(20)(3088)\\&= €19{,}488{,}000 - €18{,}528{,}000\\&= €960{,}000.\end{aligned} \tag{8.11}$$

Loss on 150 long DEC contracts: €450,000.

NOW	**Case 1**	**Case 2**
Cash market	3250	3250
Futures market SEP	3248	3298
Futures market DEC	3320	3346
HORIZON		
Cash market	3050	3050
Futures market SEP	3088	3062
Futures market DEC	3170	3090

Table 8.19

This is found by multiplying the total number of long contracts by the contract multiplier and by the DEC opening futures quote (NOW) and then subtracting the total number of long contracts by the contract multiplier and by the DEC closing futures quote (HORIZON). This is shown in equation (8.12).

$$\begin{aligned}&(150)(20)(3320) - (150)(20)(3170)\\&= €9{,}960{,}000 - €9{,}510{,}000\\&= €450{,}000.\end{aligned} \tag{8.12}$$

But since the selling price is lower than the buying price the €450,000 represents a loss.

The difference is a gain of €510,000 which, when set against the cash market loss of €600,000, provides at least some cover for the basis risk. The overall loss on the portfolio using the basis hedge is €90,000 compared to a loss of €120,000 if no basis hedge is implemented.

As a second case consider the data originally used to illustrate the effectiveness of the stack hedge. The data applied here is presented as Case 2 in Table 8.19.

Case 2

If the assumption is now made that when the hedge is initiated it is possible to buy the DEC futures at a quote of 3346 and that when the hedge is removed at the horizon the DEC contracts can be sold to close the position at 3090, the result will generate the following payout scheme:

Gain on 300 short SEP contracts: €1,416,000.

This is found by multiplying the total number of short contracts by the contract multiplier and by the SEP opening futures quote (NOW) and then subtracting the total number of short contracts by the contract multiplier and by the SEP closing futures quote (HORIZON). This is shown in equation (8.13).

$$\begin{aligned}&(300)(20)(3298) - (300)(20)(3062)\\&= €19{,}788{,}000 - €18{,}372{,}000\\&= €1{,}416{,}000.\end{aligned} \tag{8.13}$$

Loss on 150 long DEC contracts: €768,000.
This is found by multiplying the total number of long contracts by the contract multiplier and by the DEC opening futures quote (NOW) and then subtracting the total number of long contracts by the contract multiplier and by the DEC closing futures quote (HORIZON). The result of this arithmetic is shown in equation (8.14).

$$\begin{aligned}&(150)(20)(3346) - (150)(20)(3090)\\&= €10{,}038{,}000 - €9{,}270{,}000\\&= €768{,}000.\end{aligned} \tag{8.14}$$

But since the selling price is lower than the buying price the €768,000 represents a loss.

The difference is a gain of €648,000 which, when set against the cash market loss of €600,000, reduces the level of profit when no basis hedge was used. The overall gain on the portfolio using the basis hedge is €48,000 compared to a gain of €108,000 if no basis hedge is implemented.

8.2.5 *Pricing equity futures*

In the examples above the futures prices are assumed to be given by the market. It is, however, possible to calculate an indicative "fair" price for a futures contract. In essence the method for doing this relies on the calculation of an implied forward price. Equation (8.15) provides the algebraic framework in which the calculations can be performed.

$$F_{0k} = S_{0o}\left(1 + [r_{0k} - d_{0k}] \times \frac{\text{Day count}}{365}\right) \tag{8.15}$$

where

- F_{0k} represents the price at time zero (now) of a futures contract deliverable at time k;
- S_{0o} represents the price at time zero (now) of the underlying instrument in the cash market;
- r_{0k} represents the risk-free rate of interest appropriate over the period now to time k, expressed as a decimal;
- d_{0k} represents the dividend yield associated with the underlying instrument over the period now to k, expressed as a decimal.[2]

If k is taken successively to be the JUN, SEP and DEC contract maturity months then, using nominal annual risk-free rates of interest relevant for the period being considered, together with the current cash market quote and assuming a constant dividend yield of, say 2.5%, a "fair" futures price can be estimated for each of the contracts. This is demonstrated in equation (8.16).

Using the data from Table 8.11, the cash market quote is 3250 and there are 35 days until the JUN contract matures. Putting these inputs together with the assumptions that the appropriate nominal annual risk-free rate of interest for the 35-day period, on a 365[3] day count basis, is 6.5% and dividend yield is 2.5% the formula appears as:

$$F_{0_{\text{JUN}}} = 3250 \times \left(1 + [0.065 - 0.025] \times \frac{35}{365}\right) \cong 3262.5. \tag{8.16}$$

Notice that this calculation takes no account of the operation of a margin account where initial and variation margin will play a vital role in the maintenance and success of a strategy. In practice there will be arbitrage-free bounds between which the theoretical futures quote will oscillate. Nevertheless the method of calculating a futures price described in equation (8.15) provides a useful benchmark for assessing the likely benefits to be derived from more complicated hedge strategies.

It should be noted that dealing in any of the contracts offered by LIFFE – or any other exchange – is restricted to members of that exchange, large international financial institutions and some individuals, who will have paid for the privilege of membership. The normal

---

2 Where dividends are assumed to be reinvested in the index portfolio, for example, the DAX index, the dividend yield (d) is excluded from the pricing formula.

3 Note that the UK convention uses a 365 day year. The US and European markets use a 360 day count convention.

mode of operation will then be one where contracts are bought and sold on behalf of parties through brokers.

Apart from the existence of an exchange regulating and ensuring the integrity of the market, there is also a clearing house with which matched buy/sell deals will be registered and which will take on the contractual obligation of delivery. In the case of LIFFE the clearing house is the London Clearing House (LCH).

A real disadvantage of the use of futures contracts also becomes apparent if the portfolio to be hedged does not match any exchange-based contract. For example, a globally diversified equity portfolio will embrace stocks from many countries and involve many different currencies. In such a case it may be necessary to adopt a multiple regression analysis in order to establish the number of equity futures contracts required to hedge over different stock index risk exposures. But the currency risk will, almost certainly, need to be managed by entering foreign exchange forward agreements.

Entering the forward agreement currency market will facilitate the exact matching of maturity dates and the size of an exposure with no basis risk unlike the more rigidly specified futures contracts.

The dynamic policies adopted by exchanges worldwide have resulted in many new contracts being created for market consumption. Recent developments have seen the introduction of universal stock futures, futures contracts on European-wide equity indexes by exchanges worldwide, and the introduction of interest rate swap contracts as well as the more exotic weather contracts.

However, contracts also have a tendency to disappear. Those contracts which fail to attract high market volumes will disappear from the marketplace. Perhaps ahead of its time, in 1991 LIFFE introduced the FTSE Eurotrack 100 index futures contract only to withdraw it by the end of 1992. This can be particularly frustrating for an institution that has gone to the expense of creating an index-tracking equity portfolio in the belief that an exchange-based futures contract would be readily available for hedging, engineering, speculation and arbitrage activities only to find that that is no longer the case. The carefully constructed portfolio will now need to be modified to fit some new index future on offer or even worse abandoned completely – at considerable cost.

9 Equity Swaps

There are basically two types of equity swap. The first is typically a short-term, open-ended transaction which allows users to take long or short positions in individual or pairs of shares in return for a daily marking-to-market of the contracted share(s) against a LIBOR receipt (short positions) or LIBOR payment (long positions). These instruments are known as contracts for differences (CFDs). CFDs are used by a variety of end users:

- equity fund managers wishing to take positions on share price movements in a cheap, leveraged manner;
- option writers wishing to maintain a delta/gamma hedged position in a cheap and effective way;
- and as a cost effective way of creating a basket of stocks.

The second category refers to medium to long-term instruments that have a defined life span: it is this category that this chapter seeks to develop.

9.1 A basic equity swap

A basic equity swap would involve two parties entering into a contractual agreement to exchange a stream of cash flows linked to the total return of an equity index against a schedule of returns derived from a short-term interest rate index. Both of the cash flow streams will be transferred with an agreed frequency for a fixed period of time and will be calculated on an agreed notional principal.

The first equity swaps to be arranged were in 1989 and indicate, once again, the creativity invested in product innovation. Their initial impact on markets was mixed, and exact volumes, given the nature of the product, are difficult to assess. Nevertheless the large institutions all offer this type of instrument.

Equity swaps differ from the interest rate and currency swap cousins that have been discussed earlier in that at least one leg of the transaction involves the return on an equity index, a single share, or a basket of shares. There are a number of valid reasons for using this type of instrument.

1. In a domestic or international sense they allow the counterparties to such a transaction the opportunity to gain exposure to the customised returns in one market in exchange for those of another market in which they are currently invested.
2. By entering one transaction repetitive transactions and institutional costs are avoided. For example it could be argued that similar positions could be engineered simply by using either a *strip* of futures contracts or by *stacking* futures contracts to achieve the desired exposure. If either of these approaches are adopted then recurring transactions and government taxes are likely to be incurred. An alternative to the use of futures contracts would be to track the desired index or sector through direct purchase of the underlying securities. Unlike a futures construction this approach would involve high

funding costs, and transaction costs in terms of the purchase of batches of individual shares required to replicate the desired position. It also requires the purchase of the necessary shares in appropriate quantities and this may prove problematic in terms of acquisition price if the quantities involved are huge or odd amounts.

3. In the case of individual stocks or baskets of stocks a long position in stocks can be taken without actually purchasing the stock. If the position is short and stocks are not already held in the portfolio, the stocks do not need to be borrowed to cover the position. If the stocks are already held in the portfolio voting rights will still be retained.
4. In addition to the costs described in Item 2 there are other important issues involved in adopting the alternative investment strategies. Using futures will involve the maintenance of a margin account and, whether a strip approach or stack approach is used, basis will play an important role. If the underlying shares are bought considerations such as the reinvestment of dividends and rebalancing of the portfolio to adjust for changes in market capitalisation must be undertaken. So in general an equity swap should prove a more efficient instrument in so far as there would be no custodial fees and no index tracking error.
5. Yet another advantage accruing to the use of equity swaps is *market invisibility*. In the case of Items 2 and 3 large, visible transactions in the futures market could have a serious distorting influence on basis at various points during the purchase, rolling over, and closing out phases of a transaction. Similarly for the individual stock case, the invisibility of a swap transaction means that, for a limited period, a large exposure to a company's stock can be reduced or removed without announcing such a move to the market by selling stock.
6. Using equity swaps can also facilitate arbitrage against currently held equity portfolios.
7. They can also, in certain markets, allow government imposed or legal restrictions to be circumnavigated. Or, indeed take advantage of government tax or investment incentives.

Against this list of advantages must be set what is probably their greatest disadvantage, namely illiquidity. Bear in mind that equity swaps are a form of total return swap and there may be a potential difficulty in unwinding an equity swap when the reference index moves adversely. If the position is long then rises in the index will generate gains but falls will leave the long position paying both the index losses and LIBOR. The notional principals involved in these transactions are large, generally in the region of $50 million, and as with many OTC instruments, the advantage of a tailor-made product can become a disadvantage when attempting to clear the position.

In terms of construction straightforward equity swap will involve the exchange of the returns from a stock market index against some short-term interest rate index such as LIBOR. Both sets of payments involved in the transaction will be calculated on the same amount of notional principal.

However, since their introduction there have been many developments to the *plain vanilla* instrument, i.e. the type of swap described above. One such variation might be the swapping of a stream of returns from one equity index against those of another equity index. An example of this would be the exchange of returns from the DAX against returns from the S&P 500. *Variable principal equity swaps* – a swap which adjusts the notional principal according to the equity index growth or decline over the life of the swap – is another variation.

Five-year swap for DAX index-linked note with "cliquet call spread"	
Notional amount:	€50 million plus
Trade date:	To be defined
Settlement date:	To be defined
Underlying index:	DAX index ("DAX")
Participation factor:	100%
Issuer receives:	6 months Euribor – 90 bp paid semi-annually in arrears on an act/360 basis
Issuer pays:	Annual and in arrears the following: Participation factor × Min[Max$((DAX_t - DAX_{t-1})/DAX_{t-1}$, 0); 6.5%]

Table 9.1

Other examples could be the swapping of returns from an equity index against the returns from some commodity, for example, oil, or *rainbow swaps*, also known as *blended swaps*, where the weighted average returns from two or more equity indexes are used to quantify settlement; or *asset allocation swaps* where settlement is determined by comparing two equity index returns and choosing the greater. And equity swaps with *embedded options* used to set a ceiling or floor on equity index returns. For example, *capped equity swaps* – a swap which limits the amount of growth in the equity index or the interest rate index (or both) – and *knock-out equity swaps* – a swap which ceases to exist when a predefined event occurs. Table 9.1 provides an idea of the terms involved in a swap with an embedded option.

Interpreting the internal brackets first, the holder of the bond receives the DAX return over the period $t - (t - 1)$ or zero. If growth is positive that amount is then compared to the min of 6.5% and the positive DAX growth. The payout will be between 0 and 6.5%. This is the "call spread" aspect of the instrument. The cliquet component comes in the calculation of the DAX growth, which is based on the time periods t and $t - 1$. So as time progresses the t and $t - 1$ indexes change. The first growth will be determined by $(DAX_1 - DAX_0)/DAX_0$. In the next period it will be determined by $(DAX_2 - DAX_1)/DAX_1$.

Whatever returns are being contractually exchanged the mechanics of the swap remain very similar to those described in the chapters dealing with interest rate swaps. An example will serve to demonstrate these similarities but will also highlight some important differences.

9.2 Single-currency, fixed notional principal, equity index/sterling LIBOR swap

Example 9.1

An equity fund manager with £30 million in UK stocks under management takes a pessimistic view of the economy. He/she feels that international sentiment will force interest rates to rise as national governments and central banks throughout the world strive to contain inflation. Such actions will, he/she feels, impact negatively on world stock markets for many months.

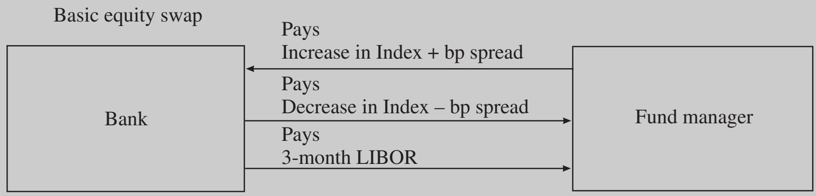

Figure 9.1

The manager does not wish to sell stock for fear of the effect that this would have on stock market prices. Nor does he/she wish to use futures or options as vehicles with which to hedge risk – the sheer volume of short futures contracts involved in hedging the portfolio would inevitably drive down futures quotes which, in turn, would lead to lower stock prices. Moreover should a downturn in stock prices materialise and continue over several quarters as expected for many months, the fund manager would need to roll over his/her positions each time incurring basis and transaction costs.

One possible solution to the problem faced lies in swapping the returns on the index, in which the portfolio shares are a component, in return for LIBOR. This strategy can, just like futures or options, be undertaken for all or part of the portfolio. However, unlike futures and options this strategy can be undertaken for multiple periods without rollover costs, basis or volatility risks, and with the transactions costs fixed at the time the contract is entered into.

The equity swap considered in the example which follows is for one year, settled quarterly on the basis of the 3-month FTSE 100 index and dividend yield return plus a 10 basis point fixed spread against 3-month sterling LIBOR. The notional principal is fixed at £30,000,000 throughout the life of the swap and, since the currency involved is sterling for both parties, there is no foreign exchange risk. The swap's structure may appear as represented in Figure 9.1.

One aspect of this structure should be immediately apparent, namely, that should the fund manager's scenario hold, not only will he/she receive the agreed LIBOR rate but he/she will also receive an amount equal to the negative index return.

To illustrate the mechanics of the process consider the following observed market outcomes over the life span of the swap.

The figures reported in Table 9.2 can be regarded as observed market outcomes at the time that the exchanges were undertaken. In other words, in this example, the quarter-by-quarter transactions are being looked at retrospectively with full knowledge of the prevailing market rates required to perform the necessary calculations.

Period (quarters)	**LIBOR (%)**	**FTSE 100 index**
Now	4.125	5300
End Quarter 1	4.125	5050
End Quarter 2	4.250	4910
End Quarter 3	4.375	5120
End Quarter 4	4.500	5250

Table 9.2

	Day count	FTSE 100	Percentage change	Fixed spread in bps	LIBOR	Dividend yield p.a.
Now		5300		10	4.125%	1.100%
3 months	91	5050	−4.72		4.125%	1.100%
6 months	92	4910	−2.77		4.250%	1.100%
9 months	91	5120	4.28		4.375%	1.100%
12 months	91	5250	2.54		4.500%	1.100%

Table 9.3

Column 2 of Table 9.2 reports the observed LIBOR rates for settlement of the interest side of the transaction at the end of each quarter. Column 3 reports the level of the FTSE 100 index used in calculating the return on the equity side of the transaction. The dividend yields, here assumed to be 1.1% p.a. in each quarter, appear in Table 9.3.

In addition Table 9.3 shows the day count for each of the quarters (Column 2) and the percentage change in the index (Column 4) rounded to two decimal places. The reported percentage change is calculated according to equation (9.1).

$$R_{ind,t} = \left[\frac{Index_t}{Index_{t-1}} - 1\right] \times 100 \tag{9.1}$$

where

$R_{ind,t}$ represents the simple return on the index at time t;
$Index_{t-1}$ represents the level of the index at t, and $t-1$.

Table 9.3 presents the quarter-by-quarter outcomes of the swap.

At the end of the first quarter the percentage change in the index, adjusted to include the quarterly dividend yield (D) is, approximately −4.44%. The adjustment is found by taking the nominal annual percentage dividend yield and scaling it by the quarter's day count divided by the number of days in the year. The index return adjusted to take account of the quarterly dividend yield (D) is presented in equation (9.2).

$$\text{Adjusted}\, R_{ind,t} = R_{ind,t} + D \times \frac{\text{Day count}}{365}. \tag{9.2}$$

	Percentage change with div. yield	Index growth	Fixed spread	LIBOR interest	Net flows	Notional principal
Now						£30,000,000
3 months	−4.44	(£1,332,820)	£7,479	£308,527	(£1,633,868)	
6 months	−2.50	(£748,505)	£7,562	£321,370	(£1,062,313)	
9 months	4.55	£1,365,370	£7,479	£327,226	£1,045,623	
12 months	2.81	£843,993	£7,479	£336,575	£514,897	
	Totals:	£128,037	£30,000	£1,293,699	(£1,135,662)	
Return from floating to equity swap:				(£1,135,662)		

Table 9.4

Multiplying the adjusted index return by the notional principal (equation (9.3)) will provide the required transfer from the institution to the fund manager: the result in terms of cash is a flow of £1,332,820 which appears as the first entry in Column 3 of Table 9.3.

$$IG_t = \text{Adjusted } R_{ind,t} \times NP \tag{9.3}$$

where IG_t represents index growth monetary value over period t, and other variables are as previously described.

In addition to this cash flow the transfer of an amount in respect of the fixed spread is required. This is calculated as indicated in equation (9.4).

$$\text{Spread payment} = NP \times (\text{Day count}/365) \times bp \text{ spread} \tag{9.4}$$

where

NP represents the swaps notional principal on which cash flow calculations are based;
bp spread represents the fixed basis point spread.

In the case of the first exchange at the end of month 3 the sum involved is:

$$(\text{£}30{,}000{,}000) \times (91/365) \times (0.001) = \text{£}7479. \tag{9.4a}$$

This sum will be paid to the institution regardless of the directional move of the FTSE 100 index.

The LIBOR indexed cash flow is calculated as presented in equation (9.5) and is clearly similar to the approach used in equation (9.4).

$$\text{LIBOR payment} = NP \times (\text{Day count}/365) \times CPL \tag{9.5}$$

where

NP represents the swaps notional principal on which cash flow calculations are based;
CPL represents the current LIBOR rate, appropriate for the period under consideration, expressed as a decimal. For the first three-month period this translates to a sum of:

$$(\text{£}30{,}000{,}000) \times (91/365) \times (0.04125) = \text{£}308.527. \tag{9.5a}$$

This figure is reported in Table 9.4 as the first entry in Column 5.

Column 6 presents the net cash flow from the perspective of the index receiver (the financial institution) involved in the swap. Rounded to the nearest pound sterling this is a negative flow of £1,633,868 in the first quarter.

The remaining columns in the table are calculated in the same way. Period-by-period net cash flows are reported in Column 6 while the sum of those flows is reported at the foot of Column 6 and again at the bottom of the table. In this particular case the fund manager's view was correct and the decision to swap the unknown, and risky, equity index returns for LIBOR was justified. The gain from the fund manager's perspective is £1,135,662.

9.3 Equity swap: fair pricing

Of course, the example developed above makes no attempt to examine the likely, or anticipated cash flows when setting up the swap. Although there appears to be no universally accepted way of valuing a swap at start-up, one logical approach would be to estimate the forward rates for LIBOR as described earlier in the text and also to estimate the theoretical price for the future by using forward rates and extrapolated dividend yields. In this way a stream of initial cash flows could be established; these streams could then be discounted and used to find a fair margin spread payable to the receiver of the equity returns.

This can be illustrated in the context of Example 9.1. If it is now assumed that at the point of the swap's initiation the market rates listed in Table 9.4 obtain then it is possible to obtain theoretical "fair" futures prices and forward interest rates.

An indicative "fair" futures price can be estimated using equation (9.6).

$$F_{0k} = S_{0o}\left(1 + [r_{0k} - d_{0k}] \times \frac{\text{Day count}}{365}\right) \tag{9.6}$$

where

- F_{0k} represents the price at time zero (now) of a futures contract deliverable at time k;
- S_{0o} represents the price at time zero (now) of the underlying instrument in the cash market;
- r_{0k} represents the risk-free rate of interest appropriate over the period now to time k, expressed as a decimal;
- d_{0k} represents the dividend yield associated with the underlying instrument over the period now to k, expressed as a decimal.

If k is taken successively to be 3, 6, 9 and 12 months then, using the spot zero rates, the current cash market quote given in Table 9.5, and assuming a constant dividend yield of 1.1%, the fair futures prices can be estimated for each of the periods. A demonstration of the approach appears as equation (9.6a). The outcomes are reported in Column 3 of Table 9.6

$$F_{03} = 5300 \times \left(1 + [0.04125 - 0.011] \times \frac{91}{365}\right) \cong 5340 \tag{9.6a}$$

while the implied three-month forward interest rates can be calculated using equation (9.7):

$$r_{ik} = \left[\frac{\left(1 + r_{0k}\dfrac{\text{dc}_k}{365}\right)}{\left(1 + r_{0k-3}\dfrac{\text{dc}_{k-3}}{365}\right)} - 1\right]\frac{365}{\text{dc}_{ik}}. \tag{9.7}$$

Period (quarters)	**LIBOR (%)**	**FTSE 100 index**
r_{03}	4.125	5300
r_{06}	4.250	
r_{09}	4.375	
r_{012}	4.500	

Table 9.5

Period (quarters)	Forward rates (%)	Future "fair" index value
End Quarter 1	r_{03} 4.1250	5340
End Quarter 2	r_{36} 4.3291	5384
End Quarter 3	r_{69} 4.5299	5430
End Quarter 4	r_{912} 4.7213	5480

Table 9.6

	Day count	FTSE 100	Percentage change	Fixed spread in bps	LIBOR	Dividend yield p.a.
Now		5300		0	4.1250%	1.100%
3 months	91	5340	0.75		4.1250%	1.100%
6 months	92	5384	0.82		4.2500%	1.100%
9 months	91	5430	0.85		4.3750%	1.100%
12 months	91	5480	0.92		4.5000%	1.100%

Table 9.7

If k is taken successively to be 3, 6, 9 and 12 months then, using the spot zero rates the implied forward rates can be estimated for each of the periods. A demonstration of the approach for the 6–9 forward rate appears as equation (9.7a). The outcomes are reported in Column 3 of Table 9.6.

$$r_{69} = \left[\frac{\left(1 + 0.04375 \dfrac{274}{365}\right)}{\left(1 + 0.04250 \dfrac{183}{365}\right)} - 1 \right] \frac{365}{91} = 0.0452985 \text{ or } 4.52985\%. \tag{9.7a}$$

These rates are collected together and reported in Table 9.7, in which the basis point spread in Column 5 is set at 0. The value that it should take is at present unknown. A calculation needs to be performed which, when discounted and added to the quarterly, discounted, equity index payments, will ensure that the value of the swap, to each of the counterparties will be zero when the transaction is initiated.

Using Excel's **Goal Seek** from the **Tools menu** on the spreadsheet entitled **Equity Swap**, it is easy to establish to what this basis point spread should be set in order to equate the sum of the discounted cash flows. Table 9.9 is essentially Table 9.7 repeated but with the *fair price* required basis point spread displayed.

The initial cash flow position prior to running "goal seek" is illustrated in Table 9.8. With the basis point spread set at zero the counterparty receiving floating LIBOR is not as well off as the equity index receiver. The LIBOR receiver loses £8058 on the swap as it is structured at present.

In Table 9.7 the percentage change in the index together with the dividend yield reported in Column 2 are calculated in exactly the same way as described in equation (9.2). The index growth values reported in Column 3 are arrived at in the same manner as

	Percentage change with div. yield	Index growth	Fixed spread	LIBOR interest	Net flows	Notional principal
Now						£30,000,000
3 months	1.0289636	£305,547	£0	£305,387	£160	
6 months	1.1012303	£323,476	£0	£320,523	£2,954	
9 months	1.1286299	£327,822	£0	£328,035	(£212)	
12 months	1.1950569	£343,079	£0	£337,922	£5,157	
	Totals:	£1,299,924	£0	£1,291,866	£8,058	
Return from floating to equity swap:				£8,058		

Table 9.8

described in respect of Table 9.4, with the modification that the sum involved is now discounted by the three-month spot rate.

$$IG_t = \frac{\text{Adjusted } R_{ind,t} \times NP}{\left(1 + r_{0k}\dfrac{\text{Day count}}{365}\right)} \tag{9.8}$$

where

IG_t represents index growth over period t;
r_{0k} represents the appropriate kth period nominal annual spot zero rate, as a decimal.

Both the fixed spread and the LIBOR flow columns will need to be discounted in a similar manner to that indicated in equation (9.8) but, in addition, the LIBOR cash flow column will need to be calculated on the basis of the implied forward rates. Equation (9.9) presents formally how this discounting process would be achieved in the context of the fixed spread column, while equation (9.10) presents the process in respect of the interest rate flows.

$$\text{Spread payment} = \frac{NP \times (\text{Day count}/365) \times bp \text{ spread}}{\left(1 + r_{0k}\dfrac{\text{Day count}}{365}\right)} \tag{9.9}$$

$$\text{LIBOR payment} = \frac{NP \times (\text{Day count}/365) \times IF_{ik}}{\left(1 + r_{0k}\dfrac{\text{Day count}}{365}\right)} \tag{9.10}$$

where

NP represents the swaps notional principal on which cash flow calculations are based;
IF_{ik} represents the implied forward rate for the ith period forward until the kth period;
r_{0k} represents the appropriate kth period nominal annual spot rate, as a decimal.

After adjusting the cash flows to take account of the calculated bp spread the value of the swap to each counterparty is zero as required. The results appear in Table 9.10.

If this revised valuation approach is adopted for the original equity index figures in Table 9.10, the overall outcome is well-balanced to both parties. The index receiver pays

	Count		Change	Spread in bps		Yield p.a.
Now		5300		−2.759046364	4.1250%	1.100%
3 months	91	5340	0.75		4.1250%	1.100%
6 months	92	5384	0.82		4.2500%	1.100%
9 months	91	5430	0.85		4.3750%	1.100%
12 months	91	5480	0.92		4.5000%	1.100%

Table 9.9

	Percentage change with div. yield	Index growth	Fixed spread	LIBOR interest	Net flows	Notional principal
Now						£30,000,000
3 months	1.0289636	£305,547	(£2,043)	£305,387	(£1,883)	
6 months	1.1012303	£323,476	(£2,043)	£320,523	£911	
9 months	1.1286299	£327,822	(£1,998)	£328,035	(£2,210)	
12 months	1.1950569	£343,079	(£1,975)	£337,922	£3,182	
	Totals:	£1,299,924	(£8,058)	£1,291,866	(£0)	
Return from floating to equity swap:				£0		

Table 9.10

a fixed spread of approximately 2.76 bps to the index payer each quarter. In this way the swap is structured to be fair to both counterparties at inception.

An alternative to this method of valuing the swap would be to obtain simulated values of the index at the relevant future dates. This type of approach, if it allowed for time-varying returns and/or stochastic volatility, would lead to scenarios where the index could fall rather than rise monotonically over time when the risk-free rate of interest is larger than the projected dividend yield.

Variations of this type of single-currency equity swap might be exchange of index growth against receipt of a fixed rate of interest, or on a zero-coupon basis or swapping index growth or decline but with dividends ignored, or a floating notional principal determined by adding or subtracting the previous quarter's growth or decline from the current level of the contract's notional principal, or using the greater growth factor from two domestic indices. Many variations are possible.

10 Equity and Equity Index Options

Equity options provide a special type of vehicle for creating and manipulating payout profiles. The profile is quite different from that used to illustrate the profit/loss regions associated with long and short futures positions. Options possess a variety of features that make them extremely attractive to many participants in financial, and indeed, other markets. It is these features that have led to the generation of a vast amount of literature in academic, marketing sources and even the popular press – particularly at times when option strategies go wrong in a big way, for example, Barings and Long Term Capital Management (LTCM). In essence, options offer market players a great deal of flexibility and enable them to construct quite complicated financial positions involving equity, interest rate, currency and commodities in an individual context, using mixtures of instruments based on products in both the cash and futures markets.

What makes the application of options so attractive is that they can be regarded as a set of easily understood building blocks that can be combined in many ways to engineer positions that will deliver a desired payout at some defined future point in time. As this chapter develops the reader will be introduced to some option strategies together with an analysis of their rationale.

Some of the characteristics that have attracted attention over recent years, and which have stimulated research and debate, relate to the following issues:

- pricing of options – selling a product
- mispricing of options – undertaking arbitrage
- using naked option positions – speculation
- risk hedging
- financial engineering.

1. Pricing of options – selling a product

There are several approaches to option pricing. In 1973 the pioneering work of Fischer Black and Myron Scholes, and Robert Merton was published. Their research into the pricing of equity options directly opened the portal to a new era in the study of finance. Since their pathfinding work many model variations have been proposed to take account of the special features inherent in interest rate products, currencies and options on futures. However, the formulae developed in their early published work are, perhaps, still the most widely known. Given the wide acceptance of the Black and Scholes (BS) pricing framework, and the ease with which the option price and important sensitivities can be calculated, their approach provides a good benchmark against which other pricing models can be compared. It must be stated, however, that there are problems associated with the BS model and particularly the way that its variants and extensions are being used in today's financial markets.

Keating *et al.* (2001) provide a good overview of the abuse to which the BS pricing framework has been subjected.

2. Mispricing of options – undertaking arbitrage

One way of judging whether an option, or for that matter any other instrument, is mispriced is to replicate its structure. This is the concept of arbitrage-free or risk-neutral valuation. Later in this chapter the payout profile of a long underlying security combined with a long put option will be shown to generate the profile of a long call option. If, for example, a call option is deemed to be too expensive then, if the over-pricing is significant enough, the call could be synthesised using other instruments. The synthetically created call would be the relatively cheap instrument held in a portfolio against a short, expensive call position. There are many examples where this type of approach could be adopted in today's markets, indeed some specialist financial boutiques and hedge funds make good profits by trading in this way.

3. Using naked option positions – speculation

Using either equity index options or options on individual shares, long or short naked option positions can be used to take advantage of an anticipated rapid move in the price of the underlying instrument. The term naked simply means that the underlying instrument is not held in the investor's portfolio. In the case of short positions this is particularly important.

In the event that the option is exercised by the option holder, the option writer must go to the market and buy the stock to deliver into the contract. This will almost certainly prove to be an expensive business. The option writer will need to pay the transactions costs associated with the option position, the purchase price of the associated shares together with their transactions costs, and on top of that any associated institutional costs. There are many texts where these ideas can be examined, for example, the classic text of Cox and Rubinstein (1985), and more recently Hull (2000).

4. Risk hedging

In this case equity index or individual options can be combined with long or short positions to protect against the risk of financial loss. If an investor suspects that the value of a share in his/her portfolio is about to fall it would make sense to insure against the potential loss. Thus a long position in the shares of, say, British Telecom plc (BT) could be hedged by purchasing a BT put. This action creates a strategy that pays out when the share price in BT falls below a specified level. An example of this type of hedging will be presented later in this chapter.

In the case of an index-tracking portfolio the investor can use exchange-based options in a similar way to that described in the BT example above. An important difference in the operational characteristics of the index and individual stock contracts lies in the mode of settlement. In the case of individual shares the security named in the contract will have to be physically delivered when the option is exercised at expiration. For index options settlement takes the form of cash rather than the underlying portfolio of shares. Eales (2000) provides a good example of this type of hedging with an accompanying Excel spreadsheet.

An alternative perspective to that of the end user seeking insurance is, of course, provided by the financial institutions, if they are providing the options with which an investor is hedging the risk of a share or a portfolio, they too will need to protect their position. In fact options will appear as long and as short positions in the institution's portfolio. The risk of this portfolio will need to be managed and to achieve this those responsible for managing the risk will use option price sensitivities about which more will be said later. Texts which provide insight into the mechanics of hedging portfolios of options can be found in books by Hull (2000), and Taleb (1997).

5. Financial engineering
In one sense financial engineering is implicit in the process of hedging the BT share discussed above. The arithmetic payout derived from a combination of a long position in the BT share and an option will generate a new set of potential rewards or losses, which will be more satisfactory, from the investor's point of view, than the original long position in the BT share. This same sort of payout engineering can be applied equally well to index options, too.

Using options in conjunction with other instruments will enable new classes of instruments to be developed and these instruments in turn can be used to satisfy the demands of the investing public.

In essence this chapter will concentrate mainly on Items 3 and 4 in the list above, i.e. the use of options as hedge instruments, but will also touch on the idea of financial engineering. Indeed, with some adjustments the points made in this discussion can be carried over almost completely to the uses of options in the management of risk involved in interest rates, bonds, currencies, and indeed commodities contexts. However, before developing the fundamental operational profiles of options, and applying them to risk management/engineering cases in detail, it is essential to examine and understand the fundamentals of options.

Basically there are two types of options available: calls and puts. A call option gives the holder (the purchaser of the option) the right to buy an underlying instrument while a put option gives the holder the right to sell an underlying instrument. Initially this section will concentrate on the idea of call options on individual shares (stocks), and on developing an understanding of option terminology. As the chapter is developed put options on individual shares and put options on indexes will be examined.

It will be useful to start by establishing a working definition of a call option.

10.1 Call options

Call options (European-style) give the holder the right, but not the obligation, to buy a fixed number of shares of a named company at a fixed price on a specified date.

The title European does not refer to where the option is traded – London, Frankfurt/Main, Paris, Amsterdam, etc. – but to the fact that the right to buy the shares at a fixed price is only applicable on the date specified in the contract. Although, over its lifetime, it may be possible to trade the option on an exchange or over the counter as an instrument in its own right, actual transfer into the underlying shares can only take place on the date specified in the contract. The specified date in a European-style contract is the *expiration date* of the contract, i.e. the date on which the contract matures. The purchaser of a call option is taking a long position and is called the *holder* of an option.

The fixed price at which exercise takes place is called either the *strike* or the *exercise* price. Making use of this right to buy is called exercising – exercising the right to buy the shares at the strike price. If an option contract is held under expiration but the holder chooses not to exercise the right to buy on that date the contract expires worthless.

Call options (American-style) give the holder the right, but not the obligation, to buy a fixed number of shares of a named company at a fixed price on *or before* the expiration date of the contract.

The title American indicates the flexibility that the holder of the option has in deciding when to exercise the right to take up the underlying instrument at the strike price. Equity options on the LIFFE exchange are American-style options. The number of options available on individual equity varies. The most up-to-date information can be found by referring to the web pages of the derivative exchanges. The LIFFE web page can be found by using www.liffe.com and EUREX by using www.eurexchange.com. Indeed, the web addresses of many exchanges throughout the world can be found by using www.site-by-site.com. Each individual equity option contract on LIFFE is normally for 1000 shares in the company concerned. A specimen contract specification for options on individual stocks appears in Table 10.1. There are, however, exceptions to this general rule. At the time of writing only option contracts in respect of AstraZeneca plc and Celltech plc are for 100 shares.

It can be seen from the contract summary in Table 10.1 that there are three main cycles listed: January, February and March. This does not mean that each and every share will have options expiring in each of those months. Shares in a particular company belong to a specific cycle. For example: Allied Domecq plc belongs to the January cycle. It has options

Unit of trading	One option normally equals rights over 1000 or 100 shares
Expiry months	**January cycle (J):** means the three nearest expiry months from Jan, Apr, Jul, Oct cycle **February cycle (F):** means the three nearest expiry months from Feb, May, Aug, Nov cycle **March cycle (M):** means the three nearest expiry months from Mar, Jun, Sep, Dec cycle
Exercise	Exercise by 17:20 on any business day, extended to 18:00 for all series on the last trading day
Last trading day	16:30 (London time) third Wednesday in expiry month
Settlement day	Settlement day is six business days following the day of exercise/last trading day
Quotation	Pence/share
Minimum price movement (tick size & value)	0.5 pence/share (£5.00)
LIFFE CONNECT™ Trading hours	08:00–16:30 (London time)
Daily settlement time	16:30 (London time)

Contract standard: Delivery will be 1000 shares (or other such number of shares as determined by the terms of the contract).

Option premium: is payable in full by the buyer on the business day following a transaction.

Exercise price and exercise price intervals: The interval between the exercise prices is set according to a fixed scale determined by the exchange.

Introduction of new exercise prices: Additional exercise prices will be introduced after the underlying share price has exceeded the second highest, or fallen below the second lowest available exercise price.

Table 10.1: Equity options

with expiry dates in January, April, July and October. Imperial Tobacco Group plc belongs to the February cycle. It has options that expire in February, May, August and November. Cable and Wireless plc belongs to the March cycle. It has options that expire in March, June, September and December. Last trading day and settlement day are clearly defined.

In the context of index options the contract takes on a similar appearance but is not assigned to a particular delivery cycle. Which contracts may be traded on an exchange at any time are defined in the contract specification, an example of which appears in Table 10.2.

An example of the way in which equity option quotes are reported by specialist providers is shown in Figure 10.1. The quotes in Figure 10.1 are reproduced from a Bloomberg screen and reflect the prices of Lloyds TSB Group plc options for different expiration months at a wide range of exercise prices.

In Figure 10.1 the underlying instrument in question is the stock (shares) of Lloyds TSB Group plc. The contract is a standard LIFFE contract which is to say that the buyer or holder of the call option is paying for the right to buy 1000 shares in the named company. As described in the contract specification the option prices quoted are all in pence per share. The price quoted for an option is known as the premium.

From Figure 10.1 there are strike prices available for the Dec contract from 390 to 900 but notice that bid–ask prices are only quoted for strikes of 600 to 900. The potential buyer is free to choose at what exercise price the contract will be struck but most of the trading

Contract size	Valued at £10 per index point (e.g. value £65,000 at 6500.0)
Delivery months	March, June, September and December plus such additional months that the three nearest calendar months are always available for trading
Exercise day	Exercise by 18:00 on the last trading day only
Last trading day	10:30:30 (London time) third Friday in delivery month[a]
Settlement day	Settlement day is the first business day after the expiry date
Quotation	Index points (e.g. 6500.0)
Minimum price movement (tick size & value)	0.5 (£5.00)
Trading hours	08:00–16:30 (London time)
Daily settlement time	16:30 (London time)

Contract standard: Cash settlement based on the exchange delivery settlement price.

Exchange delivery settlement price (EDSP): The EDSP is based on the average values of the FTSE 100 index every 15 seconds between (and including) 10:10 and 10:30 (London time) on the last trading day. Of the 81 measured values, the highest 12 and lowest 12 will be discarded and the remaining 57 will be averaged to calculate the EDSP. Where necessary, the calculation will be rounded to the nearest half index point.

Option premium: is payable by the buyer in full on the business day following a transaction.

Exercise price and exercise price intervals: The interval between the exercise prices is determined by the time to maturity of a particular expiry month and is either 50 or 100 index points. The exchange reserves the right to introduce tighter strike intervals (e.g. 25 points) where necessary.

Introduction of new exercise prices: Additional exercise prices will be introduced after the underlying index level has exceeded the second highest, or fallen below the second lowest available exercise price.

Table 10.2: FTSE 100 index option (European-style exercise)

[a] In the event of the third Friday not being a business day, the last trading day shall normally be the last business day preceding the third Friday.

GRAB **Equity CALL**

CALL SUMMARY PAGE Hit 99 <go> for option model defaults

LLOY LN LLOYDS TSB GROUP PLC 11:31 Hedge: Y

717 +5 Bid 714 Ask 717 Days 54 Fin Rate 4.16 Volat: 41.51

LLOY Dec 01	LAST TRADE INFO Bid	Ask	Last	Chnge	IMPL.VOLAT Bid	Ask	HEDGE Del	Gam	Vega	*THEORETICAL Price	Diff
1) 12 C390			324.50	unch	b.int	b.int	1000*			329.40	
2) 12 C420			295.00	unch	b.int	b.int	1000*			299.59	
3) 12 C460			255.50	unch	b.int	b.int	998*		.02*	259.90	
4) 12 C500			216.50	unch	b.int	b.int	991*		.07*	220.43	
5) 12 C550			169.50	unch	b.int	b.int	962*	001*	.24*	172.23	
6) 12 C600	122.00	128.00	125.50	unch	28.20	43.14	914	002	.44	127.14	+2.14
7) 12 C650	82.00	87.00	92.00	+6.00	34.29	40.53	787	003	.81	87.82	+3.32
8) 12 C700	49.00	54.00	58.50	+5.50	34.46	39.18	612	004	1.06	56.47	+4.97
9) 12 C750	25.00	28.50	27.50	unch	33.39	36.64	413	004	1.08	33.77	+7.02
10) 12 C800	10.50	14.50	12.50	-.50	32.07	36.76	238	003	.86	18.84	+6.34
11) 12 C850	3.50	8.00	6.50	+1.00	30.76	38.77	126	002	.58	9.86	+4.11
12) 12 C900	.50	4.00	2.00	unch	26.99	39.59	057	001	.33	4.87	+2.62
LLOY Mar 02					Days 145		Fin Rate 4.08		Volat: 39.25		
13) 3 C460			260.50	unch	b.int	b.int	985*		.17*	264.34	
14) 3 C500			224.00	unch	b.int	b.int	962*	001*	.36*	226.59	
15) 3 C550			181.00	unch	b.int	b.int	911*	001*	.70*	181.96	
16) 3 C600	136.00	142.00	140.50	unch	33.63	39.57	844	002	1.03	141.66	+2.66

Figure 10.1: BBG option quote page. © Bloomberg LP. Reproduced with permission.

activity will take place around the current value of the underlying share price. As the underlying share price moves the reported strikes will be adjusted automatically to stay in line with movements in the underlying share price. It would not be very useful for a potential hedger to know the call option price of a strike of 800 when, say, the current price of the underlying share is 200. The mechanism for creating new strikes is defined by an exchange and will be described in the contract specification.

Additionally, the potential buyer can also choose to enter into a contract for different expiration dates as shown in Figure 10.1; these could be Dec, Mar, Jun (not shown).

When a call option is purchased a long position is taken up, and this may result in the holder of that option taking delivery of 1000 shares at or before a specified date in the future. Clearly, from the discussion above, there are a number of contracts on offer at any time with different strike prices and different delivery months – however, the contract(s) selected for purchase will depend on a number of factors, not least the role being played by the buyer/speculator/hedger/arbitrageur with an eye on the current market price of the option(s) and the time remaining to expiration, and, of course, expectation about price changes in the value of the underlying Lloyds TSB share price. To provide an example, take the case where the call option purchased is the Dec contract with an exercise price of 700. Reading the quoted ASK price of 54 from Figure 10.1 and bearing in mind that one contract is for 1000 shares the basic cost of a contract will be (this ignores commissions and any institutional cost that may apply):

$(54)(1000)/100 = £540.$

The discussion so far has concentrated on exchange-based contracts that have defined strike and expiration dates. For some time CME and LIFFE have offered contracts where the purchaser has some scope for varying the standard contract specification. These contracts allow the user to specify a strike price and/or an expiration date, which does not match those laid down by the exchange. The contracts are labelled FLEX® option contracts and are European in style. The contract specification for this style of contract appears in Table 10.3.

In a more international framework many equity fund managers now track portfolios of equity diversified over Europe-wide markets. Benchmark indexes have been specially constructed by several exchanges to allow fund managers to hedge their portfolios. LIFFE, for example, offers option contracts on the Eurotop 100, 200, etc. indexes as well as Morgan Stanley (MSCI) benchmarks such as the European index, Euro-, and Pan-Euro indexes. EUREX offers option contracts on Dow Jones Euro STOXX℠ 50, Dow Jones EURO STOXX℠ sector index options as well as options on the Swiss and Finnish market indexes. The composition of each index is naturally varied but all of the contracts are denominated in euros. Typical EUREX Dow Jones and LIFFE MSCI option contracts appear in Tables 10.4 and 10.5, respectively.

Contract size	Valued at £10 per index point (e.g. value £65,000 at 6500.0)
Delivery months	Any business day within two years from the date the contract is made
Exercise day	Exercise by 18:00 on the last trading day only
Last trading day	10:30:30 (London time) on the expiry date
Settlement day	Settlement day is the first business day after the expiry date
Quotation	Index points (e.g. 6500.0)
Minimum price movement (tick size & value)	0.5 (£5.00)
Trading hours	08:15–16:15 (London time)

Contract standard: Cash settlement based on the exchange delivery settlement price.

Exchange delivery settlement price (EDSP): The EDSP is based on the average values of the FTSE 100 index every 15 seconds between (and including) 10:10 and 10:30 (London time) on the last trading day. Of the 81 measured values, the highest 12 and lowest 12 will be discarded and the remaining 57 will be averaged to calculate the EDSP. Where necessary, the calculation will be rounded to the nearest half index point.

Option premium: is payable by the buyer in full on the business day following a transaction.

Exercise price and exercise price intervals: are agreed by the parties to the contract. A FTSE 100 index FLEX® option contract may not be made with an exercise price and expiry date which are the same as the exercise price and last trading day of an existing series of the standard FTSE 100 index option (European-style exercise).

Novation: If a standard FTSE 100 index option (European-style exercise) series is made available for trading with an exercise price and last trading day which are the same as the exercise price and expiry date of an existing FTSE 100 index FLEX® option contract, then the contract shall be replaced by novation by a standard FTSE 100 index option (European-style exercise) contract.

Table 10.3: FTSE 100 index FLEX® option (European-style exercise)
FLEX® is a registered trademark of the Chicago Board Options Exchange Inc. and has been licensed for use by LIFFE Administration and Management. Source: LIFFE

Contract size	Valued at €10 per index point (e.g. value €30,000 at 3000.0)
Delivery months	March, June, September and December plus such additional months that the three nearest calendar months are always available for trading
Exercise day	Exercise by 18:00 on the last trading day only
Last trading day	13:00:30 (CET), 12:00:30 (London time) third Friday in delivery month[a]
Settlement day	Settlement day is the first business day after the last trading day
Quotation	Index points (e.g. 3000.0)
Minimum price movement (tick size & value)	0.1 (€1)
LIFFE CONNECT™ trading hours	10:00–17:00 (CET), 09:00–16:00 (London time)
Daily settlement time	17:30 (CET), 16:30 (London time)

Contract standard: Cash settlement based on the exchange delivery settlement price.
Exchange delivery settlement price (EDSP): The EDSP is based on the average values of the index[b] every 15 seconds between (and including) 11:40 and 12:00 (London time) on the last trading day. Of the 81 measured values, the highest 12 and lowest 12 will be discarded and the remaining 57 will be averaged to calculate the EDSP. Where necessary, the calculation will be rounded to the nearest half index point.
Option premium: is payable by the buyer in full on the business day following a transaction.
Exercise price and exercise price intervals: The interval between the exercise prices is determined by the time to maturity of a particular expiry month and is either 50 or 100 index points. The exchange reserves the right to introduce tighter strike intervals (e.g. 25 points) where necessary.
Introduction of new exercise prices: Additional exercise prices will be introduced after the underlying index level has exceeded the second highest, or fallen below the second lowest available exercise price.

Table 10.4: MSCI Euro index and MSCI Pan-Euro index options (European-style exercise)
[a] In the event of the third Friday not being a business day, the last trading day shall normally be the last business day preceding the third Friday.
[b] In this instance, "the index" refers to the cash value of either the MSCI Euro or MSCI Pan-Euro indexes, as applicable, so as to reflect the underlying index value of the options contract. Source: LIFFE

Before considering the call summary page (Figure 10.1) and the figures it contains in any more detail, it will be useful to examine another type of option that is available: a *put* option.

10.2 Put options

Put options (European-style) give the holder the right, but not the obligation, to sell a fixed number of shares of a named company at a fixed price on a specified date. Like their counterpart, the European call option, European puts can only be exercised at contract expiration.

Put options (American-style) give the holder the right, but not the obligation, to sell a fixed number of shares of a named company at a fixed price on or before a specified date.

Contract standard
The Dow Jones STOXX^{SM} 50 index (for Dow Jones STOXX^{SM} 50 futures).
The Dow Jones Euro STOXX^{SM} 50 index (for Dow Jones Euro STOXX^{SM} 50 futures).
Contract value
€10 per Dow Jones STOXX^{SM} 50 and Dow Jones Euro STOXX^{SM} 50 index point.
Settlement
Cash settlement based on the final settlement price, payable on the first exchange trading day immediately following the last trading day.
Quotation
In points, with no decimal places.
Minimum price movement
1 point, representing a value of €10.
Contract terms
The three nearest months within the cycle March, June, September, December.
Last trading day
The third Friday of the expiration month, if that is an exchange trading day; otherwise, on the exchange trading day immediately prior to that Friday. Trading ceases in the expiring futures contract at 12:00 noon CET on the last trading day.
Daily settlement price
The last-paid price of the trading day; or, if the last-paid price is older than 15 minutes or does not reasonably reflect actual market conditions, then Eurex will establish the official settlement price.
Final settlement price
The average price of the Dow Jones STOXX^{SM} 50 and Dow Jones Euro STOXX^{SM} 50 index calculations performed between 11:50 a.m. and 12:00 noon CET on the last trading day. The final settlement price is determined at 12:00 noon CET on the last trading day.
Trading hours
Dow Jones STOXX^{SM} 50 index: 9:00 a.m. until 5:30 p.m. CET.
Dow Jones Euro STOXX^{SM} 50 index: 9:00 a.m. until 8:00 p.m. CET.

Table 10.5: Dow Jones STOXX^{SM} 50 futures (FSTX), Dow Jones Euro STOXX^{SM} 50 futures (FESX). Source: LIFFE

The title American once again indicates the flexibility that the holder of the option has in deciding when to exercise the right to sell the underlying instrument at the selected strike price.

The interpretation of the put summary page (Figure 10.2) is analogous to that of the call table. The contract size is for 1000 shares. The quotes are all in terms of pence per share. So, by way of example, the premium to purchase a Dec 700 put is 38 pence per share and the contract for 1000 shares would cost £380.

Returning now to a more detailed look at the contents of the call and put tables, the question can be raised: what accounts for the differences in the costs of each contract?

It would appear that at £110 the Dec 600 put contract is a bargain compared to the £1280 that has to be paid for the Dec 600 call contract. The answer to this is quite simply that the instruments being compared are different. Although the strike prices and contract dates are the same the functional characteristics of the two contracts are not. They are not offering the same transaction, one instrument – the call option – is bestowing the right to buy Lloyds TSB Group shares while the put option is bestowing the right to sell Lloyds TSB Group shares.

GRAB — Equity PUT

PUT SUMMARY PAGE — Hit 99 <go> for option model defaults

LLOY LN LLOYDS TSB GROUP PLC 11:31 Hedge:Y

717 +5 Bid 714 Ask 717.5 Days 54 Fin Rate 4.16 Volat: 41.51

LLOY Dec 01	LAST TRADE INFO Bid	Ask	Last	Chnge	IMPL.VOLAT Bid	Ask	HEDGE Del	Gam	Vega	*THEORETICAL Price	Diff
1) 12 P390			.50	unch	b.int	b.int				.00	
2) 12 P420			.50	unch	b.int	b.int				.01	
3) 12 P460			.50	unch	b.int	b.int	002*		.02*	.06	
4) 12 P500			1.00	unch	b.int	b.int	009*		.07*	.35	
5) 12 P550	2.00	6.00	3.50	unch	42.19	54.30	064	001	.35	1.84	-2.16
6) 12 P600	7.00	11.00	9.00	unch	42.55	49.35	132	002	.59	6.47	-2.53
7) 12 P650	16.00	20.00	16.00	-3.50	40.39	45.07	240	003	.86	16.92	-1.08
8) 12 P700	33.00	38.00	32.50	-4.50	39.21	43.90	398	003	1.07	35.45	-.05
9) 12 P750	58.00	63.00	59.00	-2.50	37.16	41.83	578	004	1.08	62.65	+2.15
10) 12 P800	93.00	99.00	97.00	unch	35.99	42.76	739	003	.90	97.87	+1.87
11) 12 P850	138.00	142.00	140.00	unch	39.51	45.47	835	002	.69	139.24	-.76
12) 12 P900	186.00	190.00	184.50	-3.50	44.29	51.80	879	002	.56	184.84	-3.16

LLOY Mar 02 Days 145 Fin Rate 4.08 Volat: 39.25

13) 3 P460			3.00	unch	b.int	b.int	032*		.33*	2.35	
14) 3 P500			7.00	unch	b.int	b.int	064*	001*	.57*	5.38	
15) 3 P550	14.00	18.50	15.00	unch	40.82	45.34	144	001	1.02	12.46	-3.79
16) 3 P600	24.00	28.50	26.00	unch	38.86	42.31	219	002	1.32	24.49	-1.76

Figure 10.2: BBG put option page. © Bloomberg LP. Reproduced with permission.

Example 10.1: Calls versus puts

Assume in connection with the current market price of 717p, that the Dec 600 call costs the same as a Dec 600 put, namely 11p per share. If this case were to hold then a shrewd investor would buy a call contract for £110, and immediately exercise the right to buy the underlying share at 600 (£6000) and immediately sell the shares at the market price of 717 (£7170). This would result in a risk-free profit of:

£7170	Received from sale of shares
£6000	Less cost of buying shares through the call option contract
£110	Less option premium
£1060	Profit

If the same strategy is applied to a put option the investor would buy a put option with a strike of 600 and would simultaneously buy the 1000 shares at 717p per share to deliver into the contract. The outcome would be:

£6000	Received from exercising the put option
£7170	Less cost of purchasing shares at current market price for delivery into the put option contract
£110	Less cost of option contract
(£1280)	Outcome: a loss per contract of £1280, ignoring transactions and other institutional costs

Obviously this approach of using calls for getting rich quickly or poor quickly does not exist in the market. The different premiums for call and put options reflect the different characteristics of the instruments. The call is a contract to buy, in this example at 600, whereas the put is a contract to sell at 600. The premium differential between the 600 Dec call and 600 Dec put takes this into account. The fact that the market price of the share lies above the strike price of 600 implies that a call already has some value (117), which will have to be reflected in the premium, while the put has no value.

This argument is in part a satisfactory but incomplete answer to why all the premium differentials exist. Perhaps a fairer comparison would be to consider either call (or put) options with the same expiration dates but with different strikes independently.

Example 10.2: Dec 700 call versus Dec 800 call

Take the case of the Dec call contracts:

- The Dec 700 call costs 54.
- The Dec 800 costs 14.5.

Why should this 39.5 pence per share option premium differential exist?

The answer once again lies in the position of the strike price. Premium differentials are in part determined by whether the strike price lies above or below the current price of the underlying security and how distant the strike price is from the current price of the security. In this example the 700 strike is below the current market price of 717 while the 800 strike is above the current market price. Hence, because of the need to buy in value, the purchaser of the Dec 700 call contract must expect to pay a higher premium than the purchaser of the Dec 800 call contract, which, given the market price of the share, has no value.

If no premium differential existed between the Dec 700 and Dec 800 calls a lucrative strategy would be to purchase the Dec 700 for £145, exercise immediately into the underlying share at the contracted exercise price of 700 pence per share (£7000) and immediately sell the shares received to the market at the current 717p (£7170). This would enable the investor to make a risk-free profit of:

$$(\pounds 7170 - \pounds 7000 - \pounds 145) = \pounds 25 \text{ per contract}$$

once again ignoring transaction and other institutional costs.

Clearly then the premium of a call option with a strike price below the current market price of the security will be higher than the premium of a call option with a strike closer to the current market price, while an option with a strike price above the current market price will have a lower premium. Note that in the case of put options the reverse will be true. This difference between the current market price and a strike price is known as the *intrinsic value* of an option. For a call option this can be found by:

$$(\text{Current market price} - \text{Strike price}) = \text{Call option intrinsic value}$$
$$(S_T - E) = VC_T$$

where

S_T	is the current market price of the share;
E	is the exercise or strike price;
VC_t	is the value of the call at time t.

So for the Dec 700 call option the intrinsic value will be:

$$(717 - 700) = 17$$

an intrinsic value of 17 pence per share. This amount must be reflected in the premium charged for the call option.

For the Dec 800 call option the intrinsic value will be:

$$\text{(Current market price} - \text{Strike price)} = \text{Option intrinsic value}$$
$$(717 - 800) = -83.$$

But recall that the call (put) option contract gives the holder or buyer of the option the right but not the obligation to buy (sell) the underlying share at the contracted price. In this case the strike price lies well above the current market price of the share and to exercise the right to buy would result in a loss of 83 pence per share. The option has in effect no intrinsic value and a rational investor would not choose to exercise the right if this situation holds at the contract's expiration. Since there is no obligation to buy the underlying security the worst possible case that needs to be considered is where the option is allowed to expire worthless. Thus the lowest intrinsic value of an option is zero.

This leads to the formulation of a simple decision rule to determine the expiration value of a European-style call option:[1]

$$C_T = \max(S_T - E, 0)$$

where

S_T	is the current market price of the share;
E	is the exercise or strike price;
C_T	is the value of the call at time t.

An annotated graphical representation of the call profile is illustrated in Figure 10.3.

Notice that for all share prices below 700 at expiration the payout profile is flat and equal to the premium paid by the holder. After 700 has been reached the payout profile rises, breaking even at the strike + premium paid (754), moving more and more into profit as the share price climbs.

In general when the premium that has to be paid is taken into account the choices facing the holder of a call option can be summarised in the following way:

Maximum [Share price – (Strike + Premium) **or** Premium]

$$C_T = \max[(S_T - E) - \text{Premium}, -\text{Premium}].$$

[1] This formula is extremely useful and, with small modifications, can be applied to determine the payouts generated by more exotic-style options.

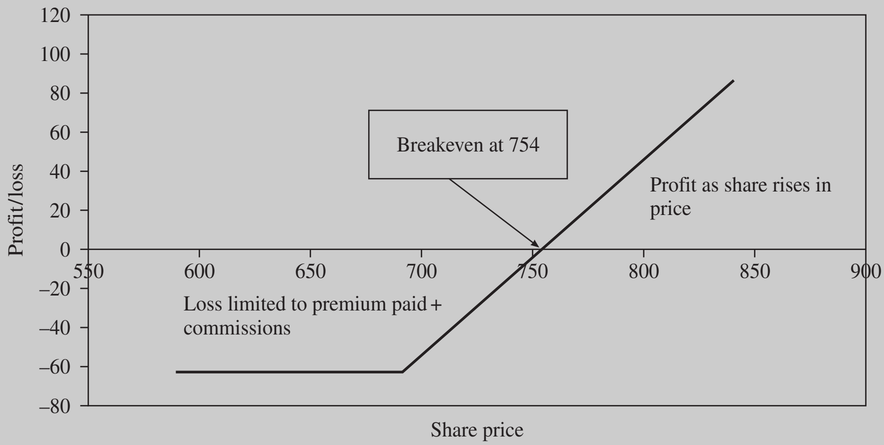

Figure 10.3: Payout profile: long call option

If [Share price – (Strike + Premium)] > [(S_T – E) + Premium] > –Premium then an option holder, acting rationally, will exercise the option and the payout will be:

$$S_T - (E + \text{Premium}).$$

If [Share price – (Strike + Premium)] < [(S_T – E) + Premium] < –Premium (i.e. more negative) then, acting in a rational manner, the holder should not exercise the option in which case the cost involved in this entire transaction will be the premium.

If [Share price – (Strike + Premium)] = [(S_T – E) + Premium] = –Premium then the holder will be indifferent. Whether exercise takes place or not the cost of the position to the holder will be the initial premium outlay.

The process of calculating intrinsic values can easily be extended to cover all strikes in a similar manner. However, another premium differential still needs to be explained: why is the Dec 700 call contract more expensive than the Mar 700 call contract? The reason for the existence of this difference lies in the fact that an option has time value. A portion of the premium that has to be paid for an option covers the time remaining until expiration. Setting up a table of call option premiums for different expiry dates facilitates the calculation of that portion of an options premium which relates to time value.

Note that each of the expiry months displayed in Table 10.6(a) has the same intrinsic value of 17 pence per share.

Having identified the intrinsic value, the time value can now be identified. To achieve this separation the intrinsic value is subtracted from each of the quoted option premiums. This is illustrated in Table 10.6(b) where the time values of 37p, 72p and 94p are calculated for the December, March and June contracts, respectively. All other strike prices could be considered in a similar manner.

In this example for a strike of 750 there is no intrinsic value, as the zeros along the bottom row of Table 10.7(a) indicate. The current price of the underlying security is below the strike price and, therefore, exercise would not be desirable.

	Calls		
	DEC	MAR	JUN
Current market price	717	717	717
Strike price	700	700	700
Intrinsic value	17	17	17

Table 10.6(a): Lloyds TSB Group plc, intrinsic values of 700 call options

	Calls		
	DEC	MAR	JUN
Option premia	54	89	111
Intrinsic value	17	17	17
Time value	37	72	94

Table 10.6(b): Lloyds TSB Group plc, time values of 700 call options

	Calls		
	DEC	MAR	JUN
Current market price	717	717	717
Strike price	750	750	750
Intrinsic value	0	0	0

Table 10.7(a): Lloyds TSB Group plc, intrinsic values of 750 call options

	Calls		
	DEC	MAR	JUN
Option premia	28.5	66	89
Intrinsic value	0	0	0
Time value	28.5	66	89

Table 10.7(b): Lloyds TSB Group plc, time values of 750 call options

The arithmetic in Table 10.7(b) demonstrates that, although no intrinsic value exists, each call option has a time value: 28.5p, 66p and 89p for December, March and June contracts, respectively.

Example 10.3: Puts

Consider now how the arguments developed for call options can be ported over to put option instruments:

- The Dec 700 put costs 38.
- The Dec 800 costs 99 (assumed).

Why should this 61 pence per share option premium differential exist?

As with calls the answer lies in the position of the strike price in relation to the current price of the underlying security and how distant the strike is from the current price of the security. In this example the 700 strike is below the current market price of 717 while the 800 strike is above the current market price.

Hence the purchaser of the Dec 700 put contract must expect to pay a lower premium than the purchaser of the Dec 800 put contract.

Example 10.4

If no premium differential existed between the Dec 700 and Dec 800 puts a lucrative strategy would be to buy 1000 of the underlying shares at the current market price of 717p per share and simultaneously purchase the Dec 800 for £380. By exercising the right to sell the shares immediately at the contracted 800p per share exercise price (£8000) and immediately deliver the 1000 shares purchased at the current 717p (£7170), this strategy would enable the investor to make a risk-free profit of:

£8000	From sale of shares at strike of 800
£7170	Less cost of purchasing shares at market price
£380	Less cost of put option contract
£450	Profit, ignoring transaction and other institutional costs

Clearly then the premium of a put option with a strike price above the current market price of the security will be higher than the premium of a put option with a strike closer to the current market price. Just as with a call the difference between the current market price and a strike price is known as the *intrinsic value* of an option. For a put option this can be found by:

(Strike price − Current market price) = Put option intrinsic value

$$(E - S_t) = VP_t$$

where

S_t is the current market price of the share;
E is the exercise or strike price;
VP_t is the value of the put at time t.

So for the Dec 700 put option the intrinsic value will be:

$$(700 - 717) = -17$$

an intrinsic value of −17 pence per share.

But recall the argument used in the case of a call option and the contractual obligation: a call (put) option contract gives the holder or buyer of the option the right but not the obligation to buy (sell) the underlying share at the contracted price. In this case the strike price lies below the current market price of the share and to exercise the right to sell would result in a loss of 17 pence per share. The option has in effect no intrinsic value and a rational investor would not choose to exercise if this situation were to hold at the contract's expiration. Since there is no obligation to sell the underlying security the worst possible case that needs to be considered is where the option is allowed to expire worthless. Thus the lowest intrinsic value of an option is zero.

For the Dec 800 put option the intrinsic value will be:

$$(\text{Strike price} - \text{Current market price}) = \text{Put option intrinsic value}$$
$$(800 - 717) = 83$$

and this amount must be reflected in the put option's value.

This leads to the formulation of a simple decision rule to determine the expiration value of a European-style put option:[2]

$$P_T = \max(E - S_T, 0).$$

An illustration of the put payout profile for hypothetical values of the underlying security at the option's expiration date is shown in Figure 10.4.

Notice that for all share prices above 700 at expiration the payout profile is flat and equal to the premium paid by the holder. After 700 has been reached the payout profile rises, breaking even at the strike – premium paid (662), moving more and more into profit as the share price falls.

In the case of a put when the premium that has to be paid is taken into account the choices facing the holder of a put option can be summarised in the following way:

Maximum [(Strike – Premium) – Share price **or** –Premium]

$$\max[(E - S_T) - \text{Premium}, -\text{Premium}].$$

If [(Strike – Premium) – Share price] > $[(E - S_T) - \text{Premium}] > -\text{Premium}$ then an option holder, acting rationally, will exercise the option and the payout will be:

$$[(E - S_T) - \text{Premium}].$$

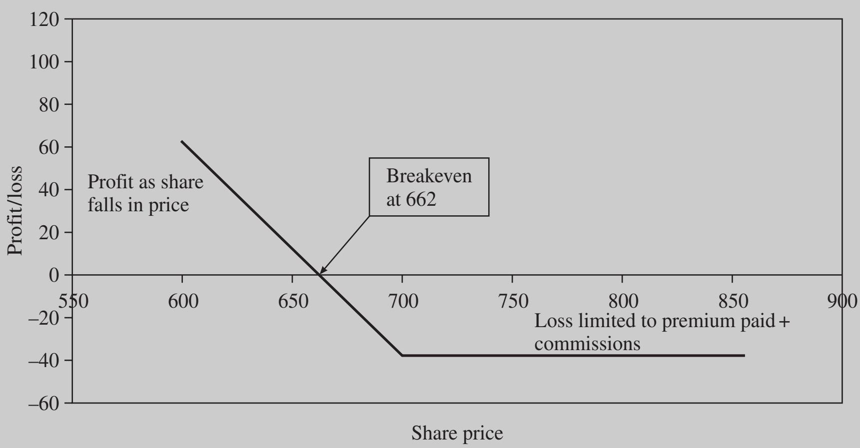

Figure 10.4: Payout profile: long put option

[2] This formula is extremely useful and, with small modifications, can be applied to determine the payouts generated by more exotic-style options.

If $[(\text{Strike} + \text{Premium}) - \text{Share price}] < [(E - S_T) - \text{Premium}] < -\text{Premium}$ (i.e. more negative) then, acting in a rational manner, the holder should not exercise the option in which case the cost involved in this entire transaction will be the premium.

If $[(\text{Strike} + \text{Premium}) - \text{Share price}] = [(E - S_T) - \text{Premium}] = -\text{Premium}$ then the holder will be indifferent. Whether exercise takes place or not the cost of the position to the holder will be the initial premium outlay.

Special labels are attached to call and put options as indicators of an option's position in relation to the underlying security. Options which have a positive large intrinsic value are said to be deep in-the-money. Options with zero (or very small positive or negative) intrinsic values are referred to as at-the-money (ATM) options. Those with high negative intrinsic values are referred to as out-of-the-money (OTM) options. This idea of in, at, and out-of-the-money options is demonstrated using call option expiration profiles in Figure 10.5(a) and (b).

The process of calculating intrinsic values for puts can easily be extended to cover all strikes in a similar manner to that already described. Calculation of a put option's time value is illustrated in Table 10.8(a).

Having identified the intrinsic value (0p), the time value can now be identified. To achieve this separation the intrinsic value is subtracted from each of the quoted option premiums. The arithmetic in Table 10.8(b) demonstrates that, although no intrinsic value exists, each put option has a time value: 38p, 60p and 76p for December, March and June contracts, respectively.

In this example for a strike of 750 there is an intrinsic value of 33p. The current price of the underlying security is below the strike price and, therefore, exercise would be desirable. The associated time value for each put option maturity is found to be: 30p, 54p and 70p for the December, March and June put option contracts, respectively.

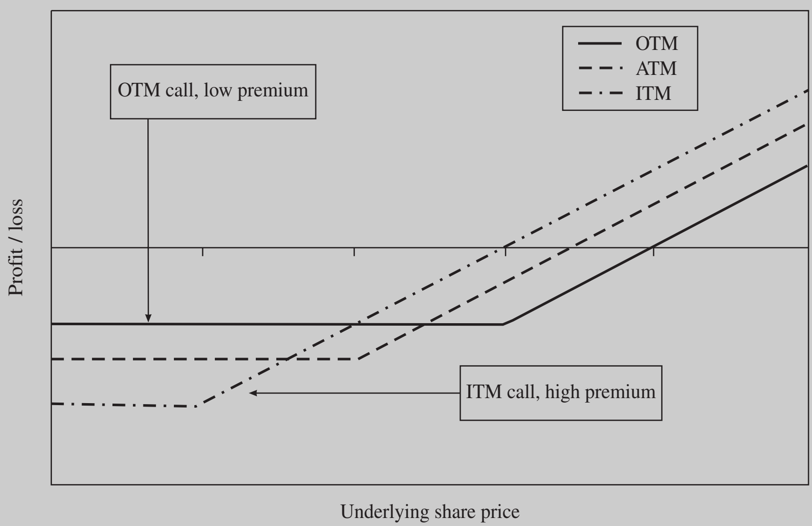

Figure 10.5(a): I, A and OTM call option profiles

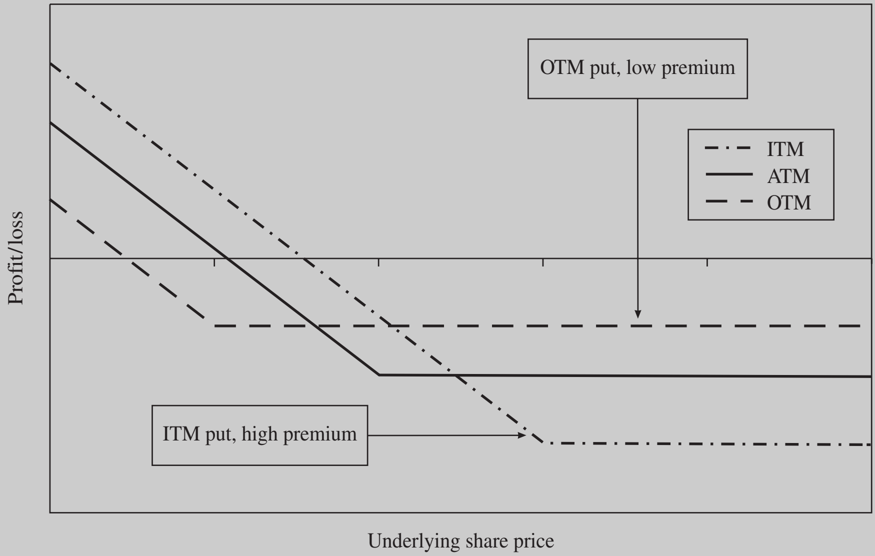

Figure 10.5(b): I, A and OTM put option profiles

	Puts		
	DEC	**MAR**	**JUN**
Strike price	700	700	700
Current market price	717	717	717
Intrinsic value	0	0	0

Table 10.8(a): Lloyds TSB Group plc, intrinsic values of 700 put options

	Puts		
	DEC	**MAR**	**JUN**
Option premia	38	60	76
Intrinsic value	0	0	0
Time value	38	60	76

Table 10.8(b): Lloyds TSB Group plc, time values of 700 put options

	Puts		
	DEC	**MAR**	**JUN**
Strike price	750	750	750
Current market price	717	717	717
Intrinsic value	33	33	33

Table 10.9(a): Lloyds TSB Group plc, intrinsic values of 750 put options

	Puts		
	DEC	**MAR**	**JUN**
Option premia	63.0	87	103
Intrinsic value	33.0	33	33
Time value	30.0	54	70

Table 10.9(b): Lloyds TSB Group plc, time values of 750 put options

All other strike prices could be considered in a similar manner.

Over the life-span of an option the underlying security's price may move up or down and in consequence may generate a positive intrinsic value for options which had previously been out-of-the-money. This possibility, however small, must be paid for. From Tables 10.6(b) through 10.9(b) it should be clear that the closer an option is to expiration the smaller will be its time value.

10.2.1 *Short calls and short puts*

Until now the discussion of options has centred on the holder's position, i.e. a long position. In order for a long position to be taken a short position has to be created; in the case of an equity call option some financial institution, company or individual must agree to the terms of the contract, namely, to deliver 1000 shares at a contracted price on or before a specified date in the future. The seller of the option is known as the option writer and, not surprisingly, the opening up of a short position is known as writing an option. There are two types of written options needed to cover the bought option positions discussed above: they are written calls and written puts. Unlike the holder of an option the writer must meet an obligation to sell or buy shares at the agreed strike price as soon as the holder decides to exercise.

The expiration payoff profiles faced by option writers are the mirror image of the profiles faced by option holders. The formula which provides the payout faced by the writer is given by:

$$(-1) \times \max[S_T - E] \text{ for a call}$$

and

$$(-1) \times \max[E - S_T] \text{ for a put.}$$

Notice that this is identical to the payout to the option holder in each case multiplied by -1.

Illustrations of short calls and short puts can be seen in Figures 10.6 and 10.7, respectively.

10.2.2 *Option strategies*

It is the kinked payout profiles that set options apart from other instruments and make them so flexible and so attractive to a wide variety of market participants. Strategies can be devised using options alone or combined with some underlying security to construct expiration payout profiles that meet end user requirements. As a concrete example use will

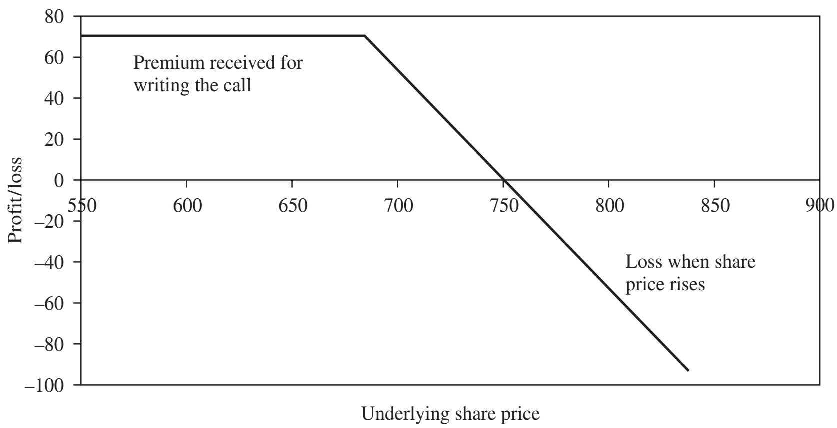

Figure 10.6: Call writers' profile

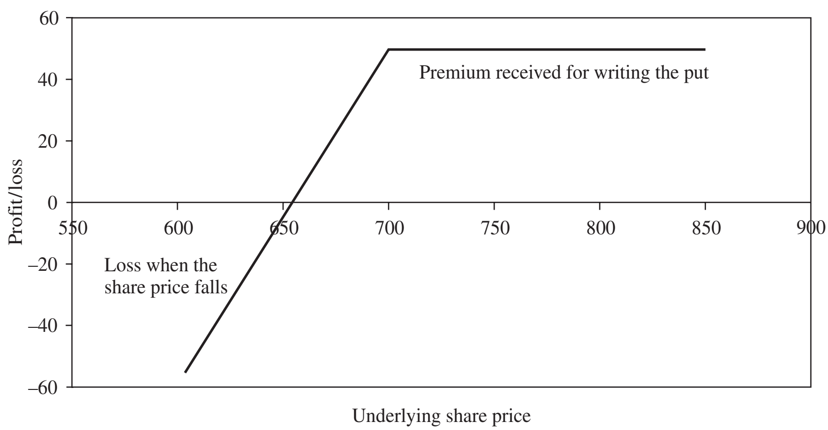

Figure 10.7: Put writers' profile

again be made of the initial call and put prices quoted on the Bloomberg screen, hypothetical option prices will be used to explain the result of the adopted strategy.

Example 10.5: Put option alone

A fund manager may decide to enter the options arena for several reasons: the price of the share is expected to rise or fall substantially. In this case a fund manager who is not holding the share in a portfolio would probably be looking to enhance the returns on the portfolio by either buying calls if it is anticipated that shares will rise in price or puts if the feeling is that the share price is set to tumble. If the share price rises (or falls) in accordance with expectations profit will be generated and realised either through an increase in the relevant option's price and selling the option on the market, or in the case of an American-style option, by exercising the right to buy or sell the underlying shares at contracted strike price.

From the Bloomberg option summary pages in Figure 10.2 the current market price of the share is 717p. Consider the case where a Dec put option with a strike of 650 is purchased today for a premium of 20. Two days later the government releases inflation statistics, which are regarded as poor by the markets. As a consequence the share price falls to 680p. The put option still has time remaining to expiration and its bid premium is now, say, 40. The option can be sold to close the position and a gain of 20 will have been made. Remember that the contract is in respect of 1000 shares so the fund manager's gain for each contract purchased will be £200 ignoring transaction costs.

Note that at this point it would not be sensible for the fund manager to exercise the option and sell the shares at the strike price of 650. Such a strategy would result in a 30 pence per share loss (650 [Strike] – 680 [Current share price]) with, of course, transaction costs on top.

Example 10.6: Call option – prepositioning

A fund manager intends to buy a particular share at some point in the near future in order to add it to a client's portfolio but currently the fund manager lacks sufficient funds to purchase the share outright. Moreover, there is an expectation in the market that the price of the share is about to move up significantly. Call options will allow the fund manager to lock into a favourable price now, and will, in a sense, provide insurance against the expected price rise.

To illustrate this idea consider the shares of Abbey National plc, which are currently trading at 1095p. To take advantage of an expected rise in the price the fund manager buys call options with a strike of 1100p. The option premium is 95p per share, which translates to a premium of £950. A week later the fund manager closes the option position when the Abbey National plc share price reaches 1150p and the option bid price is 117 – a profit of £220 per contract.

Here the gain in terms of the underlying security will be 50 (1150 [Current share price] – 1100 [Strike]). In other words for each call option contract bought the fund manager will make a profit of £500 but remember that the call option cost £950 to buy so adopting this strategy will actually result in a loss of £450 per contract.

Example 10.7: Put option – portfolio insurance

A fund manager may suspect that a share that he/she holds in a portfolio is about to fall in price and wishes to insure the current value of the stock. By purchasing put options with a strike close to the current market price of the share, the fund manager is able to lock into a figure close to the value required.

As an example of this situation consider the Abbey National plc shares. The market price is currently at 1095p and the fund manager wishes to insure against a fall from this level. Put options can be used to achieve this by buying an appropriate number of contracts with a strike price of 1100p. For example, if 10,000 Abbey National plc shares are held in the portfolio 10 put contracts would be required to hedge the position. The

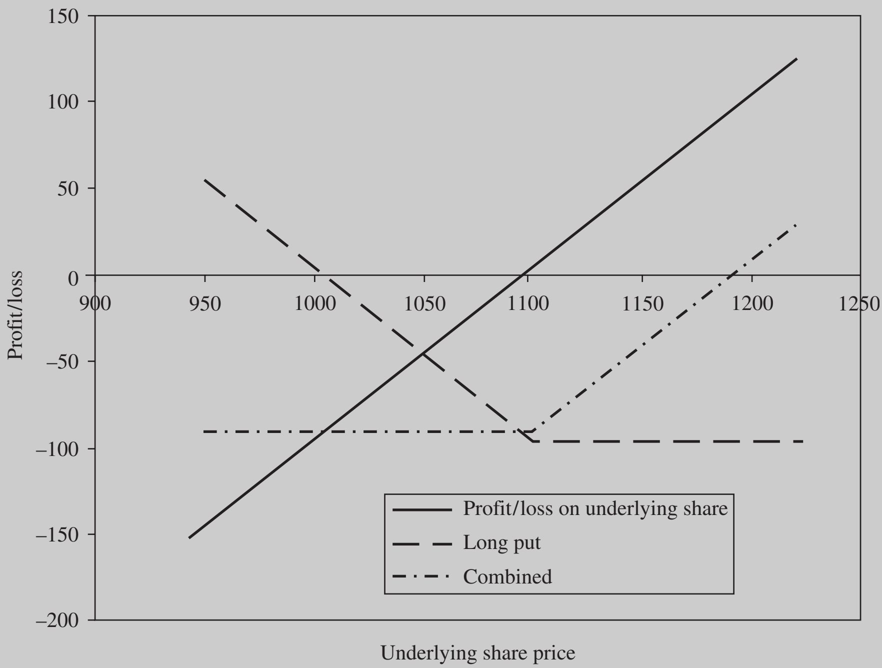

Figure 10.8: Portfolio insurance

puts can be purchased for 96, that is to say a premium of £960. Figure 10.8 illustrates the idea of portfolio insurance – notice in particular the fund manager's participation in a rise in the value of the Abbey National plc shares and limited downside risk.

If the assumption is made that the share price falls to 1000p one week after the options have been purchased and the price of a put contract rises to 152 a profit of 56 or £560 per contract can be made by closing the put position. If the put is exercised 100 per contract is made but the cost of buying the puts needs to be taken into account and this reduces the profit to £4 per contract.

Example 10.8: Arbitrage – long call plus short put versus short universal stock futures (USF)

To explore the idea of arbitrage consider the following: An investor buys a Dec 700 call @ 54 and writes a Dec 700 put @ 33. The net outlay is 21. Simultaneously the investor shorts a Dec Lloyds TSB Group plc USF @ 718.5. These two strategies are illustrated and annotated in Figure 10.9.

To analyse what happens to the position consider the payout at expiration when the underlying security closes at 750.

The short USF is bought back @ 750 for a loss of 31.5. The long call has a payout of 50 but cost 54 to buy when the strategy was entered into. There is therefore a net loss

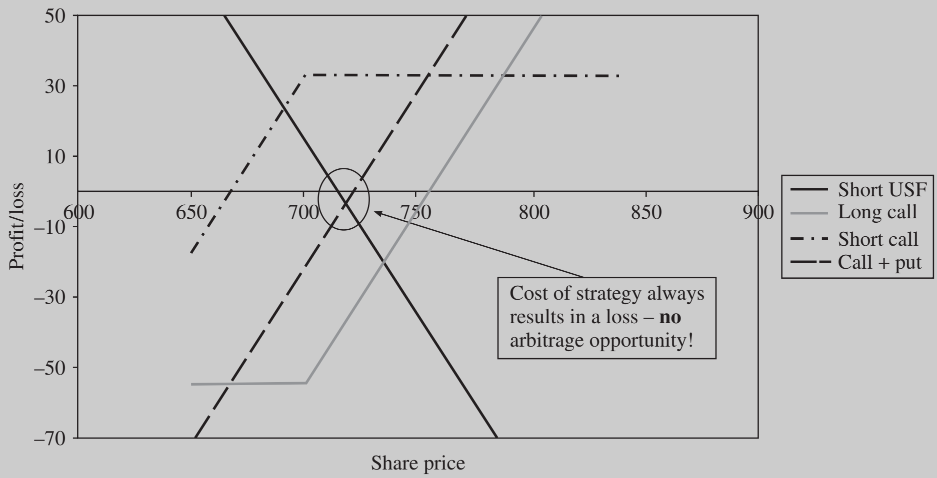

Figure 10.9: Arbitrage strategy

of 4 on the call leg of the strategy. The put leg of the strategy brings in a gain of 33, since the put has expired worthless to the holder. The net gain to the investor by adopting this option strategy is 29 (33 – 4). These activities and cash flows are presented in Table 10.10. The overall outcome of this strategy, taking the USF and option strategies together, is a loss of 2.5 per contract.

What if the underlying security closes at 650 and not 750?

Under this scenario the short USF is bought back @ 650 and will make a profit of 64.5. The long call will expire worthless but cost 54 to buy when the strategy was entered into. The put will finish in-the-money and with a strike of 700 will generate a gain of 50 to the holder. However, when the put was written it brought in a premium of 33. Combining these sums yields a net payout of –6.5 (i.e. 64.5 – 54 – 50 + 33). Once again a loss of 2.5 per contract is experienced by the investor and no arbitrage opportunity exists.

		Short USF		**Short put**		**Long call**
Now	Sell @	+718.5	Sell @	+33	Buy @	–54
Expiration	Buy @	–750.0	Value at T	0	Value at T	+50
	Result	–31.5		+33		–4
Outcome	–2.5					

Table 10.10

		Short USF		**Short put**		**Long call**
Now	Sell @	+718.5	Sell @	+33	Buy @	–54
Expiration	Buy @	–650.0	Value at T	–50	Value at T	0
	Result	68.5		–17		–54
Outcome	–2.5					

Table 10.11

The interested reader is left to verify for his or herself that the same will be true for a *long put plus short call versus long USF strategy*.

As demonstrated above the potential payoff profiles that can be constructed to reflect the scenarios listed depend upon whether the options are used by themselves – without the backing of an underlying security (a naked or uncovered position) – or in conjunction with an underlying security (a covered position). Other popular strategies will now be developed using an options-only framework.

This section concentrates on the use of options as stand-alone instruments. The stand-alone idea is built upon in order to demonstrate how the kinked shape of options' payoff profiles can be used to structure special positions.

The four payout graphs (long call, long put, short call, short put) introduced above can be regarded as a basic set of building blocks that facilitate the engineering of positions designed to exploit expectations about price behaviour of an instrument. The instrument in question might be an equity index of some construction or an individual stock. It should be emphasised that one of the big differences between derivative instruments such as futures and forward contracts is that for the holder of an option the downside risk is limited. Should the market move in the opposite direction to that expected, the maximum loss involved in a transaction will be the initial contract premium that was paid to engineer the position, multiplied by the number of contracts bought.

An important point here is that the holder of an option – a call or a put – must have held a view about where the market was going when the transaction was entered into. However naive, some method of forecasting is likely to have been used to gain that view. In the futures and forwards arena this might have been the case but any forecast is, to an extent, irrelevant because the price at some future time is fixed at the point in time when the contract is formally agreed. After that point whatever happens to the market price – whether it rises or falls – the future price is assured. Any gain (loss) in the spot market will be offset by a loss (gain) in the futures or forward market, ensuring absolute price stability. Options, on the other hand, protect against anticipated downside risk but allow participation in the market at favourable prices, if the anticipated downside risk disappears.

The world's leading exchanges recognise a number of option strategies which combine calls and puts into popular and useful payout profiles. The exchange-recognised strategies are available across several derivative markets and cover bonds, short-term interest rate products (STIRs), equity index options and individual equity options. Moreover if the separate legs of a strategy are booked to a single account they qualify for reduced transaction fees. It must be stated, however, that not all strategies enjoy recognised status across all market arenas. For example, at the time of writing although a synthetic futures position can be created using calls and puts on individual stock options this is not a LIFFE-recognised strategy.

To gain a more thorough understanding of some of the basic strategies that can be constructed use will be made of the spreadsheet **Strategies and Time** on the accompanying CD. This spreadsheet allows the user to combine calls and puts across different strike prices and different volatility inputs. By defining the number of days before expiration of the contract(s), the spreadsheet allows comparison of the payout profiles a number of days prior to expiration with the familiar expiration payout profile of the strategy.

Stock index option pricer **Input the following values:**		**Volatility p.a. (%)**
Days to expiration	90.00	
Strike Price CALL High	750	41.51
Strike Price CALL Middle	700	41.51
Strike Price CALL Low	650	41.51
Strike Price PUT High	750	41.51
Strike Price PUT Middle	700	41.51
Strike Price PUT Low	650	41.51
Current underlying price	717	
Rate of interest (%) p.a.	4.16	
Approx. dividend yield (%) p.a.	0.00	

Table 10.12: Stock index option prices

Strategy 1: Straddle

Assume that a portfolio manager does not hold shares in Lloyds TSB Group plc in a portfolio under his or her control. He or she feels – perhaps on the basis of a controversial press report – that the price of that share is about to move dramatically but is not sure in which direction. To make the most of this situation the portfolio manager decides to buy one call and one put option with the same strike price and the same expiration date. This strategy is known as a straddle and is illustrated in Figure 10.10. In this example an ATM call and an ATM put have been combined.

Notice that the value of the strategy 90 days from expiration is quite high, just dipping slightly into the loss quadrant of the graph. This is because both of the options have considerable time value remaining. Only when the options reach expiration in 90 days' time and the share price stays close to the strike price will the maximum loss be experienced.

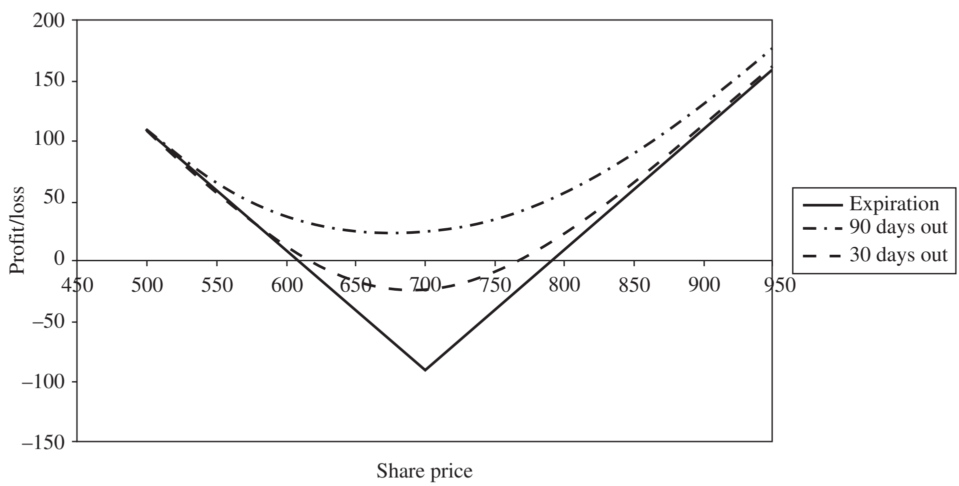

Figure 10.10: Option strategies – long straddle

Case 1		
Today (volatility 41.51%)	Buy call and put for	118
Next day (volatility 50%)	Sell call and put for	140
	Profit	22
Case 2		
Today (volatility 41.51%)	Buy call and put for	118
Next day (volatility 60%)	Sell call and put for	168
	Profit	50

Table 10.13

To demonstrate the attractiveness of this strategy consider what would happen if the volatility increases the day after the strategy has been purchased. Using theoretically calculated mid-prices[3] assume that the call was purchased at a premium of 71 and the put at a premium of 47, giving a total of 118. The next day the underlying security remains at 717 but volatility rises to 50% p.a. On this basis the call can be sold for 82 and the put for 58. This will yield a profit of 22 (£220) on the strategy. Should the volatility rise to 60% the call and put premia will rise to 96 and 72, respectively. This scenario will result in a 50 (£500) profit for the strategy. Table 10.13 summarises this example.

Strategy 2: Long call spread

Assume in this example that a portfolio manager feels that the price of a share is likely to rise and wishes to profit from this view. The fund manager feels, however, that the premium for a call is too high. By sacrificing some of the upside potential in the share price movement he/she can write a call option and use the premium received to offset the premium that must be paid to purchase the long call.

The strategy then will be to purchase a call with a "low" strike price and write a call with a higher strike price. Once again using the Lloyds TSB Group plc call option quotes the strategy might be: buy a Dec 01 700 call @ 54 write a Dec 01 750 call @ 25. The net cost of establishing this position is 29. Figure 10.11 presents the strategy diagrammatically.

The figure illustrates quite clearly that once the share price has reached the strike price of the written call any further gain resulting from the rise in price of the share will be passed on to the holder in possession of the call with the higher exercise price. On the other hand because no underlying security is involved in this transaction the maximum loss faced by the strategy holder will be the premium paid for the long call.

It is interesting to note at this point that the payoff profile exhibited in the long call spread strategy can be replicated in several ways, for example by combining a long underlying position with a long call with an OTM strike price and a long put with an OTM strike or by using puts – selling an ITM put and buying an OTM put. It is also worth pointing out that this profile is one frequently encountered in other market arenas where names such as *cylinder*, *range forward* and *collar* are used to describe it.

---

3 The model used to calculate these prices is the Black and Scholes model about which more will be written in the next chapter.

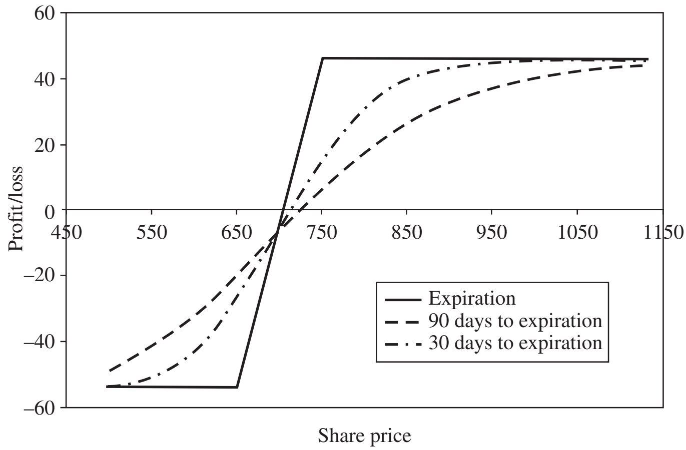

Figure 10.11: Option strategies – long call spread

Strategy 3: Long combo

This strategy is very similar in its pre-expiration payout profile to that exhibited by a short futures position. The fundamental rationale for its use is that the instrument in question is expected to decline in value. If this expectation materialises then the gains are potentially unlimited. Note, however, the plateau that is formed at expiration where – in this example – a positive payout is generated when the share price finishes between A and A′ in Figure 10.12. To some extent this plateau provides a cushion in the event that the price of the instrument does not fall as expected and/or an increase in volatility occurs.

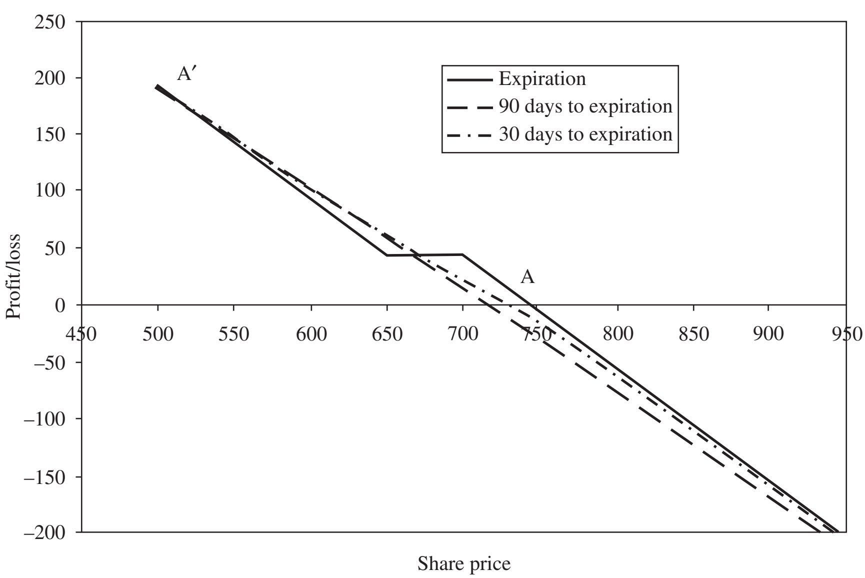

Figure 10.12: Option strategies – long combo

Drawing on the Lloyds TSB Group plc data the position can be created by shorting a Dec 700 call @ 49 and buying a Dec 650 put @ 20. The income from entering this strategy is 29. The risk faced by the investor is a rise in the price of the underlying security beyond 729p.

Strategy 4: Two by one ratio call spread

A glance at the time value associated with this strategy and illustrated in Figure 10.13 explains the appeal that this strategy holds for a wide range of investors and fund managers. Consider the payout profile 90 days from expiration – the fact that an ITM call option has been written and has taken in premium enables two OTM calls to be purchased at very little expense. It is only because of the passing of time and an underlying share price that remains static, around its value when the strategy was initiated, that the spread returns its largest loss.

Strategy 5: Long calendar spread

Strategies 1 to 4 are often referred to as vertical strategies. This terminology simply indicates that each option in the strategy will reach expiration at the same moment in time. The figures depicting the strategies all show what the strategy will look like before the options used to create the strategy expire, i.e. while they have some time value remaining. Figure 10.14 illustrates the behaviour of time value for a call option over a period of one year. Notice that as the option approaches its terminal date the option's time value begins to decrease at an increasing rate. It is this time value behaviour that an investor using a calendar spread seeks to exploit.

To analyse how such a strategy might function consider the following example. A fund manager wishes to purchase a put option to protect against the fall in the price of a share in his/her portfolio. The current share price is 622p and there are just 30 days to the expiration

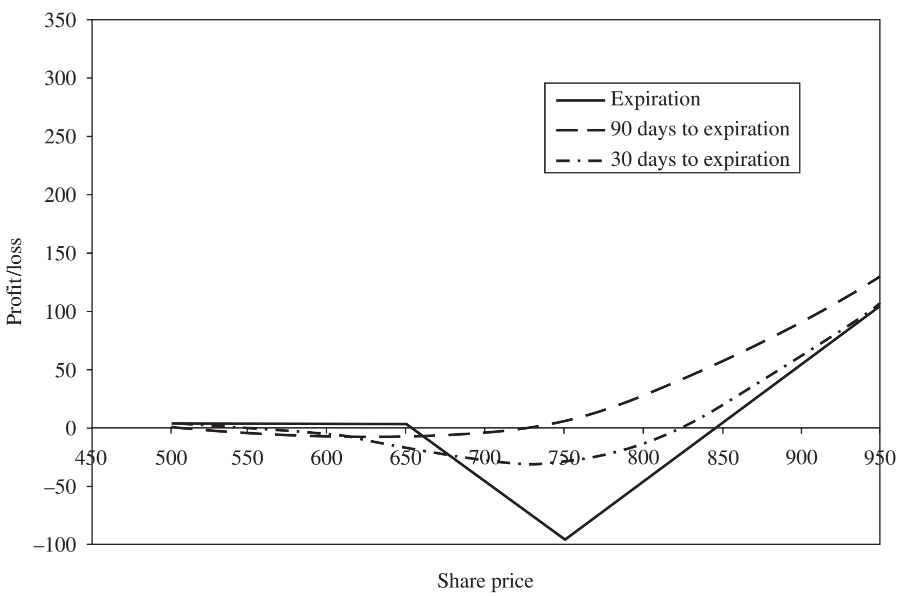

Figure 10.13: Option strategies – long 2 × 1 ratio call spread

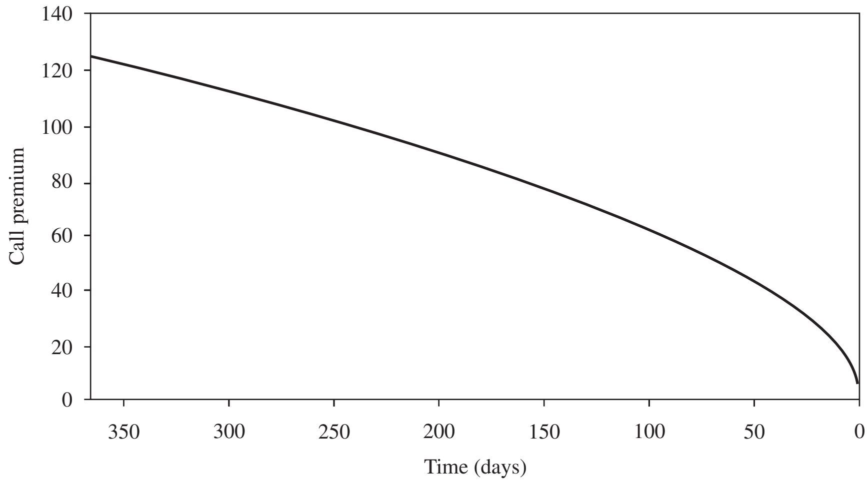

Figure 10.14: Option time value

of the nearby contract. If the nearby Jan 600 put contract can be sold for 12 and the Apr 600 contract can be purchased for 29.5 the net outlay will be 17.5. Since the time value of the Jan contract will decay faster than the time value of the Apr contract the option value of the nearby contract will fall faster than the option value of the Apr contract.

Under one illustrative scenario let the share price fall from 622p to 600p by expiration of the Jan contract. This means that the Jan contract will expire worthless. From the analysis of the components of an option's premium discussed earlier it should be clear that the Apr contract also has zero intrinsic value. The Apr contract, though, will have a substantial amount of time value remaining and it is this time value that the investor has purchased at a cheaper rate by initially writing the short dated contract.

Strategy 10: Long volatility trade

This strategy seeks to exploit the expectation that the volatility of the underlying instrument is about to rise dramatically. This is a LIFFE recognised strategy for equity options and involves buying calls and selling a futures contract or buying puts and buying a futures contract with the same time to maturity and in such quantities that a combined value of zero for delta is achieved and maintained dynamically. Its appearance is similar to the straddle discussed in Strategy 1.

Figure 10.15 illustrates what happens to the value of this position at expiration. With no passage of time and no change in the price of the underlying security the increase in volatility will lead to an increase in the value of the position. So once again this is a strategy to be employed when expectations are rife that something is going to happen to the value of the underlying security but nobody is sure what that move will be.

Strategies – general comments

Clearly there are many strategies to choose from and a natural question would be which strategy is best? Figure 10.16 presents some popular strategies organised according to anticipations about changes in volatility and price of the underlying security.

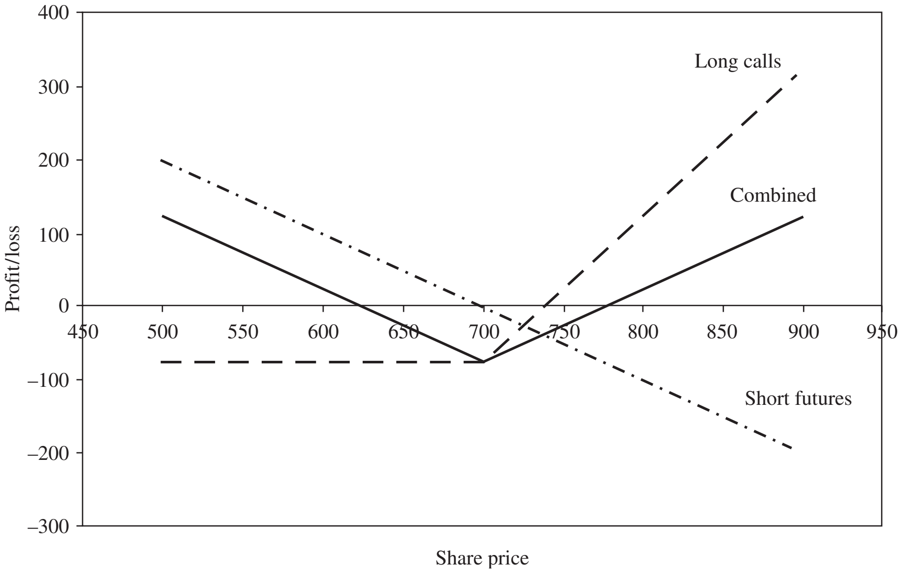

Figure 10.15: Long calls plus short futures position. Note that long calls and short futures are combined to give a position delta of zero and that the strategy needs to be managed dynamically to ensure that a delta neutral position is maintained over time

The grid is based on volatilities across values of 10% to 40% and an underlying price from 660p to 740p. Throughout it has been assumed that each strategy has been initiated at the centre point (labelled Box) where volatility is 25%, the underlying security is 700p, the financing rate is 4.16%, and the time to option expiration is 90 days. To assess the performance of each strategy, it has been assumed that 10 days have elapsed since its implementation. In all cases Black and Scholes premiums have been used with no bid–ask spread and no transaction costs and under the assumption that interest rates and dividend yield remain constant.

Scrutiny of the positive quadrant of Figure 10.16 shows that some strategies appear to outperform others. For example, when volatility rises but there is little or no movement in the price of the underlying security, in terms of absolute payout a long OTM call (750 strike) combined with a long OTM put (650 strike) creating a strangle will outperform a long ATM straddle (700 strike). Likewise when both volatility and the underlying share price are expected to rise a long OTM call would outperform a long ATM call.

One problem with this, however, is that the market does not often behave exactly as an investor – small or large – expects. Tables 10.14 and 10.15 compare the relative performance of a variety of option strategies under scenarios of changing volatility and changing price of the underlying security. In Table 10.14 percentage performance is determined by

$$\left[\left(\frac{\text{Payout}}{\text{Outlay}}\right) - 1\right] \times 100$$

which permits a comparison of the percentage returns from the ATM straddle and the OTM strangle.

It is clear from Table 10.14 that at levels of volatility below 25% the straddle will lose less than the strangle at all share price levels. However, once volatility rises above 25% the

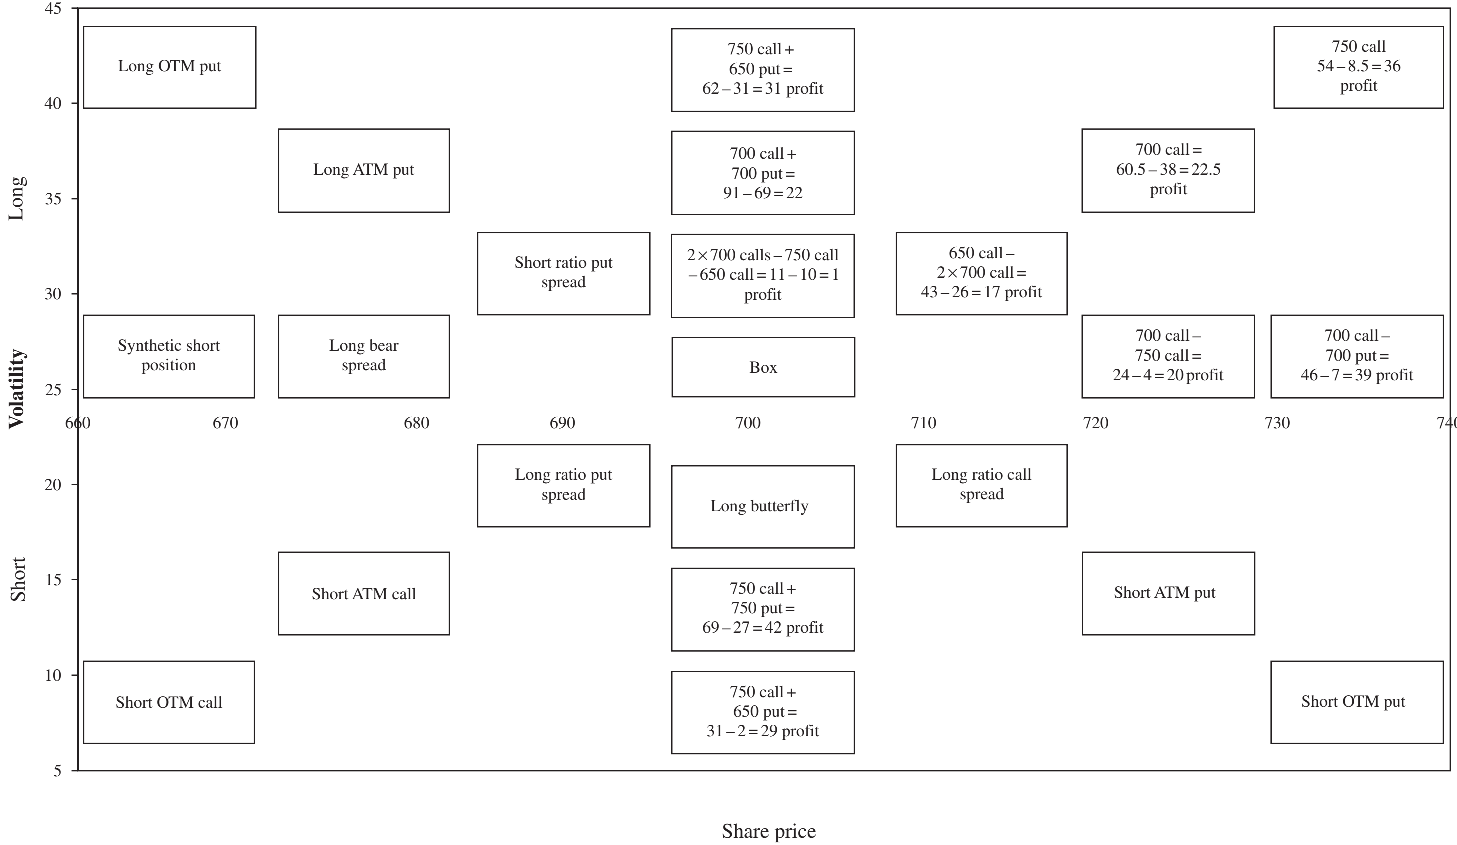

Figure 10.16: Option strategy compass

Strangle	Purchased @		31		
Volatility/share price	660	680	700	720	740
10	−78%	−89%	−90%	−80%	−57%
15	−57%	−70%	−70%	−58%	−33%
20	−31%	−42%	−42%	−29%	−6%
25	−1%	−10%	−8%	3%	25%
30	32%	25%	27%	39%	59%
35	67%	62%	65%	76%	95%
40	103%	99%	103%	114%	132%
Straddle	**Purchased @**		**69**		
Volatility/share price	**660**	**680**	**700**	**720**	**740**
10	−43%	−58%	−60%	−49%	−28%
15	−30%	−41%	−42%	−33%	−17%
20	−15%	−23%	−23%	−16%	−3%
25	1%	−4%	−4%	2%	13%
30	18%	14%	15%	20%	30%
35	35%	32%	34%	39%	48%
40	53%	51%	52%	57%	66%

Table 10.14

strangle delivers a higher positive payout. Only for a 25% constant volatility is the payout result mixed. Table 10.15 illustrates the payout regimes under the selected share price levels and volatilities after 10 days have elapsed. The zeros represent states where the straddle loses less of the original investment than the strangle. The ones represent states where the strangle delivers a better relative positive payout than the straddle.

Almost all of the examples developed above make the assumption that volatility is constant across all exercise prices. However, inspection of the data contained in the Bloomberg option tables reveals that this is not the case. The option summary pages for any of the companies discussed in this chapter report different volatilities for calls and puts, for bid and ask quotes, and for different expiration dates. For example, in Figure 10.2, the Lloyds Group plc Dec 01 put with a strike of 700 has a volatility of 39.21% associated with the bid price of 33 and 43.90% associated with the ask price of 38. These volatilities go by a special name: they are called *implied volatilities*. They are in fact the volatilities that have been used to arrive at the market quoted prices for that option. Since volatility is non-constant in the real world this will have implications, and indeed serious consequences, for

Volatility/share price	660	680	700	720	740
10	0	0	0	0	0
15	0	0	0	0	0
20	0	0	0	0	0
25	1	0	0	1	1
30	1	1	1	1	1
35	1	1	1	1	1
40	1	1	1	1	1

Table 10.15: Comparison of strangle and straddle payouts

all aspects of option usage: risk hedging, constructing synthetic instruments, arbitrage, and trading strategies.

The effect of volatility on option premia highlights, perhaps, their most important disadvantage for many end users. Like most investment decisions timing is crucial. In some cases markets react sluggishly to news, good or bad, giving a shrewd, or simply lucky, investor the opportunity to get in or out of a position before the true impact of that news is felt. More often than not though markets will have anticipated events and the market prices prevailing for affected instruments will, at the moment the news is released, already reflect the fact. Given the high gearing of options this can be most unfortunate for the hedger since premiums will rise very rapidly.

Take the following as a case in point: during the breakfast briefing session a fund manager is advised by his/her economists that the latest round of international arms talks is likely to conclude in an agreement that expenditure on armaments be reduced worldwide. It is felt that subsequent cuts will be large and that shares in several arms-related companies, already trading nervously, will fall dramatically. One share which the fund manager holds in the portfolio, and about which he/she is concerned, is currently priced at 544p. At a briefing session the fund manager's fundamental and technical analysts suggest an initial price support level of 520p but fear that once that level is breached a drop to as low as 425p is quite probable.

Immediately following that briefing session the fund manager checks the put option price for the contract with the closest expiry date. The premium is 1.5p, calculated on the basis of an interest rate of 5%, 25 days to expiration of the contract, a strike of 525, and annualised volatility of 15%. Later in the morning the share price has fallen slightly to 538p but the option premium has now risen to 7p. Two things have worked against the fund manager here: the option has moved closer to an at-the-money position and option traders, having anticipated the outcome of the talks, have raised volatility from 15% to 25%. If the bid–offer spread is also taken into account the position will be even worse. Initially at a premium of 1.5p the spread might have been 0.5–2.5p, later in the day the bid–offer spread might look more like 5–9p. Taking a negative view and assuming that the decision to buy a put option is delayed, one contract, at a later point in the day, may cost £90 as opposed to £15 at breakfast time. Of course, from a more positive point of view, a large profit would have accrued very quickly had the option been bought immediately at breakfast time!

EXCEL demo: Strategies and time[1]

Exercise prices and volatilities

This spreadsheet requires several inputs. In total three strikes for calls (cells B6, B8, B10) and three strikes for puts (cells B12, B14, B16) can be entered with a volatility associated with each of the strike prices (C6 for the strike entered in cell B6, C8 for the strike entered in cell B8, etc.).

Number and type of contracts used

The user needs to input the number of each type of option position taken. For example, long positions in calls are inputted in cells F2, F4, F6, while short call positions are entered in cells G2, G4, G6. For long put option positions the entries are required in cells F8, F10, F12 and G8, G10, G12 for short positions.

[1] These spreadsheets are in the accompanying CD-ROM.

Other inputs
The current price of the underlying security is required in cell B18, while the appropriate rate of interest and dividend yield are required in cells B19 and B20, respectively.

Graph scale
To adjust the scale for a graph of the desired strategy the start value for computation is needed in cell B24.

Selected bibliography and references

Cox, J., Rubinstein, M., *Options Markets*, Prentice Hall 1985
Eales, B. A., *Financial Engineering*, Palgrave 2000
Hull, J., *Options, Futures and Other Derivatives*, 4th edition, Prentice Hall 2000
Keating, C., Shadwick, B. A., Shadwick, W. F., "Unfocused Model Thinking", *Futures and Options World* 358, March 2001, pp. 76–83
Taleb, N., *Dynamic Hedging: Managing Vanilla and Exotic Options*, Wiley 1997

11 Option Pricing

11.1 Introduction

In several of the preceding chapters reference has been made to the use of options in creating engineered payoff profiles. In each of the examples used option premia were taken as given by the market. This chapter looks at how those prices, at least to an extent, can be calculated. In addition it considers and explains several different frameworks that are available for the calculation of option premia.

In Chapter 6 reference was made to the well-known and oft-cited Black and Scholes (B&S) option pricing model. The current chapter will explain how the B&S model can be applied and will develop alternative pricing approaches: lattice frameworks and Monte Carlo simulation. The overall objective in introducing these alternative pricing frameworks is to develop a good intuitive understanding of their strengths and weaknesses, whilst at the same time developing an operational framework that will enable the reader to generate payoff profiles, under a variety of market outcomes for a variety of instruments. The framework will also provide a springboard from which more complex models can be analysed and from which hedge parameters can be obtained.

The B&S model is introduced in Section 11.2 and will be examined in the context of pricing regular stock or equity options. The binomial and trinomial model will be used to examine the pricing of regular stock options, too, but will also demonstrate how the flexibility offered by these frameworks might lend itself to pricing options with non-standard payoff profiles. The Monte Carlo method will be introduced to demonstrate pricing, once again, under a non-standard option payoff scenario.

11.2 The Black and Scholes model

The option pricing model developed by Fischer Black and Myron Scholes first appeared in the academic literature in 1973. It now forms the basis for calculating the benchmark premia for regular and many forms of exotic options. The original model that Black and Scholes introduced to the world was based on some quite restrictive assumptions, subsequently researchers have been able to relax some of those underlying assumptions and develop their own less constrained option pricing models.[1] It is not the intention here to derive the B&S option pricing model from first principles and then to compare it algebraically with other models. It is, however, worth getting an intuitive appreciation of the B&S option pricing model, the inputs it requires, and the assumptions which underpin it.

[1] For a thorough coverage of recent options pricing models see: Hull, J., *Options, Futures, and Other Derivatives*, 4th edition, Prentice Hall, 2000.

11.2.1 *Inputs*

The inputs required to price an option in a basic B&S world are listed below:

1. The current price of the share;
2. The strike price;
3. The time remaining until expiration;
4. The risk-free rate of interest;
5. A measure of the standard deviation of the continuously compounded annual rate of return on the share – this input is crucial to the pricing model and is known as volatility.

Extensions of the B&S model require additional inputs such as dividend yield, foreign as well as domestic interest rates (currency options), resetting of strike prices, and barriers (exotic options). Some of these variations will be discussed and introduced in the sections that follow.

11.2.2 *Assumptions*

The assumptions on which the basic model is built are, *prima facie*, formidable:

1. The option is only exercisable at expiration;
2. The market operates continuously;
3. The share pays no dividend over the lifetime of the option;
4. The risk-free rate of interest is constant over the lifetime of the option;
5. There are no transaction costs, zero taxes and no bid–offer spread;
6. The underlying share can be shorted without penalty and short-sellers receive the cash benefits from the short sale in full;
7. Share prices are continuous and are not subject to precipitous changes in prices either up or down – shares are, in fact assumed to follow an *Ito process*.

11.2.3 *Stock price movements*

A fundamental premise of the B&S model is that equity prices move according to a Wiener process. Moreover when there are a series of many small random movements in the share price the track that it is tracing can be assumed to be geometric Brownian motion. This process can be defined as:

$$dS = \mu S dt + \sigma S dz \tag{11.1}$$

where

μ	is constant and represents the expected return on the share reported as an annualised rate;
σ	is constant and represents the share's volatility reported as an annualised rate;
dt	represents a minute passage of time;
dz	represents a term which generates randomness into the movement of the share price (S).

In this process notice that, if there is no randomness, the *dz* term can be set equal to zero and the change in a share's price from period to period can be expressed as:

$$dS = \mu S dt \text{ or}$$

$$\frac{dS}{S} = \mu dt. \tag{11.2}$$

Integrating equation (11.2) over 0 to *T*,

$$\int_0^T \frac{dS}{S} = \int_0^T \mu dt \tag{11.3}$$

the following is derived:

$$\ln S = C\mu T \tag{11.4}$$

where *C* is the constant of integration. Taking logarithms of both sides and setting *C* equal to the starting value of the share price, equation (11.4) can, in turn, be rewritten as:

$$S_T = S_0 e^{\mu T}. \tag{11.5}$$

This equation implies that knowledge of the expected rate of return together with a starting value for the share price will enable the investor to calculate the value of the share at some future time period *T*. The process involves continuous compounding as implied by the use of the exponential applied to the expected rate of return.

Of course in reality share price movements will not be deterministic; there will be fluctuations on a period-by-period basis, the magnitude and direction of which will be unpredictable. For this reason the term σSdz is included in equation (11.1). Breaking this term down into its component equations will help to show how uncertainty is captured by the model. On an intuitive level the parameter σ represents the volatility of the share price movements – more will be said about how this parameter can be measured in the sections which follow. The variable *S* represents the share's current market price, while the *dz* component drives the directional uncertainty and, in part, the magnitude of share price changes that occur. Technical requirements specify that the *dz* term possesses two properties:

1. Over an infinitely small passage of time (*dt*) the change in *z* (*dz*) will be random;
2. The changes in *z* that occur over any two infinitely small sections of time will be independent.

To turn equation (11.1) from a theoretical construct into a practical model the assumptions are adopted that the mean of *dz* is zero and the variance of *dz* is one. This allows the model to be reformulated as:

$$dS = \mu S dt + \sigma S \varepsilon \sqrt{dt} \tag{11.6}$$

where

ε is a random number drawn from the normal distribution with mean zero and variance (standard deviation) equal to one, $N(0,1)$.

Rearranging equation (11.6) yields:

$$\frac{dS}{S} = \mu dt + \sigma \varepsilon \sqrt{dt} \tag{11.7}$$

which, as demonstrated earlier, is the logarithmic return rather than the share price itself which is being modelled by using "known" values of σ and μ, and given a minute change in time. The ratio dS/S, the returns, will be normally distributed[2] with mean μdt and variance $\sigma^2\sqrt{dt}$.

The Black and Scholes pricing formula makes use of Ito's lemma. This result can be used to consider the process for a share price, S. If the changes that occur in the share's price can be modelled as equation (11.6), and μ and σ are functions of the share price and time, then:

$$dS = \mu(S,t)dt + \sigma(S,t)dz. \tag{11.8}$$

Obviously this list presents quite a mixture of inputs and a very demanding set of assumptions some of which will seldom, if ever, be satisfied.

Considering first the required inputs it is clear that over some variables the purchaser/seller of the option will have a great degree of control, for example, the strike price that is to be adopted. The strike could be OTM, ATM or ITM as the strategy to be constructed requires.

The purchaser/seller is also, within limits, able to decide for what period of time the option will be held. In addition to having the possibility of trading the option (buying or selling the option before expiration), he/she can decide which expiration date is to be purchased or sold.[3]

On the other hand there are also inputs required over which the option user can exercise little or no influence. The current price, which has to be paid for the security, will be determined in the marketplace by the law of supply and demand. It is an input whose value will normally vary throughout each and every day and over which the option purchaser or writer will have no direct control.

Associated with this is the fact that movements in a security's price determine returns. These returns will in turn be used to gain an estimate of the security's volatility. The option purchaser/writer can do little about the value chosen to represent volatility apart from, perhaps, in the choice of formula used to estimate the input's value. In the previous chapter mention was made of the concept of implied volatility and the implied volatilities reported in the Bloomberg tables illustrate clearly the problem of trying to identify one number to represent volatility in an option pricing model.

Figure 11.1(a) illustrates this point. The implied volatilities of call option bid prices quoted on the Bloomberg CAC 40 summary page, Figure 11.1(b), when plotted against their associated exercise price present a lop-sided smile. A similar plot of the ask implied volatilities depicts a slight grimace. The volatility calculated on the basis of historical observations provides one isolated point which lies outside the plots of current market implied volatilities.

---

[2] There has been a considerable amount of debate about the distribution of returns. A good summary of this is available as a collection of readings: *Return Distributions in Finance*, edited by John Knight and Stephen Satchell, Butterworth-Heinemann, 2001.

[3] In some cases blue chip equity options in the US can have much longer time periods to expiration. In 1991 the Chicago Board Options Exchange (CBOE) introduced Long-term Equity AnticiPation Securities (LEAPS) which are in fact American-style call and put options with expiration dates up to two years ahead.

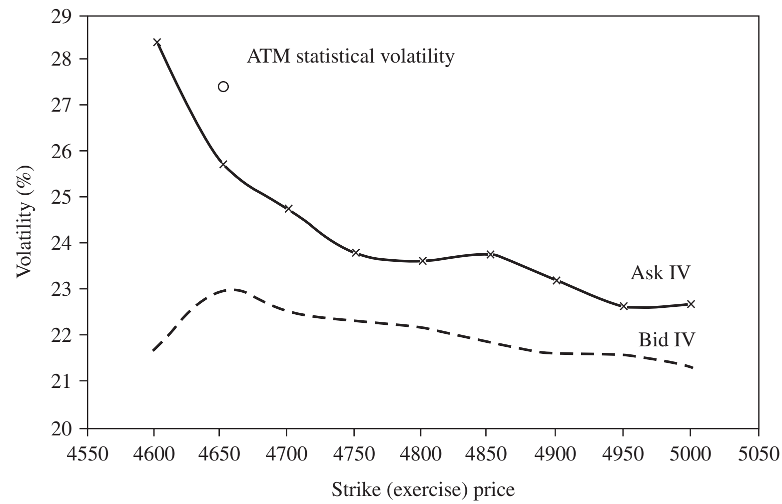

GRAB Index CALL

CALL SUMMARY PAGE Hit 99 <go> for option

CAC CAC 40 INDEX 13:23 Hedge:M model defaults

4666.7 -12.67 Div pt Days 21 Fin Rate 3.34 Volat: 27.46

	LAST TRADE INFO				IMPL.VOLAT		HEDGE			*THEORETICAL	
CAC Dec 01	Bid	Ask	Last	Chnge	Bid	Ask	Del	Gam	Vega	Price	Diff
1) 12 C4200			495.9	unch	b.int	b.int	951*		1.2*	481.23	
2) 12 C4250			448.7	unch	b.int	b.int	930*		1.6*	434.66	
3) 12 C4300			375.0	-25.6	b.int	b.int	903*	001*	2.0*	389.41	
4) 12 C4350			355.0	unch	b.int	b.int	870*	001*	2.4*	345.78	
5) 12 C4400			311.0	unch	b.int	b.int	830*	001*	2.9*	304.14	
6) 12 C4450			269.5	unch	b.int	b.int	783*	001*	3.3*	264.80	
7) 12 C4500			210.0	-20.1	b.int	b.int	730*	001*	3.7*	228.08	
8) 12 C4550			169.0	-23.5	b.int	b.int	672*	001*	4.1*	194.24	
9) 12 C4600	138.7	166.4	135.0	-22.9	21.66	28.14	619	001	4.3	163.47	+10.92
10) 12 C4650	116.0	128.5	112.0	-15.2	22.99	25.80	549	001	4.4	135.90	+13.65
11) 12 C4700	89.5	99.5	92.0	-7.4	22.53	24.77	475	001	4.5	111.55	+17.05
12) 12 C4750	68.0	74.5	67.0	-9.2	22.32	23.82	399	001	4.3	90.39	+19.14
13) 12 C4800	50.5	56.5	52.0	-4.6	22.17	23.65	327	001	4.1	72.28	+18.78
14) 12 C4850	36.0	43.0	41.3	unch	21.87	23.79	262	001	3.7	57.02	+17.52
15) 12 C4900	25.0	30.0	26.0	-3.3	21.65	23.24	200	001	3.2	44.38	+16.88
16) 12 C4950	17.2	20.0	17.4	-2.7	21.59	22.67	148	001	2.7	34.07	+15.47
17) 12 C5000	11.2	14.0	10.9	-2.0	21.38	22.72	108	001	2.2	25.80	+13.20
18) 12 C5100			5.1	-1.6	b.int	b.int	100*	001*	2.0*	14.20	

Australia 61 2 9777 8600 Brazil 5511 3048 4500 Europe 44 20 7330 7500 Germany 49 69 92041210
Hong Kong 852 2977 6000 Japan 81 3 3201 8900 Singapore 65 212 1000 U.S. 1 212 318 2000
I666-32-0 07-Dec-01 13:24:04

Figure 11.1(a) & (b): Implied volatility smile. © Bloomberg LP. Reproduced with permission.

Calculating volatility from past data

1. The basic standard deviation approach

This method of obtaining an estimate of volatility is based on obtaining the standard deviation from a set of returns (usually these returns are reported as natural logarithms (logs to base e) to comply with the notion of continuous compounding). As a demonstration of how this process works in practice take the data set on the CAC 40 over a period of 40 days depicted in Column 2 of Table 11.1. The third column reports the natural log returns calculated as:

$$R_t = \log_e \left[\frac{S_t}{S_{t-1}} \right] \tag{11.9}$$

where

R_t represents return at time period t;
S_{t-i} represents the security's price at time period $t - i$, $i = 0, 1$.

In this example the observations have been recorded daily so the returns are calculated on the basis of today's price divided by yesterday's price.

Column 4 reports the most recent annualised volatility calculated as 26.59%. The volatility value is based on the variance of returns calculated as:

$$\sigma^2 = \frac{\sum_{i=1}^{T} \left(R_t - \bar{R}\right)^2}{(T-1)} \tag{11.10}$$

where

σ^2 represents the daily variance expressed of returns;
R_t represents return at time period t;
T represents the number of observations used in the calculation of volatility and the square root of number of days in the year – here 252 trading days have been assumed.

The reported annualised volatility (*vol**) is calculated as equation (11.10a):

$$vol^* = 100\sqrt{\sigma^2 252} = 26.59\%. \tag{11.10a}$$

From Table 11.1 it can be seen that over the period under consideration a rolling window of volatilities has been established. The first recorded volatility based on closing prices from the previous 30 days is estimated to be 26.59%. The next calculation, reflecting the previous day's volatility, again based on 30 observations but now starting from the return on Day 3, yields a volatility of 27.26%. This method is simple to work with and is readily understood. It is, however, affected by factors such as the number of observations used to generate the standard deviations. If, for example, 60 closing prices are used to estimate the annualised volatilities the figure would rise to approximately 30.0%. This can be accounted for, in this case, by the sad, and devastating events of September 11th and the unsettling effect that they had on financial markets worldwide. Data sets which only cover a short period of time will fail to pick up influences which affected trading and pricing of the security in the more distant past. This may be a good thing insofar as some influences two years in the past may no longer be significant for calculating today's volatility. The impact of that event will have disappeared.

Date	CAC 40	Log$_e$ returns	Annualised volatility
07 Dec 01	4667.87	−0.002463	26.59%
06 Dec 01	4679.38	0.000391	27.26%
05 Dec 01	4677.55	0.032220	28.33%
04 Dec 01	4529.24	0.014545	26.84%
03 Dec 01	4463.84	−0.002734	27.56%
30 Nov 01	4476.06	0.006768	27.96%
29 Nov 01	4445.87	0.000083	28.43%
28 Nov 01	4445.50	−0.009029	28.88%
27 Nov 01	4485.82	−0.017361	29.51%
26 Nov 01	4564.38	−0.002066	29.56%
23 Nov 01	4573.82	−0.005664	
22 Nov 01	4599.80	0.009342	
21 Nov 01	4557.03	−0.007976	
20 Nov 01	4593.52	−0.014474	
19 Nov 01	4660.49	0.015829	
16 Nov 01	4587.30	0.002187	
15 Nov 01	4577.28	0.000479	
14 Nov 01	4575.09	0.002059	
13 Nov 01	4565.68	0.042300	
12 Nov 01	4376.58	−0.030978	
09 Nov 01	4514.28	−0.012932	
08 Nov 01	4573.04	0.017597	
07 Nov 01	4493.27	0.007053	
06 Nov 01	4461.69	−0.005287	
05 Nov 01	4485.34	0.026189	
02 Nov 01	4369.40	0.000563	
01 Nov 01	4366.94	0.005891	
31 Oct 01	4341.29	0.020799	
30 Oct 01	4251.93	−0.030415	
29 Oct 01	4383.24	−0.021529	
26 Oct 01	4478.63	0.022622	
25 Oct 01	4378.45	−0.024345	
24 Oct 01	4486.35	0.006842	
23 Oct 01	4455.76	0.026026	
22 Oct 01	4341.29	0.017755	
19 Oct 01	4264.89	−0.017018	
18 Oct 01	4338.09	−0.016783	
17 Oct 01	4411.51	0.023630	
16 Oct 01	4308.49	0.022197	
15 Oct 01	4213.91		

Table 11.1: Rolling window volatility

2. Exponentially weighted moving average (EWMA) method

This approach sets out to estimate volatility by adopting a weighting process. One way in which this is carried out is to find today's variance by weighting the previous time period's variance and adding to that figure a weighted version of the square of the difference between the previous period's model of forecasted volatility and the realised volatility. This can be expressed algebraically as:

$$\sigma_t^2 = \lambda\sigma_{t-1}^2 + (1-\lambda)u_{t-1}^2 \tag{11.11}$$

where

σ_{t-i}^2 represents the variance of returns at period $t-i$, $i = 0, 1$;
λ is a weighting factor and $0 \leq \lambda \leq 1$.

The spreadsheet entitled **EWMA** on the accompanying CD provides a demonstration of how the method works. The initial setting for λ is 0.9, which yields a forecast volatility for 7 December of 26.11%, but the reader is encouraged to use different values and see the impact this has on the one step ahead forecast of volatility. On an intuitive level it can be seen that a high value of λ will place high emphasis on the previous period's realised volatility, a low value will accentuate the role of the error term, u. The implication of the latter is that volatility will react slowly to surprises or shocks that have occurred in the market.

A moment's reflection will reveal a possible advantage of the EWMA approach. This is that once the process has been initialised only the last period's realised volatility and error term are required to project one period forward. So data storage and complicated algorithms are not required for obtaining estimates of volatility.

3. GARCH modelling

The General Autoregressive Conditional Heteroscedasticity (GARCH) model was first proposed by Bollerslev in 1986. In its simplest form the GARCH(1, 1) model is very similar to the EWMA model discussed in the previous section.

The GARCH(1, 1) model can be expressed formally as:

$$\sigma_t^2 = \bar{\omega} + \alpha\sigma_{t-1}^2 + \beta u_{t-1}^2 \tag{11.12}$$

where

u_{t-1}^2 represents the difference between the last period's observed and predicted values, the same term referred to as the surprise in the EWMA model.
$\bar{\omega}$ represents an important constant term and clearly marks a difference between the current volatility modelling approach and the EWMA approach. The constant term $\bar{\omega}$ can be broken down into two component parts: a long-run average variance part (V) and a constant weight γ which determines the influence that V has in determining today's variance.
α and β are independent coefficients whose estimated values also serve as weights when determining the impact that the previous period's volatility and surprise will have on today's volatility, and $0 < \alpha + \beta \leq 1$.[4]

[4] For a comparison of the exponential smoothing and GARCH approach to estimation of variance refer to Giannopoulos, K. and Eales, B., "Educated Estimates", *Futures and Options World*, Issue 299, April 1996, pp. 45–47.

In the GARCH model the weights γ, α and β must sum to unity:

$$\gamma + \alpha + \beta = 1. \tag{11.13}$$

In this model the constant term, $\bar{\omega}$, and the coefficients α and β are jointly determined using advanced econometric procedures. Once they have been estimated the long-run variance can be readily identified. From equation (11.13):

$$\gamma = 1 - \alpha - \beta$$

and

$$V = \frac{\bar{\omega}}{\gamma}. \tag{11.14}$$

Whilst this type of equation is more costly to obtain than its exponential smoothing or simple standard deviation counterparts, research suggests that following periods of high activity in the marketplace GARCH models may provide better estimates of current variance and instrument volatility.

Although requiring large data sets and the electronic storage capacity that goes with them, recent developments in econometric software support packages have provided researchers and market practitioners alike with reliable tools with which to estimate the variance (volatility) parameter. The reader though is encouraged to use specialist software when estimating GARCH-family models for serious purposes.

In the same way the risk-free rate of interest is an input whose value is determined by market factors outside the option user's control.

Turning now to the assumptions on which the model is based, it is clear that in some respects they are untenable. Markets do not operate continuously; they close at weekends and for bank, religious and other public holidays, and national exchanges are not open for twenty four hours a day. The question then arises: in the option valuation process should this variable be counted as the number of calendar days to expiration or the number of *trading days* to expiration?

The assumption concerning constancy of the risk-free rate of interest is also a problem. The tendency for interest rates to change frequently and sometimes with little forewarning is a well-documented feature of modern, developed economies. Even if the user is prepared to overlook this shortcoming in the case of equity options, when the underlying security on which an option is being valued is an interest rate product, making use of a constant risk-free rate of interest to price an option, whilst at the same time taking out an option to insure against loss arising from a change in interest rates, would seem to be a dubious practice.

Assuming away dividends is also questionable; owners of shares do receive dividends at certain times in the year, in the UK this normally takes place on a semi-annual basis, but in the United States it occurs on a quarterly basis so the chances are that for some options series dividends will fall due and will be paid during the life-span of the option. If this is the case the share price could experience a jump bringing into question the assumption about the behaviour pattern that share prices are supposed to follow and, in consequence, the way in which the volatility input is estimated. Moreover, the timing of dividend payment will also have a direct impact on the estimation of an option's premium.

Clearly question marks can also be raised against the remaining inputs: zero transactions cost, no bid–offer spread, etc.

However, despite the many criticisms that can be made – and some documented systematic biases that have been shown to exist in certain applications of the formula – from a practical and a philosophical standpoint, the fact remains that the B&S option pricing model as an estimator of options premia is robust and provides a good benchmark against which other pricing models can be compared.

The B&S formula for the pricing of a call option based on the assumptions outlined above and the inputs described above can be written as:

$$C(S, E, t, r, \sigma) = SN(d_1) - Ee^{-rt}N(d_2). \tag{11.15}$$

While, making use of the put–call parity theorem, the price of a put option (P) can be written as:

$$P = C - S + Ee^{-rt} \tag{11.16}$$

where

- S represents the current share price;
- E represents the strike (exercise) price;
- t is the time to expiration as a proportion of a year;
- r represents the period effective risk-free rate of interest as a decimal;
- σ represents the standard deviation of the continuously compounded annual rate of return on the share volatility;
- $\log_e$ represents logarithms to base e (natural logarithms).

$$d_1 = \frac{\left[\log_e\left(\frac{S}{E}\right) + \left(r + \frac{\sigma^2}{2}\right) \times t\right]}{\sigma\sqrt{t}} \tag{11.17a}$$

$$d_2 = \frac{\left[\log_e\left(\frac{S}{E}\right) + \left(r - \frac{\sigma^2}{2}\right) \times t\right]}{\sigma\sqrt{t}} \tag{11.17b}$$

or

$$d_2 = d_1 - \sigma\sqrt{t}. \tag{11.17c}$$

Example 11.1

To demonstrate how this works consider the following example where the input data is as presented in Table 11.2.

Using this data to calculate the required d_1 and d_2 inputs:

$$d_1 = \left[\frac{\log_e\left(\frac{8}{8}\right) + \left(0.0425 + \frac{0.30^2}{2}\right) \times \frac{365}{365}}{0.30\sqrt{365/365}}\right] = 0.2916666 \tag{11.17a$'$}$$

and using equation (11.17c) $d_2 = 0.29167 - 0.3 = -0.00833$.

Stock index option pricer **Input the following values:**	
Days to expiration	365
Strike Price CALL	8
Strike Price PUT	8
% Volatility p.a.	30.00
Current underlying price	8
Rate of interest (%) p.a.	4.25
Approx. dividend yield (%) p.a.	0.00

Table 11.2

These need to be converted into normal distribution values and this can be achieved easily on Excel with the NORMSDIST() function.

$N(d_1)$ and $N(d_2)$ are found to be 0.61473 and 0.49668, respectively. Inserting these values into equation (11.15) yields:

$$C = 8(0.61473) - 8e^{(-0.0425)(1)}(0.49668) = 1.10973. \tag{11.15$'$}$$

Finally adopting the put–call parity formula in equation (11.16) the put premium is found to be 0.77685. The put and call premia, rounded to three decimal places, are displayed in Table 11.3 and are found from the basic Black and Scholes Excel spreadsheet (**B&S Basic**) on the accompanying CD.

An important extension to the basic formula presented above allows for the introduction of dividends (q). The formulae introduced in equations (11.15–11.17b) need to be modified to reflect the addition of this input variable. Equation (11.15) becomes:

$$C(S, E, T, r, q, \sigma) = Se^{-qt}N(d_1) - Ee^{-rt}N(d_2). \tag{11.18}$$

Equation (11.17a) becomes:

$$d_1 = \frac{\left[\log_e\left(\dfrac{S}{E}\right) + \left(r - q + \dfrac{\sigma^2}{2}\right) \times t\right]}{\sigma\sqrt{t}}. \tag{11.19}$$

Equation (11.17b) becomes:

$$d_2 = \frac{\left[\log_e\left(\dfrac{S}{E}\right) + \left(r - q - \dfrac{\sigma^2}{2}\right) \times t\right]}{\sigma\sqrt{t}} \tag{11.20}$$

and the put–call parity formula, Equation (11.16) is modified to become:

$$P = C - Se^{-qt} + Ee^{-rt}. \tag{11.21}$$

Option type	**Put**	**Call**
Theoretical premium	0.777	1.110
Strike	8	8

Table 11.3

Option type	Put	Call
Theoretical premium	0.816	1.049
Strike	8	8

Table 11.4

If a continuous dividend yield of 1.25% is paid on this stock, using the **B&S Basic** spreadsheet on the CD, the European-style put and call will be priced as shown in Table 11.4.

11.3 Alternative pricing frameworks

Having developed the Black and Scholes continuous time pricing framework in the preceding sections of this chapter, attention is now focused on the use of trees and simulation approaches to the valuation of options.

In modern life many financial instruments have special features, which make obtaining an accurate price – in a standard, continuous time method of pricing – difficult. The previous section of this chapter examined the Black and Scholes method of pricing a regular call or put option. This section will introduce some alternative pricing frameworks and will apply those approaches to pricing a selection of exotic options.

The first pricing framework to be considered is the binomial approach. This method provides an intuitively attractive approach to the pricing of options. In its most basic form it is simple to use and yet affords an insight into the potential power and flexibility it can offer. Having gained an understanding of the principles of the binomial approach to pricing a third branch to the tree is added. The trinomial method is then looked at in detail. It is compared to its binomial counterpart and its advantages discussed. Finally some very flexible and appealing but often slow and computer-memory-hungry simulation methods are discussed.

11.3.1 *Binomial pricing*

The binomial pricing method provides a useful vehicle for gaining an insight into option pricing and hedging. It also provides a good way of breaking down the path of the underlying security – and hence the behaviour of the option price – through time and therefore an opportunity to study what happens at discrete points in time.

A one-period example will demonstrate how this works in practice and will serve to identify the inputs required to generate an option price and put in place a risk-neutral hedge. Imagine that, over the one time period being considered, the price of the underlying security can rise by a factor equal to u or fall by a factor d.

The interpretation of Figure 11.2 is that from a starting price (S) the underlying security will rise to a value of S multiplied by u or fall to a value of S multiplied by d.

The evolution of the option price is similar. The case of a call option appears as illustrated in Figure 11.3. Here the initial price of the call (C) will take a value of C_u if S rises to uS. The actual numerical value of C_u will be determined by using a modified version

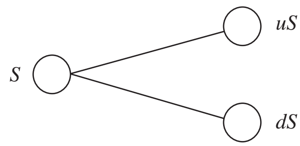

Figure 11.2

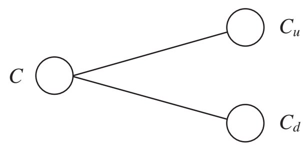

Figure 11.3

of the decision rule introduced in Section 11.2, which described the Black and Scholes model:

$$\max[uS - E, 0].$$

The payoff C_d is determined similarly by using the decision rule:

$$\max[dS - E, 0].$$

The path followed by a put option is presented in Figure 11.4 and will be found by applying the rule: $\max[E - uS, 0]$ if S moves to uS or $\max[E - dS, 0]$ if S moves to dS by the end of the period.

One way in which the call option premium can be found is to set up a hedged portfolio which effectively results in the establishment of a risk-neutral portfolio. Such a portfolio requires a quantity of the underlying security (h) to be held against a written option position. The value of h is chosen such that the value of the portfolio remains unaffected by movements up or down over the time period under consideration.

Algebraically this suggests a relationship of the form:

$$huS - C_u = hdS - C_d$$

which isolating h yields:

$$h = \left(\frac{C_u - C_d}{(u-d)S}\right). \tag{11.22}$$

Introducing the risk-free interest rate allows the development of two more equations:

$$huS - C_u = e^{rT}(hS - C) \tag{11.22a}$$

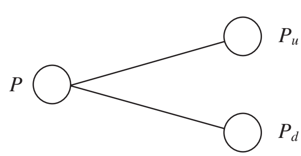

Figure 11.4

or equivalently:

$$hdS - C_d = e^{rT}(hS - C) \tag{11.22b}$$

which implies that the return on the portfolio is growing at the risk-free rate r over time period T. Given values for u and d, defining T, and observing S and r at start up, h can first be calculated using equation (11.22) and then the option premium calculated using either (11.22a) or (11.22b).

The parameters u and d are crucial in the process of determining the option premium even for this simple example. In the simple binomial model the assumption is often made that:

$$u = 1/d \tag{11.23}$$

and that:

$$u = e^{\sigma\sqrt{\Delta t}}. \tag{11.24}$$

Adopting that convention here the one-period option premium can be calculated.

Example 11.2: Pricing a call option

Assume that the period concerned is exactly one year, that r is 4.25% p.a., the option's strike (E) is set at £8, the price of the underlying security (S) currently prevailing in the market is £8, annualised volatility (σ) is 30%.

Using equation (11.24) the value of u is found to be 1.3499 and d, using equation (11.23), 0.7408 to four decimal places. These values now enable the risk-neutral hedge to be established through equation (11.22).

The numerator in equation (11.22) is found using the decision rules and will be: £10.7989 − £8.00 = £2.7989.

The denominator in equation (11.22) is given by $(1.3499 - 0.7408)(8) = 4.8728$. Taking the ratio of these two values yields $2.7989/4.8728 = 0.5744$, which is the instantaneous hedge ratio (h).

Thus to hedge a written call option 0.5744 of the underlying security needs to be held in the portfolio. Now solving equation (11.22b) for the unknown option premium C:

$$C = hS - \left[e^{-rT}(hdS - C_d)\right] \tag{11.25}$$

$$C = (0.5744)(8) - \left[e^{-0.0425}((0.5744)(0.7408)(8) - 0)\right] = 1.3328.$$

Thus the call option premium is found to be approximately £1.33.[5]

Multi-period binomial models can be built on this approach. It is, however, more convenient to develop a more elegant formulation of equation (11.25) when solving for C (or P). Taking as a starting point equation (11.25) and rearranging yields:

$$C = hS - \frac{hdS}{e^{rT}} + \frac{C_d}{e^{rT}}. \tag{11.26}$$

5 The reader is left to verify that application of equation (11.22a) leads to the same result.

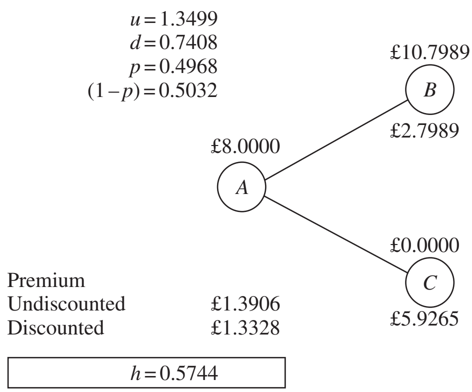

Figure 11.5

Substituting equation (11.22) for h gives:

$$\frac{(e^{rT} - d)(C_u - C_d)}{e^{rT}(u - d)} + \frac{C_d}{e^{rT}}$$

$$= e^{-rT}\left[\frac{(e^{rT} - d)}{(u - d)}C_u - \frac{(e^{rT} - d)}{(u - d)}C_d + C_d\right]$$

$$= e^{-rT}[pC_u - pC_d + C_d]$$

$$= C = e^{-rT}[pC_u + (1 - p)C_d] \tag{11.27}$$

where

p $= (e^{rT} - d)/(u - d)$;
e is the exponential.

Equation (11.27) provides a convenient first step in the generalisation of option pricing using the binomial process. In particular equation (11.27) establishes the formula for pricing a call option over one period. In a two-period case the underlying price starts from either uS or dS. Likewise the option value starts from either C_u or C_d. Figure 11.6 imparts clearly the implications that this simple extension has on the pricing mechanism.

Nodes D and F can only be reached if every movement in the underlying price is up (D) or down (F). Node E though can be arrived at in two ways:

A–B–E

or

A–C–E.

In other words the end nodes *D–E–F* have 1, 2, 1 approach routes, respectively. If the problem were to be extended to a three-period case with four end nodes *G–H–I–J*, the approach routes would number 1, 3, 3, 1.

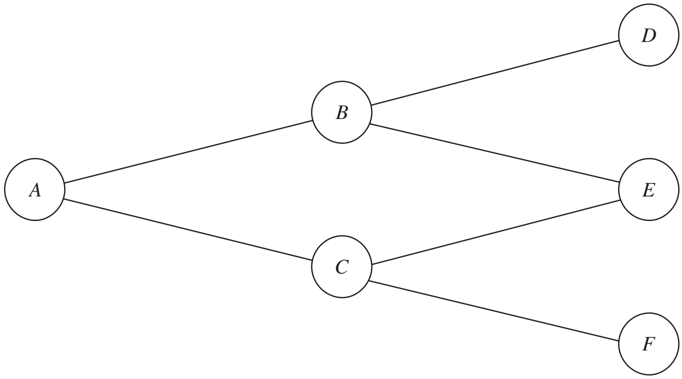

Figure 11.6

Solving explicitly for the option price at node B yields a similar equation to that obtained in equation (11.27):

$$C_u = e^{-r\Delta t}[pC_{uu} + (1-p)C_{du}] \tag{11.27a}$$

where

C_{uu} and C_{du}	are found at nodes D and E, respectively;
Δt	represents the time slice used in the analysis. Thus if T in the one-period case was set equal to one year Δt now represents 0.5 of a year.

$$C_d = e^{-r\Delta t}[pC_{ud} + (1-p)C_{dd}]. \tag{11.27b}$$

Substituting (11.27a) and (11.27b) into equation (11.27) solves for the call option price at start-up. The result is shown in equation (11.28).

$$\begin{aligned} C &= e^{-nr\Delta t}\left[e^{-r\Delta t}(p(pC_{uu} + (1-p)C_{du}) + (1-p)(pC_{ud} + (1-p)C_{dd}))\right] \\ &= e^{-2r\Delta t}\left[\left(p^2C_{uu} + 2(1-p)(p)C_{ud} + (1-p)^2C_{dd}\right)\right]. \end{aligned} \tag{11.28}$$

Notice that the discount factor in (11.28) reflects the whole period (T) over which the option is to run and n is set equal to 2 to reflect the number of segments into which the total period is being divided. This implies that p must also be modified to reflect the shortened time span that each node connecting arc now covers. In the one-period case p was set equal to 0.4968 based on a u value of 1.3499 and d 0.7408. To reflect that two time slices are now involved over the one-year period p becomes 0.4974 based on a u of 1.2363 and a d of 0.8089.

The general formula for solving for an option price in this way appears as equation (11.29). Note that the calculations involved can be dramatically reduced by recognising that not all end node ($S-E$) values will be greater than zero. If an end node value is not positive it, and all subsequent nodes, can be omitted from further calculations in arriving at the option price. The general formula for calculating a call option price using the binomial framework with n time slices is presented in equation (11.29).

$$C = e^{-nr\Delta t}\left[\sum_{q=i}^{n}\frac{n!}{(n-q)!q!}p^{q}(1-p)^{(n-q)}\left(u^{q}d^{(n-q)}S-E\right)\right] \tag{11.29}$$

where

- n represents the number of sub-periods into which the life of the option is divided;
- q represents the number of successes (i.e. the number of times that the underlying share price finishes in-the-money) from the total of n sub-periods;
- i represents the lowest number of up moves in the share prices that lead to the call option finishing in-the-money. (Note that for $q \geq i$ there will be a positive payout to the option, but that for $q \leq i$ the payout will be set to zero.) ! denotes the factorial operator.

This is the approach adopted in the spreadsheet entitled **BINOM.XLS** on the accompanying CD. Using this spreadsheet the reader can gain an insight into the operation of the binomial method of option pricing. In particular, setting the number of time slices to 4, and viewing the **Trees** sheet the reader can see how the underlying price and the option price unfold over four periods. Figure 11.7 gives an illustration of this based on the data inputs in Table 11.5.

The one-year life of the option is therefore broken down into four three-month sub-periods. Equations (11.23) and (11.24) provide vehicles for the estimation of u and d, these are 1.1618 and 0.8607, respectively. After three months the underlying is valued at £9.29 if the price rises or £6.89 if the price falls. The instantaneous delta hedge parameter is found to be, approximately, 0.53743 using equation (11.22).

Only at expiration nodes K and L does the option have intrinsic value. Using $\max[0, S_T - E]$ to determine payout at expiration at node K the intrinsic value is £14.58 − £8.00 = £6.58 while at node L the value is £10.80 − £8.00 = £2.80, all other end nodes display a value of zero: node M has values £8.00 − £8.00 = 0, N has £5.93 − £8.00 and O has £4.39 − £8.00. The last two values for $S_T - E$ are negative and so the investor will let the option expire worthless and receive a zero payout. Working back through the tree from nodes K and L to node G generates a "p" "$1-p$" weighted option value.

Inputs required	**Data**
No. of sub-periods	4
Risk-free rate of interest p.a.	4.25%
Volatility p.a.	30%
Life of option (months)	12
Current price of the underlying	£8
Strike price	£8

Table 11.5

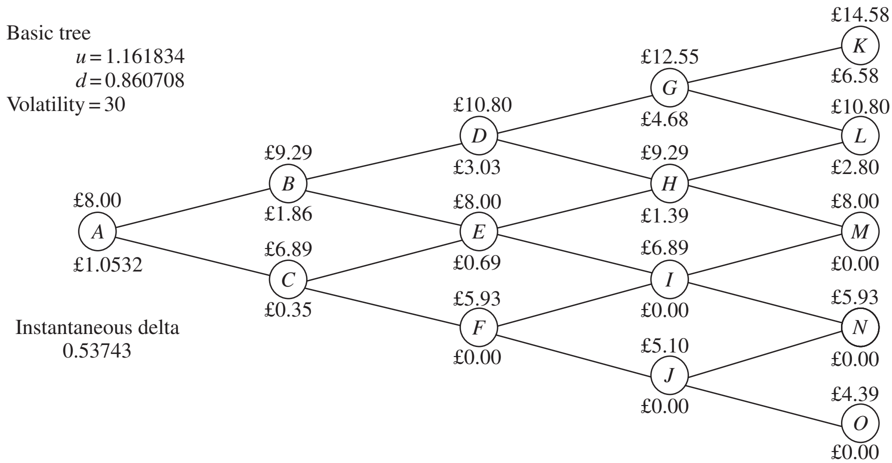

Figure 11.7

Explicitly:

$$(\text{£}6.58p + \text{£}2.80(1-p)) = Call_G$$

$$(\text{£}6.58)(0.4980) + (\text{£}2.80)(0.5020) = \text{£}4.68.$$

Similarly for node H:

$$(\text{£}2.80p + \text{£}0(1-p)) = Call_H$$

$$(\text{£}2.80)(0.4980) + (\text{£}0)(0.5020) = \text{£}1.39.$$

Clearly the values for I and J will be zero. This backward process is repeated until node A is reached where the undiscounted call option price is found to be £1.10. This value must be discounted to take account of the time span of the option's life: £1.10 $e^{(-0.0425)(1)} =$ £1.05.

In the case of a European-style put option the premium can be found using put–call parity, namely:

$$P = C + Ee^{(-r)(t)} - S \tag{11.30}$$

$$P = \text{£}1.05 + \text{£}8e^{(-0.0425)(1)} - \text{£}8 = \text{£}0.72.$$

Thus adopting the put–call parity equation the put's premium is found to be 72p.

This option pricing methodology can be used to take into account payment of dividends, varying interest rate regimes over time, etc., and provides a powerful tool in both the analysis and pricing of certain types of options, particularly those where a closed-form solution is not available. The binomial method, though, does suffer from two particular weaknesses. One weakness is located in the finite number of outcomes that even a large n-step binomial tree produces compared to the infinite number of outcomes that can be drawn from a continuous distribution. The second weakness is a function of the truncation of the binomial tree after a finite number of steps. An example will make this clear.

Example 11.3: Binomial pricing over many branches

Consider the following scenario:

Current share price	£8
Strike	£9
Volatility	30%
Risk-free rate of interest	5%
Time to expiration	12 months

In the basic Black and Scholes model these inputs would result in a call premium of £0.7092. Using this same data in conjunction with the BINOM.XLS spreadsheet the results in Table 11.6 were obtained.

Truncating the tree after 4 steps the call option premium is found to be £0.7564, after 5 steps it is over 10.0% lower at £0.6791. At the 6th and 7th steps there is some evidence of convergence to the Black and Scholes price when compared to the percentage difference in price obtained between steps 4 and 5. The difference is now approximately 9.2%. In either case rounding is not an answer, taking an average between adjacent end nodes may improve the result but could still leave a pricing error. For example, the average price given by the 4th and 5th nodes is 0.7177 compared to a Black and Scholes premium of 0.709. Note that after 98 and 99 steps the difference between the two prices has shrunk to 0.064 and the average of the two premia is 0.7095, which is very close to the Black and Scholes benchmark price of 0.7092.

Number of steps	4	5	6	7	98	99
Call premium	0.7564	0.6791	0.7449	0.6764	0.7076	0.7113

Table 11.6

This see-sawing effect is typical of the binomial pricing method and is still evident even at high branch levels in a tree expansion. Figure 11.8 illustrates how, in the case of an OTM call, the call price oscillates, converging to and then diverging from the Black and Scholes benchmark as the number of branches in the tree is increased.

A very good reason for using this discrete time approach to pricing an option becomes evident when an element of time dependency is involved in the option specification or an event affecting the option's existence occurs at defined points in time. Two examples will help to illustrate why this is so.

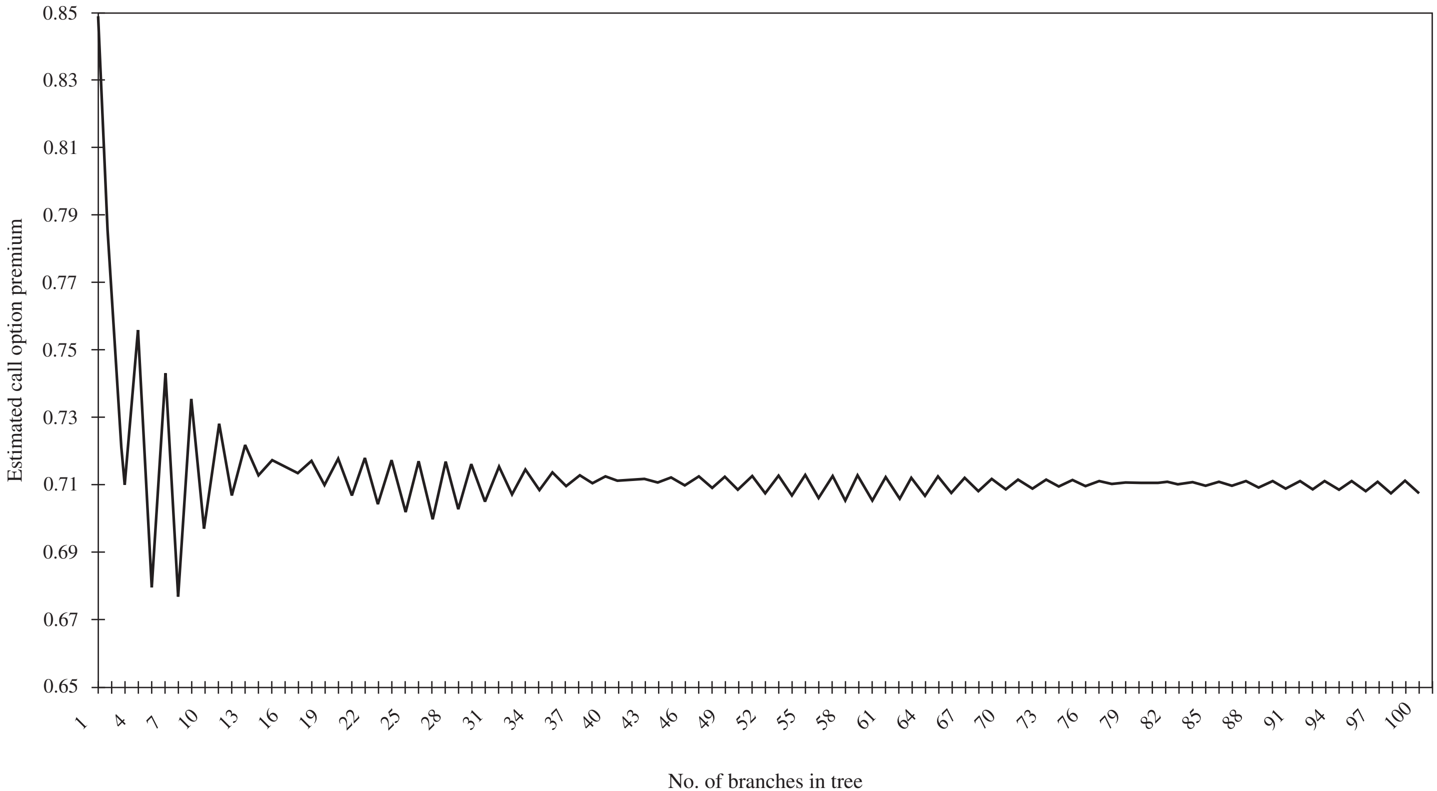

Figure 11.8: Binomial see-sawing effect

Example 11.4: A callable bond

A 4-year, coupon-bearing bond is issued with an embedded call option. The bond is callable over a short window of time at Year 2 at a price, per nominal 100, of 108, similarly in Year 3 the bond can be called by the issuer but now at a price of 106 per nominal 100. The investor or holder of the bond has in effect written a call option on the bond. An intuitively appealing way of pricing this type of structure is to use a tree

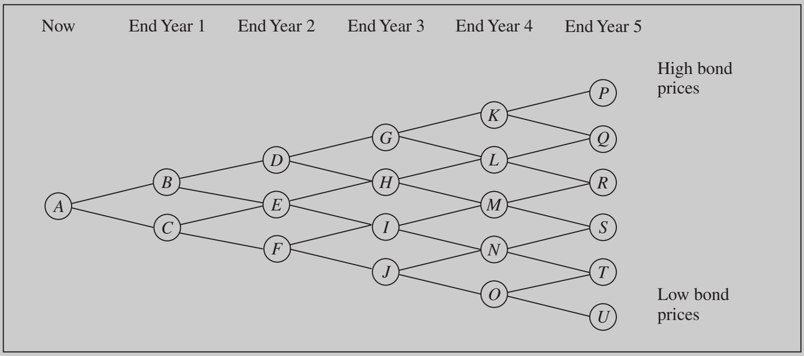

Bond not called

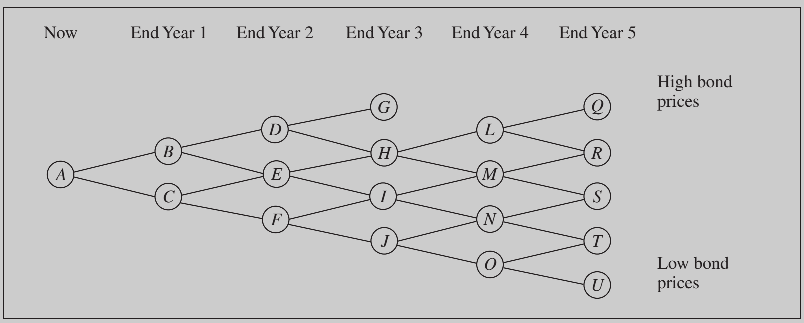

Bond called end of Year 3

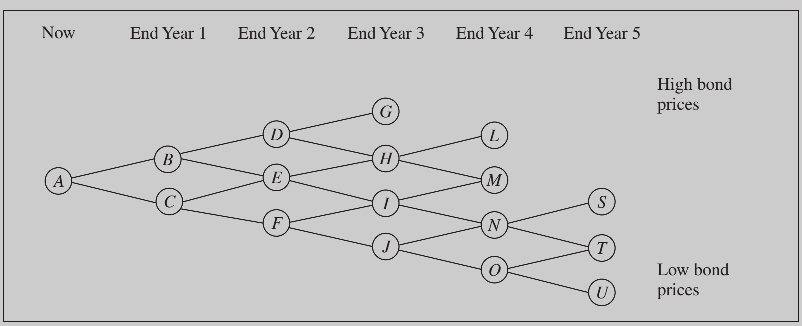

Bond called end of Years 3 and 4

Figure 11.9

approach. Figure 11.9 provides a diagrammatic insight into the way a binomial pricing can be adapted to account for some path dependency in an option's payout calculation.

Notice that several branches of the tree have been lopped so that the option premium calculated back to node *A* will be based on fewer values than if the tree, and the instrument, had run the full course.

Example 11.5: Barrier options

Another and more involved example can be found in the form of barrier options, sometimes called knock-in or knock-out options. There are many types of barrier option available, double strike caps or N-cap options (which in the case of a knock-out style option will replace one cap with a second cap during that period of the transaction), double barriers that define a range for the underlying instrument which must not be breached during the lifetime of the option for a payout to be received, and knock-in/knock-out floors. However, in their simplest form barrier options can be categorised as indicated in Table 11.7.

As Table 11.7 illustrates these calls or puts can be specified in such a way that they either cease to exist when a defined barrier is touched (out) or come into existence when a barrier is touched (in). In many cases barrier options are sold with a rebate clause which effectively returns a portion of the option premium to the holder in the event that the underlying security fails to touch or breach the barrier during the instrument's lifetime.

Calls	**Puts**
Down-and-in	Down-and-in
Down-and-out	Down-and-out
Up-and-in	Up-and-in
Up-and-out	Up-and-out

Table 11.7

The tree approach to pricing offers a number of advantages over its continuous time counterpart. One of the disadvantages of continuous time models concerns the monitoring of the track of the underlying security in order to determine whether the barrier has at some point been touched. The Black and Scholes closed-form pricing formula developed by Rubinstein and Reiner (1991) assumes that the track is monitored continuously. Using binomial (or trinomial) trees permits the more realistic assumption that monitoring takes place at discrete but defined intervals of time by reference to Reuters, Telerate or Bloomberg quote rates.

Figure 11.10 shows a down-and-in barrier put option with a barrier set at 680p, a value which lies below the strike of 690p which in its turn is below the current underlying security quote at 740p. Clearly in this discrete time framework a pricing approach which allows discrete time monitoring over a 20-day period is going to provide a more appropriate price for the option.

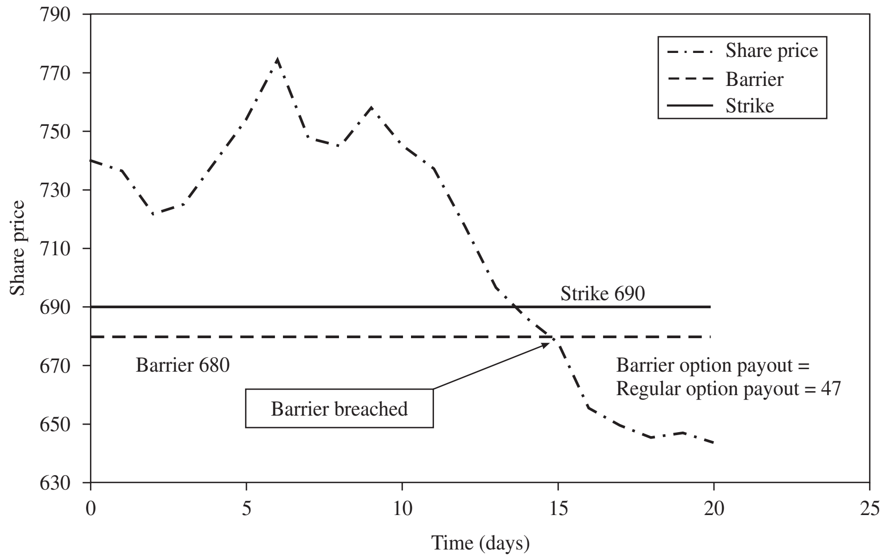

Figure 11.10: Barrier option

Inputs required	**Data**
No. of sub-periods	4
Risk-free rate of interest p.a.	4.25%
Volatility p.a.	30%
Life of option (months)	12
Current price of the underlying	£8
Strike price	£8
Barrier	£9

Table 11.8

Taking the data used earlier to develop the idea of binomial pricing consider now the pricing problem for the case of an up-and-out put option. If the barrier is set at nine then for any nodes where the underlying security is priced at or above nine subsequent arcs will be pruned from the tree.

Figure 11.11 shows the effect that this will have on the tree and the price of the specified barrier option.

The price of an up-and-in barrier put option using this four-branch tree is found to be 0.25p compared to the 72p required for a regular put option. Although four branches will not be sufficient to arrive at a reasonably accurate option price the tree nevertheless illustrates the effectiveness of the binomial pricing method.

11.3.2 *Trinomial pricing*

A trinomial tree introduces more flexibility into the analysis of option prices and hedge parameters. Instead of limiting movements in the underlying security to up (u) and down (d) an additional in-between movement is permitted, m. In many cases the m move is

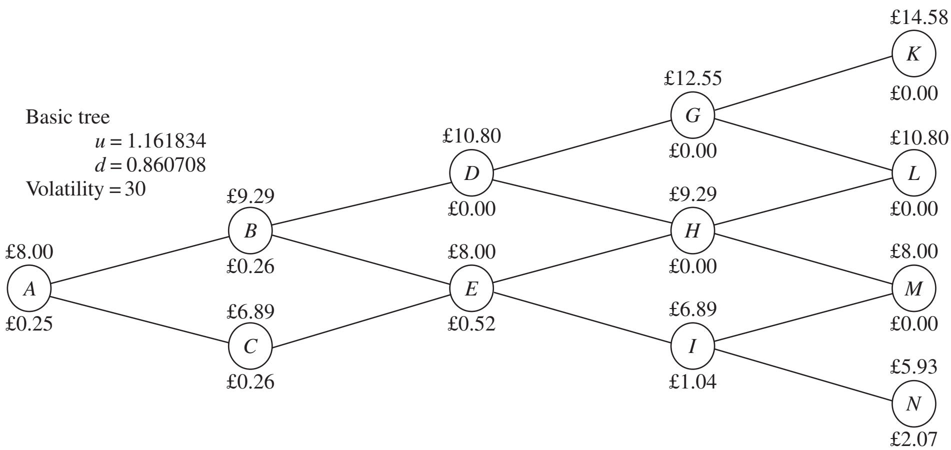

Figure 11.11: Four-branch binomial tree with branches pruned

defined as being no change in price, but this limitation can be removed in practice.[6] The probabilities associated with the three specified moves are p_u, p_d and p_m, respectively.

The addition of this third branch at each level of the tree clearly has the effect of increasing the number of nodes at all steps throughout the tree. Figure 11.12 illustrates this point graphically.

If node A is defined as $n = 0$, B, C, D as $n = 1$ etc., then the number of end nodes at any step can be calculated using $(2n + 1)$, i.e. at $n = 1$ there will be three nodes (B, C, D) while at step $n = 2$ there will be five nodes. This compares to the $n + 1$ nodes that will be constructed when using the same step designation in the binomial model. With the exception of $n = 0$ there will, of course, also be an increase in the total number of nodes involved in the tree. For $n = 1$ the binomial tree has two end nodes and three nodes in total compared to the trinomial's three end nodes and a total of four nodes. At step $n = 5$ the increase becomes more marked – the binomial will have only five end nodes and 21 nodes in total compared with the 11 end nodes and total of 36 nodes generated in the trinomial framework.

If N is used to denote the total number of nodes in the tree then the binomial model will generate a total of $(N^2 + 3N + 2)/2$ nodes compared to $(N^2 + 2N + 2)$ in the trinomial case. In general the aggregate extra nodes required using the trinomial model can be calculated as shown in equation (11.30):

$$\frac{N^2 + N}{2}. \tag{11.30}$$

It is self-evident that when fine-tuning the calculation of option premia and their associated hedge parameters computation time will increase greatly.

Note that $\Delta \ln S$ is assumed to be driven by $(r - d - \sigma^2/2)\Delta t + \sigma \Delta z$ where d denotes the annualised dividend yield, and r, σ, Δt and Δz are as defined earlier in this chapter.

6 For a demonstration of this refer to Hull, J., *Options, Futures and Other Derivatives*, 4th edition, Prentice Hall, 2000, pp. 405–406, 578–596.

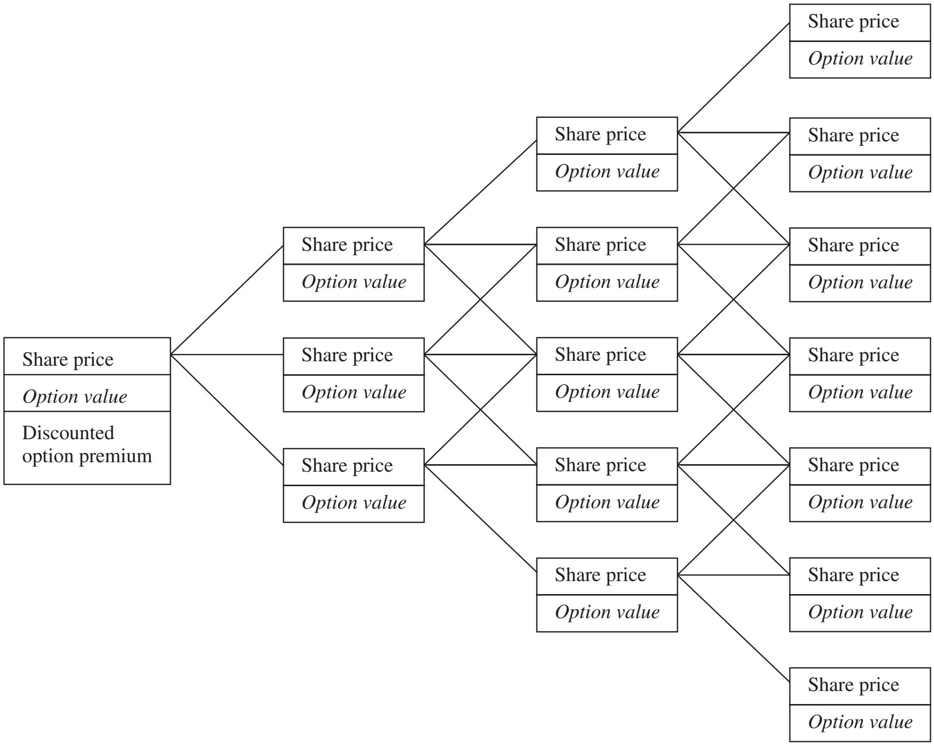

Figure 11.12: Trinomial tree

A standard approach to determining the probabilities of up, down and no change moves is to define a parameter λ such that:

$$\lambda = \frac{h^2}{\sigma^2 k} \tag{11.31}$$

where $k = T/n$, i.e. the total time period divided by the number of steps required.

The probabilities are then found using the following equations:

$$p_u = 0.5\left[\sigma^2 \frac{k}{h^2} + \gamma^2 \frac{k^2}{h^2} + \gamma \frac{k}{h}\right]$$

$$p_d = 0.5\left[\sigma^2 \frac{k}{h^2} - \gamma^2 \frac{k^2}{h^2} - \gamma \frac{k}{h}\right]$$

$$p_m = 1 - p_u - p_d = 1 - \frac{1}{\lambda} \tag{11.32}$$

where $\gamma = (r - d - \sigma^2/2)$.

Note that the number of ways in which the up, down or no change moves can occur is given by the equation:

$$\frac{n!}{i!j!(n-i-j)!}$$

where

- n represents the number of steps;
- i represents the number of up moves;
- j represents the number of down moves.

$$C = e^{-nr\Delta t}\left[\sum_{i=0}^{n} \frac{n!}{i!j!(n-i-j)!} p_u^i p_m^{(n-i-j)} p_d^j ((\exp(\ln\ S) + (i_u - j_d)h)(S-E))\right]. \qquad (11.33)$$

Example 11.6: Trinomial pricing over many periods

Assume the same scenario as that adopted for the binomial model (Table 11.9). Convergence in the trinomial model is much faster in terms of the number of branches required to obtain an accurate option price when measured against the Black and Scholes benchmark figure. Truncating the tree after four steps the call option premium is found to be 0.7385, after five steps it is 0.7338, a difference of −0.0047. At the 19th and 20th steps the difference is 0.0006. These differences are of a similar order of magnitude to those obtained in the binomial model after many more branches were included.

Taking the data used earlier to develop the idea of binomial pricing, consider now the pricing problem for the case of an up-and-out put option. If the barrier is set at nine then for any nodes where the underlying security is priced at or above nine subsequent arcs will be pruned from the tree.

Figure 11.13 shows the effect that this will have on the tree and the price of the specified barrier option.

The price of an up-and-in barrier put option using this four-branch tree is found to be 0.51p compared to the 72p required for a regular put option. Although four branches will not be sufficient to arrive at a reasonably accurate option price the tree nevertheless illustrates the effectiveness of the trinomial pricing method. The trinomial approach has an additional strength in that λ can be chosen to ensure that one of the nodes has an underlying security value which is exactly equal to the specified barrier. This should help to facilitate a more rapid and accurate convergence to the barrier option's premium.

Current share price	£8
Strike	£9
Volatility (%) p.a.	30
Risk-free rate of interest (%) p.a.	4.25
Time to expiration	12 months

Number of steps	4	5	19	20
Call premium	0.7385	0.7338	0.6975	0.6981

Table 11.9

Inputs required	Data
No. of sub-periods	4
Risk-free rate of interest p.a.	4.25%
Volatility p.a.	30%
Life of option (months)	12
Current price of the underlying	£8
Strike price	£8
Barrier	£9

Table 11.10

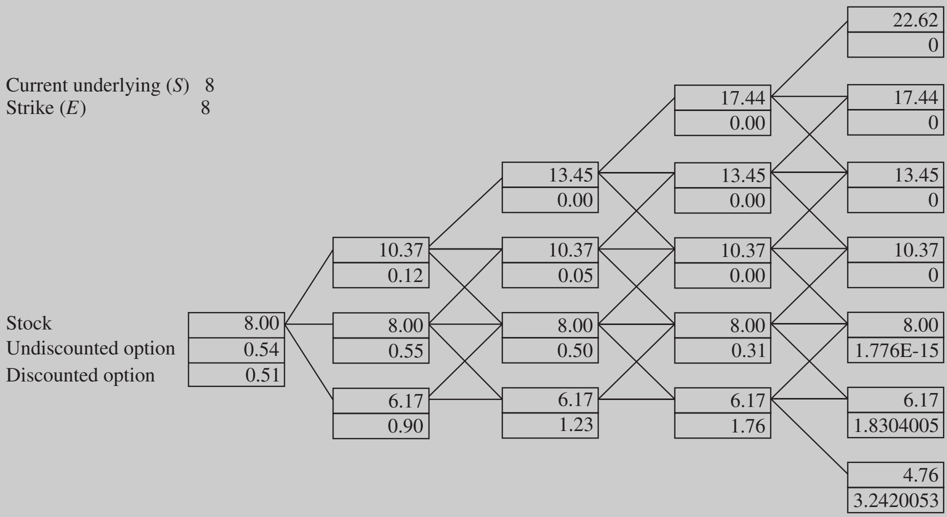

Figure 11.13: Put option: four-branch trinomial tree with branches pruned. Note that in this example λ is set equal to 3. If λ is set to give a value of the stock closer to the £9 barrier than the current £10.37 the option premium falls to £0.26

11.4 Monte Carlo simulation

In the derivation of the B&S option pricing formulae it is assumed that a non-dividend-paying share follows a continuous time stochastic process of the form:

$$dS = \mu S dt + \sigma S dz \tag{11.34}$$

or

$$\frac{dS}{S} = \mu dt + \sigma dz \tag{11.35}$$

and that dS/S is normally distributed as described in equation (11.36):

$$\frac{dS}{S} \approx N\left\{\mu dt, \sigma\sqrt{dt}\right\} \tag{11.36}$$

where the variables and parameters are as previously defined.

Making use of the above definition of dz and rewriting equation (11.35) in terms of discrete time, equation (11.37) is obtained:

$$\frac{\Delta S}{S} = \mu \Delta t + \sigma \varepsilon \sqrt{\Delta t}. \tag{11.37}$$

Equation (11.37) provides a way of generating time paths that the underlying security might follow in order to analyse payoff accruing to various option strategies using both regular and exotic options.

11.4.1 *Generating simulated security time paths*

In order to generate some time paths from equation (11.37) a number of explicit assumptions are required. The time path will be created for a period of 20 days and will be assumed to reflect the closing price of the underlying security – in the interests of consistency the same data will be used as that used to demonstrate the pricing of a barrier option in the binomial and trinomial sections of this chapter. It is further assumed that the time interval involved will be one trading day (i.e. $1/252 = 0.0039683$ of a year), that the annualised return on the underlying is 1.5% with an annualised standard deviation of 30%.

The equivalent daily return on the underlying can be found by multiplying the annualised return by the time interval. Thus the daily return is found to be $(0.0039683) \times (0.015) = 0.00005952$. Similarly the adjusted standard deviation is found to be $(0.30) \times (\sqrt{}0.0039683) = 0.018898$.

$$\frac{\Delta S}{S} N(0.0000592, 0.018898). \tag{11.38}$$

In other words $\Delta S/S$ is normally distributed with mean 0.0000592 and standard deviation 0.018898.

Excel has the facility to generate random numbers under the **Tools menu** in its **Data Analysis** add-on. To run this facility select **Tools, Data Analysis** followed by **Random Number Generator**. At this point a dialogue box appears and allows the user to input the data necessary to generate a set of random numbers. The necessary data to replicate this demonstration will be:

Number of variables	1
Number of random numbers	20
Distribution	Normal
Mean	1
Standard deviation	0

Pressing OK will give 20 random numbers satisfying the specified distribution, which will be placed into a new worksheet. In this example the random numbers are presented in Column 2 of Table 11.11.

In conjunction with Equation (11.37) the random numbers in Column 2 of Table 11.11 are used to create an evolution of the share's path over a 20-day period. The data generation process being undertaken is a simple Monte Carlo simulation. The change in the share price is calculated using the following approach:

$$\Delta S = [(0.015)(0.003968\ldots) + (0.3)(0.062994\ldots)(0.300232\ldots)800] = 4.5.$$

Day	Random number	Share price (p)	Price changes	Query
0		800.0		0
1	0.3002	804.6	4.6	0
2	1.2777	824.1	19.5	0
3	−0.2443	820.3	−3.8	0
4	−1.2765	800.6	−19.7	0
5	−1.1984	782.5	−18.1	0
6	−1.7331	756.9	−25.6	0
7	2.1836	788.2	31.3	0
8	0.2342	791.7	3.5	0
9	−1.0950	775.4	−16.3	0
10	1.0867	791.4	16.0	0
11	0.6902	801.7	10.4	0
12	1.6904	827.4	25.7	0
13	1.8469	856.3	28.9	0
14	0.9776	872.2	15.9	0
15	0.7735	885.0	12.8	0
16	2.1179	920.5	35.5	1
17	0.5679	930.4	9.9	1
18	0.4040	937.5	7.2	1
19	−0.1349	935.2	−2.3	1
20	0.3655	941.7	6.5	1

Table 11.11: Monte Carlo simulated track

This value is added to the starting price of the share (800) which will then become the appropriate share price for calculating the next change, and the base to which the new change is added.

The final column in Table 11.11 poses the question has the barrier been breached or touched. If, when compared to the current share price, the answer is yes, a "1" is inserted and the down-and-in put will be triggered. If the answer is no, a "0" is entered and the barrier option remains unactivated. Over this track the barrier option is activated but the payout will still be zero since the last recorded value of S, assuming the option expires at that point, is £9.42 which is greater than the strike of £8.00. So the put, regular and barrier all expire worthless.

This illustrates the way in which a basic Monte Carlo can be used to analyse and value certain types of option. Clearly the one track generated to illustrate this methodology is insufficient from which to draw sensible conclusions. In practice many hundreds, even thousands of tracks will need to be generated before a realistic conclusion can be drawn. The process is computer intensive and to speed up the method many variations have been developed: the use of antithetic variables (generates an evolution of the share price which is the mirror image of the original set of random numbers); control variates (where benchmark information from an outside source is used in conjunction with a simulation); and bootstrap simulation (this approach does not require estimation of parameters such as mean and standard deviation from a historical time series of an underlying security). In essence the user specifies a period of data from which repeated sampling with replacement is undertaken. The drawings from the sample are made at random and from these drawings tracks are produced from which inference can be drawn about the likely path that the underlying security will follow.

Overall the integrity of an option price found using simulation depends on the number of simulated tracks used to find the price. An advantage that the method has is that repeated experiments can be performed and statistics, such as the mean and standard error, estimated. This will give an indication of how the accuracy of the price of the derivative is improving when the number of experiments is increased. The standard error of the mean result of these experiments is a well-documented statistic presented as equation (11.39):

$$\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{M}} \tag{11.39}$$

where

- $\sigma_{\bar{x}}$ represents the standard error of the sample means;
- σ represents the standard deviation of the experimental results;
- M represents the number of independent experiments performed.

Good sources of literature on these approaches to pricing can be found in Hull (2000), Clewlow and Strickland (1999) and Eales (2000).

Selected bibliography and references

Black, F., Scholes, M., "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy* 81, May–June 1973, pp. 637–659

Clewlow, L., Strickland, C., *Implementing Derivative Models*, Wiley 1999

Cox, J., Rubinstein, M., *Options Markets*, Prentice Hall 1985

Eales, B. A., *Financial Engineering*, Palgrave/Macmillan 2000

Giannopoulos, K., Eales, B., "Educated Estimates", *Futures and Options World* 299, April 1996, pp. 45–47

Hull, J., *Options, Futures and other Derivatives*, 4th edition, Prentice Hall International 2000

Keating, C., Shadwick, B. A., Shadwick, W. F., "Unfocused Model Thinking", *Futures and Options World* 358, March 2001, pp. 76–83

Merton, R. C., "Theory of Rational Option Pricing", *Bell Journal of Economics and Management Science* 3, 1976, pp. 3–54

Rubinstein, E., Reiner, M., "Breaking Down the Barriers", *Risk* 4 (8), September 1991, pp. 28–35

Taleb, N., *Dynamic Hedging: Managing Vanilla and Exotic Options*, Wiley Financial Engineering 1997

12 Equity-linked Structured Products

12.1 Introduction

The reasons for using equity-linked structured products are many and varied. Corporates may be motivated by the need to raise capital at a competitive rate of interest, to fund a merger or acquisition or to manage their balance sheets. Whatever an individual company's reason for using such instruments, cost of capital, capital management and tax considerations are usually of prime importance. These days though even individual investors are using these products as tools that can help to attain their investment goals. In times when equity markets are bullish this might be to forego interest on a deposit account in order to participate in stock market growth. When stock markets fall, investors' funds may be divided between a bond paying a high coupon and equity index participation with a guaranteed return of investment. The paramount goals for the individual investor will be to obtain a higher rate of interest than that available from their local building society/mortgage or bank deposit account, or to participate in rises in a stock market index or maybe both.

A good starting place, and perhaps one that can be described as traditional, is the well-known convertible bond. These instruments, however, have moved on considerably from their original function, namely exchange of a debt instrument for equity, and can possess features which make them difficult to price and hedge effectively. This complexity, however, can in turn lead to potential arbitrage opportunities for other market players to exploit.

12.2 Convertible bonds

Convertible bonds (CBs) represent an exciting but rather neglected section of the fixed income markets. This type of bond endows on its holder the right to purchase a defined quantity of shares at a defined price. This is achieved by returning the bond to the issuer at maturity, or in some cases earlier, and receiving in return the specified amount of equity. In the absence of a company's bankruptcy or default there will be a floor value to the convertible bond, for example, its redemption value at maturity. In the event that the price of the company's shares rises, the value of the bond will also rise but not necessarily in a perfectly correlated way. Expressed simply the bond has features of both fixed income and equity instruments and these features combined create a profile very much akin to that of a call option. This profile is identified in Figure 12.1, which illustrates the traded profile for a Deutsche Bank CB and in Figure 12.3, which presents a theoretical payout structure of a CB at expiration. The basics of CBs will be covered in more detail in Section 12.2.1.

The hybrid equity/fixed income features of CBs make them attractive to a wide cross-section of market participants. At one end of the scale are risk-averse, fixed-income

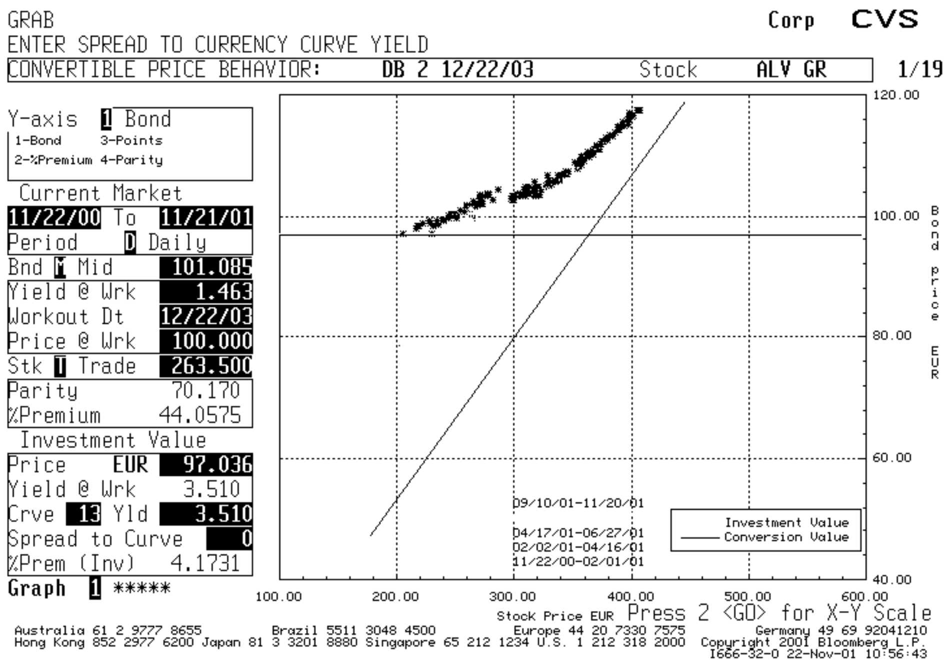

Figure 12.1: The traded profile for a Deutsche Bank convertible bond. © Bloomberg LP. Reproduced with permission.

portfolio managers looking for returns linked to equity exposure while at the other end of the scale are risk-loving investors speculating that distressed paper will recover and, in the process of recovery, deliver a high return. Section 12.3.1 develops an example of the potential of CBs by breaking down the bond into its separate exposures.

12.2.1 *Basics of convertible bonds*

Even in their most basic form convertible bonds can be difficult to analyse and hedge. Nevertheless they offer a wide cross-section of market participants the opportunity to invest indirectly in the shares of a company with limited downside risk. Before discussing these instruments in detail it is essential to get to grips with some market jargon and standard descriptive statistics which are used to describe and analyse CBs. To help in the process it is useful to consider the terms describing a typical convertible bond. Table 12.1 presents the specification, or indenture, of such a bond.

By now expressions such as coupon, maturity and par value will be familiar and need no further explanation. The issue size of £50,000,000 is self-explanatory and is broken down into 50,000 bonds each with a par value of £1000. Other terms such as conversion price, conversion ratio, parity and premium need explanation.

Conversion price refers to the price that the holder of the bond will need to pay when converting the bond into shares. In this illustrative example the conversion price is fixed at £12.50. The number of shares that the holder of the bond can buy at the specified conversion price is fixed by the conversion ratio, here 80.

Coupon (annual)	3.00%
Maturity	26 January 2012
Issue size	£50,000,000
Denomination or par value	£1000
Number of bonds	50,000
Issue price	100%
Conversion price	£12.50
Conversion ratio (shares/bond)	80
Current share price	£8.50
Parity	68%
Conversion premium	53.85%

Table 12.1

As an equation the conversion ratio is defined as:

$$\text{Conversion ratio} = \frac{\text{Denomination}}{\text{Conversion price}} \tag{12.1}$$

$$\text{Conversion ratio} = \frac{£1000}{£12.50} = 80. \tag{12.1$'$}$$

Parity is an important concept in the analysis and understanding of convertible bonds. In a sense parity reflects the in-, at-, or out-of-the-moneyness of the bond at currently prevailing share prices in the market. In the case of a single currency bond, parity can be defined as:

$$\text{Parity} = \text{Share price} \times \text{Conversion ratio} \tag{12.2}$$

but is normally referred to as a percentage of the CB's par value:

$$\text{Parity} = \frac{(\text{Share price}) \times (\text{Conversion ratio})}{\text{Bond par value}} \times 100. \tag{12.3}$$

Thus if the conversion ratio is set at 80 (Denomination/Conversion price) and the market price of the stock at the bond's maturity date is £8.50, parity, expressed as a percentage, will be 68%.

$$\text{Parity} = \left[\frac{(£8.50) \times (80)}{£1000}\right] \times 100 = 68\%. \tag{12.3$'$}$$

As mentioned above, in practice the value found using equation (12.3′) is expressed as a percentage, so that parity would be quoted as 68% and referred to as 68 bond points. At this share price a rational investor will not wish to convert the bond into shares. The bond has a redemption value of £1000 while the market price of the shares is only £680. Viewed from an equity perspective the bond is out-of-the-money with a current parity of 68 bond points.

Repeated application of equation (12.3′) for a continuum of stock prices will generate a parity line. This line appears as depicted in Figure 12.2. The parity line itself indicates the total current market value of the shares associated with the bond were it to be converted at that share price.

Continuing the examination of a CB from the equity point of view, there is evidently a premium involved in gaining exposure to the company's equity through the purchase of the bond rather than a direct purchase of the company's equity. Using the same example as

Figure 12.2: Convertible bond parity line

above an intuitively appealing calculation would be to find the difference between the value of the equity through conversion (£1000) and current parity (£680) yielding a premium of £320. Alternatively premium can be calculated as the difference between the CB's price and parity multiplied by 100 divided by parity. The resulting statistic defines, in percentage terms, the excess that an investor needs to pay in order to gain exposure to the number of shares specified in the conversion ratio. A more formal representation of this statistic is presented in equation (12.4) while equation (12.4′) calculates 47.06% as the premium attached to the CB under examination.

$$\text{Premium} = \frac{(\text{CB price} - \text{Parity})}{\text{Parity}} \times 100 \tag{12.4}$$

$$\text{Premium} = \frac{(1000 - 680)}{680} \times 100 = 47.06\%. \tag{12.4′}$$

This implies that the current share price will have to rise by 47.06% for a zero premium to exist, i.e. Parity = (1.4706)(680) = 1000.

The interplay between rising/falling bond and share prices will lead to continually changing parity and premium statistics as time progresses. To illustrate this point and to consider different ways in which a CB can be quoted consider the following example.

Example 12.1: CB quoting conventions

Consider a convertible bond offering a coupon of 3.00%, a conversion ratio of 800 shares per nominal £5000 bond at a conversion price of £6.25 per share. The convertible is currently trading at £102.20 (no accrued interest) and the price of the share into which the CB can be converted stands at £5.80. Table 12.2 shows different ways in which the CB can be valued.

	Cash value (1)	Market quote in bond points (2)	Value per share (3)
Nominal value	£5000	100.00	£6.25
Conversion price	£5000	100.00	£6.25
Conversion ratio	800		1
Coupon	£150	3.00	£0.1875
Current price	£5110	102.20	£6.3875
Parity	£4640	92.80	£5.8000
Premium	£470	9.40	£0.5875
Premium (%)	10.13%	10.13%	10.13%

Table 12.2

Consider the figures in Column 1 first. The nominal value of the CB is £5000 and its conversion price is also £5000. The conversion ratio is given as 800 shares for each nominal £5000 bond. The coupon is 3.00% p.a. and is thus worth £150. The current market price of the bond is £5110. The figure of £5110 can be determined by multiplying the market quote of £102.20 by the CB's nominal value of £5000 and dividing by 100.

$$\text{Current market price (cash value)} = \left[\frac{(£102.20)(5000)}{100}\right] = £5110.$$

The CB's parity expressed as cash value will be:

$$\text{Parity (cash value)} = (800)(£5.80) = £4640.$$

Premium expressed as cash value will be the difference between the market price of the bond and its cash value:

$$\text{Premium (cash value)} = £5110 - £4640 = £470$$

which as a percentage of parity will be:

$$\text{Premium (\%)} = \left[\frac{£470}{£4640}\right] \times 100 = 10.13\%.$$

Which indicates that the current market price of the share is below the conversion price specified in the CB's term sheet. The investor is paying more for the potential right of buying shares in the company at the specified conversion price some time in the future. Turning now to Column 2, the value of the CB in terms of bond points can be found. This is simply expressing the £5000 unit of trading as a conventional bond quoted off £100 blocks. Reporting a CB in these terms will make direct comparisons with other types of fixed income instruments much easier.

The nominal value of the £5000 CB expressed in bond points will be:

$$\frac{£5000}{50} = 100.$$

The conversion price can be treated in the same way. The conversion ratio is 625 for a £5000 unit of trading. It will not have an interpretation if scaled down to 100 bond points so this cell has been left blank.

In bond points the coupon will be 3. The coupon is reported as a percentage figure in the bond's term sheet.

The parity figure expressed in percentage terms can be found as follows:

$$\text{Parity} = \left[\frac{(800)(5.80)}{5000}\right] \times 100 = 92.80\%.$$

While the premium in bond points expressed as a ratio of nominal value will be:

$$\text{Premium} = \left[\frac{(5110 - 4640)}{5000}\right] \times 100 = 9.40$$

and the premium as a percentage of parity will be:

$$\text{Premium} = \left[\frac{9.40}{92.80}\right] \times 100 = 10.13\%.$$

Finally Column 3 provides an insight into the CB's value but reported on a per share basis.

The nominal value will be £5000 divided by the conversion ratio. Thus 5000/800 yields a value per share of £6.25. This is the same as the defined conversion price and is quoted on a per-share basis. The cash value of the coupon, £150, when divided by the conversion ratio gives a per-share coupon of £0.1875. The price paid for each share if conversion were to take place now would be £6.3875. Parity is the current market price of the share, £5.80. The difference between the current conversion price and parity provides the premium measure in value terms and this is £0.5875 per share. As a percentage this is once again 10.13% (£0.5875/£5.8000).

If that same bond trades one month later at £100.90 (again – unrealistically – assuming no accrued interest) and the share price has now fallen to £5.40 the revised statistics will be as shown in Table 12.3.

Note that the fall in the bond's price and the fall in the share price have combined to increase the premium to 16.78%. This is presented in the last row of Table 12.3.

	Cash value (1)	**Market quote in bond points (2)**	**Value per share (3)**
Nominal value	£5000	100.00	£6.25
Conversion price	£5000	100.00	£6.25
Conversion ratio	800		1
Coupon	£150	3.00	£0.1875
Current price	£5045	100.90	£6.3063
Parity	£4320	86.40	£5.4000
Premium	£725	14.50	£0.9063
Premium (%)	16.78%	16.78%	16.78%

Table 12.3

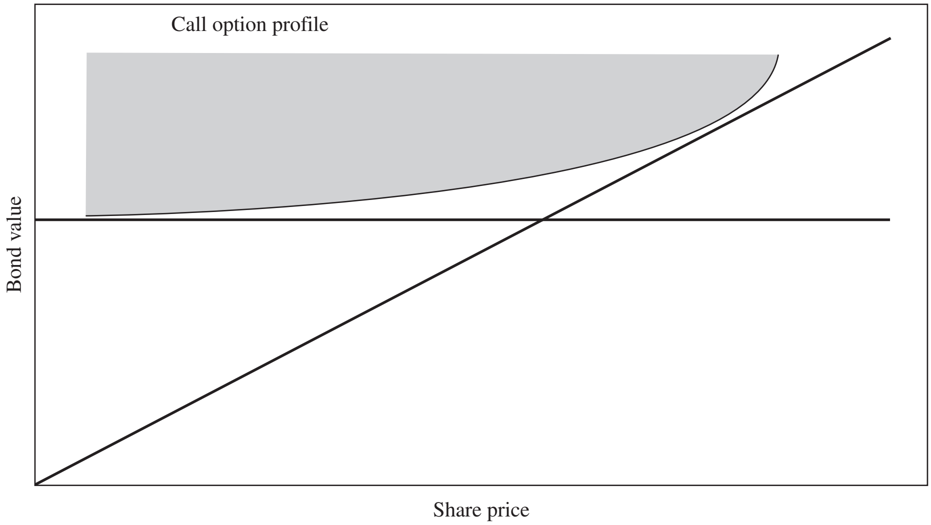

Figure 12.3: Parity curve with bond floor

12.2.2 *The CB floor*

A very attractive feature of CBs, from an investor's view point, is the floor that they offer. To illustrate this idea consider a bullet bond redeemable at 100 in 2010. If this redemption value is plotted together with the parity line an interesting feature becomes evident. The shaded area on Figure 12.3 indicates the call option profile created by the CB's parity line and the bond's floor at redemption. The floor in this example is established by the bond's redemption value while the rising portion of the parity line above and to the right of the floor provides the positive payoff emanating from high equity prices. Some CB issues do provide a protection against this eventuality by having the bond underwritten or guaranteed by a third party with a high, or higher, credit rating than the issuer. This added degree of protection is paid for in the form of a higher bond price which in turn means a lower yield to the bond's holder.

Of course the redemption value will only be realised by the investor should conversion not take place, in effect leaving the bond to be redeemed at maturity for its face value plus final coupon. However, between now and maturity many things can happen. If, for example, the company goes through a bad patch resulting in the price of its stock falling precipitously, the bond price will inevitably suffer. A hypothetical example of this scenario is illustrated in Figure 12.4 where, as the stock price falls below a certain level the bond's price follows and, in the worst case scenario, the investor's floor collapses completely as the value of the bond falls rapidly to zero.

In the example above the floor at maturity was illustrated as 100; in practice the CB's floor can be taken to be the investment value of the bond.

To establish the convertible's floor it is necessary to calculate the straight corporate bond's price and this can be achieved using the familiar bond pricing

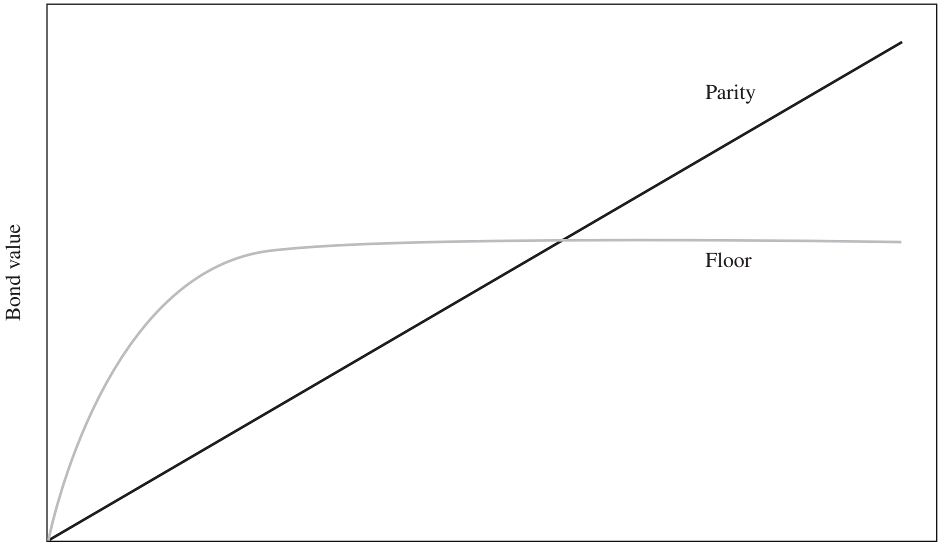

Figure 12.4: Parity line, CB floor and default

formula presented in earlier chapters and repeated here for convenience as Expression (12.5).

$$P_o = \sum \frac{C}{(1+r)^t} + \frac{RV}{(1+r)^T} \tag{12.5}$$

where (for a corporate bond):

C	represents coupon paid annually and fixed throughout the lifetime of the bond;
r	represents the discount rate which includes a risk spread, and r is expressed as a decimal;
RV	represents time to maturity in years;
t	is an index of time in years ($t = 1, 2, \ldots, T$);
T	represents the years to maturity.

In reality the CB's floor at any time will depend on interest rates prevailing in the market and will therefore be subject to change even in the absence of company default or threat of default. As an example consider a 5-year, 6% coupon corporate bond priced off a yield-to-maturity of 7.00% plus credit spread of 150 basis points.[1] In the absence of any accrued interest the floor defined by the price of this bond will be £90.15, calculated as shown in Table 12.4.

If, as a result of government policy, benchmark yields rise to 10% and corporate discount rate rises to 11.5% then the level of the floor will fall to around £79.93, calculated as shown in Table 12.5.

[1] More will be said about pricing CBs in Section 12.2.3.

Year	Cash flows	Discounted cash flows
1	6	5.52995392
2	6	5.09673172
3	6	4.69744859
4	6	4.32944571
5	106	70.4948149
		90.1483948

Table 12.4

Year	Cash flows	Discounted cash flows
1	6	5.38116592
2	6	4.82615778
3	6	4.32839262
4	6	3.88196648
5	106	61.507989
		79.9256718

Table 12.5

Year	Cash flows	Discounted cash flows
1	6	5.68720379
2	6	5.39071449
3	6	5.10968199
4	6	4.84330046
5	106	81.1042415
		102.135142

Table 12.6

Should the reverse happen with required yield falling to 4% + 150 bps for the corporate bond, the floor will rise to approximately £102.14, calculated as shown in Table 12.6.

The instability of the floor can, of course, also be affected by pessimism about the perceived creditworthiness of issuers falling into a particular category or market segment. Such a view could lead to increases in the credit spread component of the discount factor used to value bonds in this credit band.

It is fair to assume that investors must pay for the privilege of owning the long call option feature of a CB. One way in which the cost of this privilege can be assessed is to calculate the extent to which the traded price of the convertible lies above the floor defined by the appropriate straight corporate bond. As an equation this can be shown as presented in equation (12.6):

$$\text{CB premium} = \frac{(\text{CB price} - \text{Corp. floor})}{\text{Corp. floor}}. \tag{12.6}$$

Thus if, at a point in time, the CB is trading at 102.60 and the bond floor is 84.84 the CB premium will be 20.93%, using equation (12.6). An important question poses itself at this point, namely: how is the convertible bond to be priced?

Number of periods	4
Interest rate (%)	4
Credit spread	1.5
Share price volatility (%)	20
Expiration date (months)	48
Current price of the underlying share	7
Conversion ratio	10.5
Conversion price	9.524
Coupon (annual)	0
Par value of bond (denomination)	100

Table 12.7: Convertible bond: binomial pricing

12.2.3 *Pricing a convertible bond*

On occasions in this text reference has been made to the use of trees for instrument pricing. In particular the binomial and trinomial pricing methods have been introduced. An intuitively appealing starting point when analysing the problem of pricing a CB would therefore be to make use of the binomial framework. Bearing in mind the caveats raised earlier concerning the use of the binomial and trinomial approaches to pricing, this section will make use of a four-period tree to illustrate CB pricing using this type of lattice approach. The hypothetical convertible bond to be analysed is described in Table 12.7.

Thus in this example a 4-year CB is being priced and analysed over the four-year period using one-year time slices. The current risk-free rate of interest is 4%, the bond's coupon is 0%, the par value of the bond is £100, on conversion each £100 bond can be converted into 9.524 shares of the company at a nominal price of £10.5 per share. The company's share price is currently standing at £7.00, it has a volatility of 20% p.a., and no dividends will be paid over the lifetime of the bond.

Digression: Pricing a bullet bond in a binomial framework

Figure 12.5 demonstrates how, in principle, a binomial tree can be used to arrive at the price of a coupon-bearing or a zero-coupon bond. Note that a binomial pricing framework is not really required in order to calculate the price of a straight corporate bond – unless trying to establish an appropriate implied credit spread using option adjusted spread analysis is the object of the exercise.[2] In the present example, however, it is assumed that the credit spread is known and is given as being 150 bps above the risk-free rate. Thus using 5.50% (4.00% + 1.50%) as the appropriate discount rate for each node and working backwards applying the p and $(1-p)$ parameter values at nodes *K, L, M, N, O* through to *A* the price of the bond is found to be 80.72.[3] This value represents

2 For a description of this process see Eales, B. A., *Financial Engineering*, Palgrave (2000).

3 To arrive at the figure of 94.79 which appears at nodes *G, H, I, J*, the end nodes *K, L, M, N, O* (100) are discounted by 1.055. Any coupon due at the end of year 3 is then added to the result. The same steps are followed for nodes *D, E, F*, etc. until the end node price of £80.72 per £100 nominal is obtained at the start node *A*.

the CB's floor. Since the CB in Table 12.7 is a zero-coupon bond this price could quite easily have been found using the standard pricing approach:

$$100/(1 + 0.055)^4 = 80.72$$

or, in the case of a coupon-bearing bond by using the pricing approach adopted when establishing the hypothetical floors which appear in Tables 12.4 to 12.6.

Using the methodology described in earlier chapters the up (u), down (d), and parameter values p and $(1 - p)$ are calculated as 1.2840, 0.7788, 0.4879, 0.5120, respectively. These parameters appear in Figure 12.6. Using the u and d values the progress of the company's share price can be calculated at each node of the tree.

Figure 12.6 illustrates a constant credit adjusted spread method of pricing the CB. The discount factor used here is the risk-free rate 5.50% (4.00% risk free rate + 1.50% credit spread) and is applied to each node in the tree. Starting again at the maturity end nodes (*K, L, M, N, O*), the procedure involves comparing the bond redemption value plus final coupon against parity with the larger of the two values being adopted as the price of the CB at that node. The parity values are found in the usual way by applying u (up) and d (down) values calculated according to the formulae presented in the discussion of the binomial methodology in Chapter 11. Here u equals 1.2214 and d equals 0.8187. The values £7.00 (node *A*), £8.55 (node *B*), £5.73 (node *C*), £10.44 (node *D*), £7.00 (node *E*), £4.69 (node *F*), etc. trace out the possible paths followed by the share price given the input data.

At each final node at the end of year 4 a comparison is made between the maturity value of the bond and the value associated with converting the bond into shares. In this case each bond can be converted into 10.5 of the company's shares so that at each individual node conversion

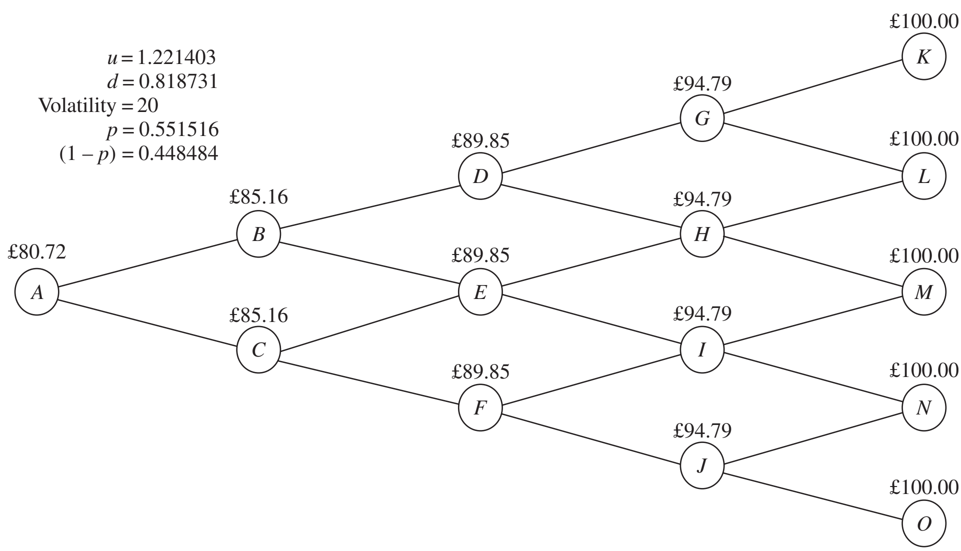

Figure 12.5: Bullet corporate bond

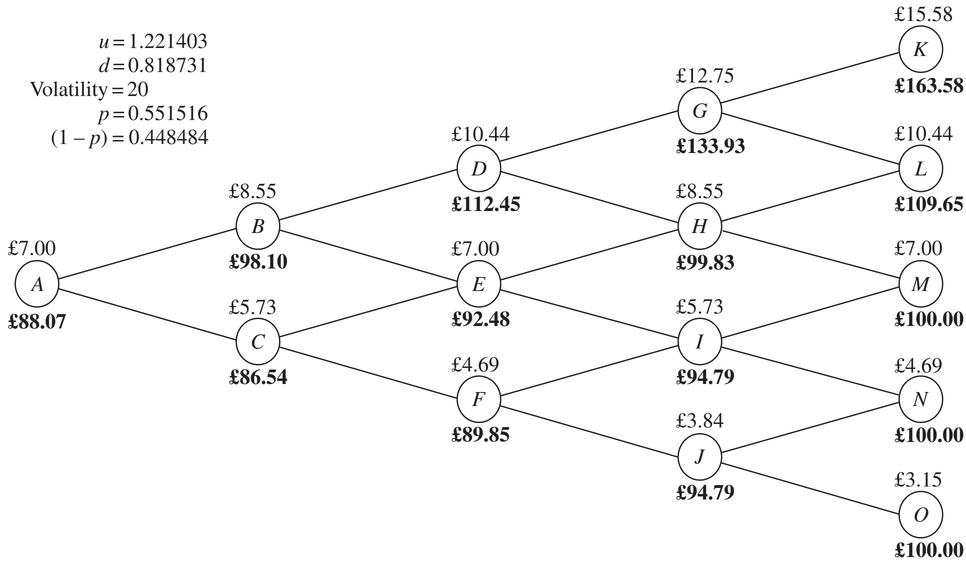

Figure 12.6: Convertible bond priced off constant credit spread

would result in a total value of final share price multiplied by 10.5. The larger of the two figures at each node is taken to be the end value for that node. To clarify this point consider node *K* where the final share price is calculated via the tree as £15.58. On multiplication this generates a figure of £163.58 while the bond's redemption value plus final coupon will be £100.

Assuming that a rational decision will be taken by an investor the figure adopted as the end node value will clearly be £163.58. In other words the bond will be converted.

The same will hold true for node *L*. Since £109.65 is greater than the £100 redemption value of the bond conversion will take place. At nodes *M*, *N* and *O* the bond will be redeemed for par value plus final coupon. At some points where redemption value plus coupon and terminal value of shares are very close the decision to convert may become more complex. Institutional and transaction costs, the dilution factor and investor preferences may come into play and influence the final decision.

Once these end node values have been established the familiar process of sliding back through the tree is used to arrive at values for nodes *G*, *H*, *I* and *J*, in the first instance. Applying the estimated parameters ($p = 0.5515$ and $(1 - p) = 0.4485$) to the values at nodes *K* (£163.58) and *L* (£109.65) a value of £133.93 is found. This value needs to be discounted at the credit adjusted rate of interest (5.50%), i.e. £133.93/1.055 = £132.12 to obtain the present value of the convertible bond. This value must now be compared to the value of the shares at that node: (10)(£12.75). From node *G* in Figure 12.6 an accurate value for the total value of shares is £133.92. Notice that no coupon needs to be added since the CB is a zero-coupon bond. The value of the CB at that point is thus parity £133.92 since it is greater than the backward calculated bond price of £132.12. Node *H* is found in a similar way but now taking £109.65 and £100.00 as results to be weighted and discounted. This yields £99.83 which compared to parity of (10.5)(8.55) = £89.77 would lead to a decision to hold on to the bond.

Repeating this process back to nodes B and C the values of £98.10 and £86.54, respectively, are obtained. By inserting these values in equation (12.6) a CB price of £88.07 is obtained.

$$CB\ price = \frac{[(98.10)(0.5515) + (86.54)(0.4485)]}{1.055} = 88.07. \tag{12.7}$$

Philips (1997) suggests, however, that this method of pricing a CB will generally lead to mispricing. He argues that when the company's share price is high and the CB is highly in-the-money the risk of default will be minimal and therefore a more appropriate discount rate to apply would be the risk-free rate with the basis point credit risk add-on being dropped. To facilitate the incorporation of this factor he proposes a delta-based technique which effectively adjusts the discount rates used at each node. The formula he advances appears as equation (12.8).

$$r^* = \Delta r_1 + (1 - \Delta) r_2 \tag{12.8}$$

where

- r^* represents the delta adjusted discount rate;
- r_1 represents the risk-free rate;
- r_2 represents the credit adjusted rate;
- Δ represents the (rate of change of the CB price/rate of change of parity).

The automatic question that then arises is how can the delta be calculated at each node. One approach would be to use the standard formula for calculating a hedge ratio (h) for written calls (puts) that was introduced in Chapter 11, namely:

$$h = \left(\frac{C_u - C_d}{S(u-d)}\right) \text{for calls or } h = \left(\frac{P_u - P_d}{S(u-d)}\right) \text{for puts.}$$

An alternative method would be to perturbate the starting price of the share by some small amount and find the small finite change that has occurred to the parity and bond figures in the binomial tree. This can be represented as follows in equation (12.9):

$$CB\ delta = \frac{\partial CB}{\partial Parity}. \tag{12.9}$$

Either approach allows the discount rates to be adjusted to reflect the in- or out-of-the-moneyness of the bond. In the example illustrated below using the method of tree perturbation and the data in Table 12.8 (which for convenience has been rounded to three decimal places) the delta at node A is found to be 0.533.

The change in the convertible bond's price when the price of one share is reduced by 0.01 to £6.99 is £88.015. This compares to the £88.071 established for a share price of £7.00 using the binomial tree. At the same time the change in parity is 0.105 [10.5(7.00–6.99)]. Delta can thus be calculated by subtracting the perturbed pass value of the bond from the originally calculated bond price and dividing this result by the change that has taken place in parity. Equation (12.9′) demonstrates this approach for node A:

$$\left[\frac{88.071 - 88.015}{73.5.0 - 73.395}\right] = \left[\frac{0.056}{0.105}\right] = 0.533. \tag{12.9′}$$

All other deltas can be calculated in the same way.

(a) First pass

Bond price	Parity	Bond price	Parity	Bond price	Parity	Bond price	Parity	Bond price	Parity
163.577	163.577	133.926	133.926	112.450	109.649	98.099	89.773	88.071	73.500
109.649	109.649	99.831	89.773	92.482	73.500	86.540	60.177		
100.000	73.500	94.787	60.177	89.845	49.269				
100.000	49.269	94.787	40.338						
100.000	33.026								

(b) Second pass

Bond price	Parity	Bond price	Parity	Bond price	Parity	Bond price	Parity	Bond price	Parity
163.344	163.344	133.734	133.734	112.315	109.492	98.011	89.645	88.015	73.395
109.492	109.492	99.749	89.645	92.439	73.395	86.517	60.091		
100.000	73.395	94.787	60.091	89.845	49.198				
100.000	49.198	94.787	40.280						
100.000	32.978								

Table 12.8: Bond price and parity changes

The deltas calculated using the perturbed tree method are reported in Table 12.9. Once the deltas for each node have been calculated the revised discount rates can be estimated. For node *K*, for example, the delta adjusted discount rate will be:

$$r^* = (1)(4.00\%) + (1-1)(5.50\%) = 4.00\% \tag{12.8$'$}$$

whereas for node *H* the delta adjusted rate will be:

$$r^* = (0.641)(4.00\%) + (1-0.641)(5.5\%) = 4.5385\%. \tag{12.8$''$}$$

Figure 12.7 demonstrates the impact that this approach will have on a CB's price. Using the delta adjusted rates as discount factors at each relevant node the price of the CB is found to be £90.61 compared to the £88.07 price obtained using more conventional methods. Of course these findings are based on a four-period tree which will not reflect the CB's true value, nevertheless when price differences exist in bi- or trinomial trees with many branches an arbitrage opportunity may well present itself. Since the premium calculated using the straightforward binomial model appears high compared to the adjusted discount approach the CB can be synthesised by buying an appropriate quantity of the company's equity. To complete the process sell the CB. In this way a position is engineered which will ensure an arbitrage profit.

Delta	Node	Delta	Node	Delta	Node	Delta	Node	Delta	Node
1.000	K	1.000	G	0.860	D	0.688	B	0.533	A
1.000	L	0.641	H	0.410	E	0.267	C		
0.000	M	0.000	I	0.000	F				
0.000	N	0.000	J						
0.000	O								

Table 12.9: Binomial tree deltas

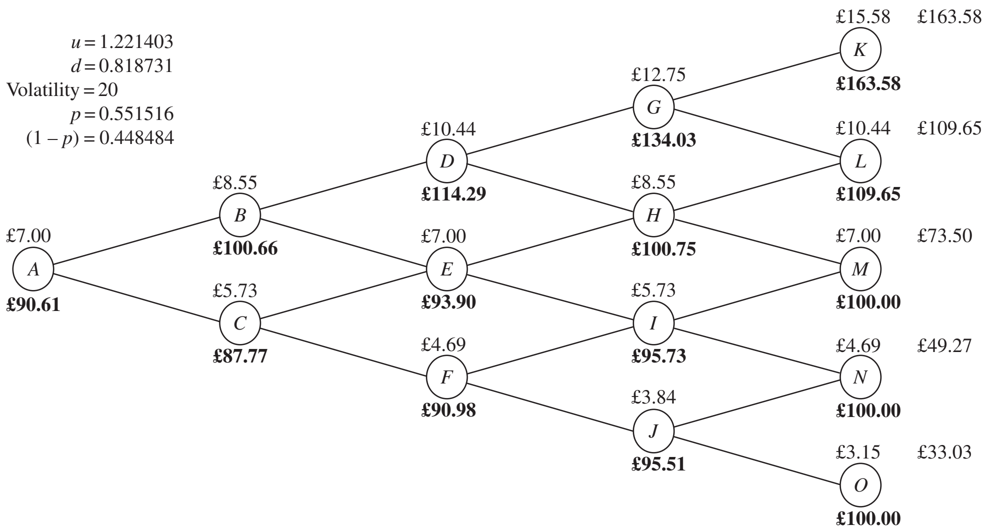

Figure 12.7: Convertible bond priced off constant credit adjusted rates

Turning to the trinomial pricing approach a similar pattern emerges. Figure 12.8(a) depicts a trinomial tree where the bond has been discounted at a constant 5.5% rate at each node. The price derived using this method is £87.90. Whereas using the approach suggested by Philips, and a lambda set arbitrarily at 3, Figure 12.8(b) reveals a price of £90.32. The

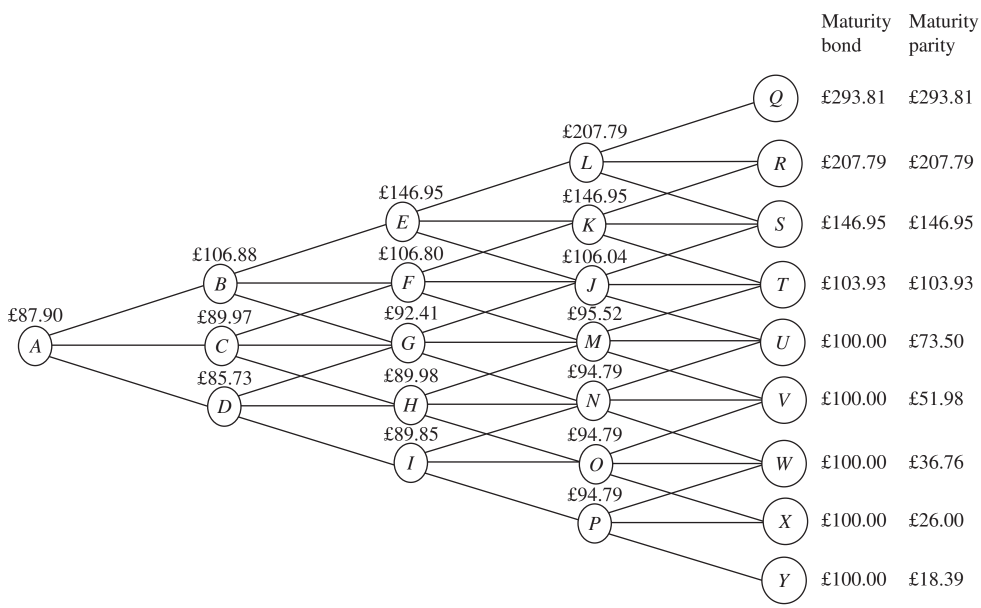

Figure 12.8(a): Convertible bond priced off constant credit spread

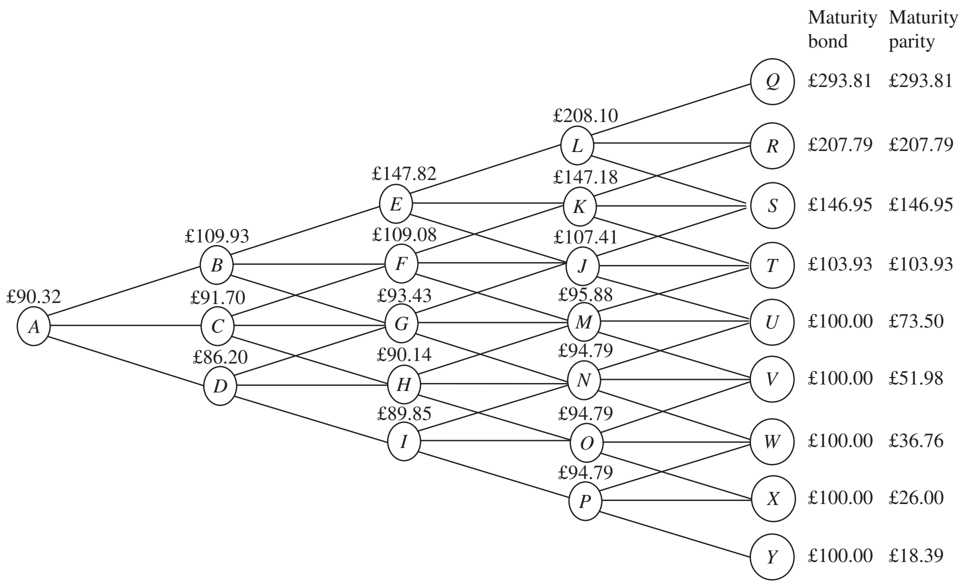

Figure 12.8(b): Convertible bond priced off delta adjusted interest rates

empirical deltas used to estimate this adjusted discount rate price and the adjusted rates themselves found using the perturbation method are shown in Tables 12.10 and 12.11, respectively.

The reader is reminded once again of the caveats raised in Chapter 11 concerning pricing using lattice frameworks. Here the results presented are for illustrative purposes. To gain an accurate price far more branches would need to be incorporated into the tree, and for the trinomial tree an appropriate value for lambda would need to be established.

12.2.4 *Additional CB features*

Another interesting modification that can be introduced is the provision for a downward reset, or refix, minimum conversion rate at some specified future date. This feature is incorporated into a CB as a means of forcing conversion: an additional amount of the company's stock will be created but the bond will not be redeemed for cash. An example of this type of refinement is provided by the Mitsui Trust and Banking 0.5% CB, 1 August 2007. Provision is made for this bond to be converted at ¥920 with a minimum conversion price of ¥500. Figure 12.9 depicts the general effect that a reset will have on the instrument's profile. It should be noted that the effect of a reset clause may very well give rise to negative convexity in the CB's payout profile prior to expiration. The reset feature resembles an option whose strike price is reset when predefined conditions have been satisfied, for example, ladder options, or in the case of specified future dates cliquet options described earlier. The diagram also provides a reminder that as stock price falls the bond will also decline and may trade at a discount to par.

From the point of view of the investor downward resets may be attractive if they anticipate share prices to perform badly over the medium term but to recover by the time that the bond reaches maturity or is called. The investor then enjoys acquisition of shares

Period 4 Bond	Parity	Period 3 Bond	Parity	Period 2 Bond	Parity	Period 1 Bond	Parity	Period 0 Bond	Parity
293.808	293.808	205.184	207.787	145.384	146.952	106.884	103.928	87.895	73.500
207.787	207.787	145.111	146.952	106.805	103.928	89.972	73.500		
146.952	146.952	106.045	103.928	92.414	73.500	85.728	51.981		
103.928	103.928	95.521	73.500	89.982	51.981				
100.000	73.500	94.787	51.981	89.845	36.762				
100.000	51.981	94.787	36.762						
100.000	36.762	94.787	25.999						
100.000	25.999								
100.000	18.387								
Period 4 Bond	**Parity**	**Period 3 Bond**	**Parity**	**Period 2 Bond**	**Parity**	**Period 1 Bond**	**Parity**	**Period 0 Bond**	**Parity**
293.388	293.388	204.891	207.491	145.179	146.742	106.759	103.779	87.838	73.395
207.491	207.491	144.903	146.742	106.678	103.779	89.921	73.395		
146.742	146.742	105.912	103.779	92.371	73.395	85.717	51.907		
103.779	103.779	95.493	73.395	89.977	51.907				
100.000	73.395	94.787	51.907	89.845	36.709				
100.000	51.907	94.787	36.709						
100.000	36.709	94.787	25.962						
100.000	25.962								
100.000	18.361								

Table 12.10: CB bond price and parity changes

Adjusted rate	Node	Adjusted rate	Node	Adjusted rate	Node	Adjusted rate	Node	Adjusted rate	Node
0.0400	Q	0.0402	J	0.0403	E	0.0424	B	0.0468	A
0.0400	R	0.0402	K	0.0422	F	0.0477	C		
0.0400	S	0.0416	L	0.0489	G	0.0527	D		
0.0400	T	0.0510	M	0.0540	H				
0.0550	U	0.0550	N	0.0550	I				
0.0550	V	0.0550	O						
0.0550	W	0.0550	P						
0.0550	X								
0.0550	Y								
Risk-free		0.04							
credit spread adjusted rate		0.055							

Table 12.11: Delta adjusted discount rates

at a cheaper rate but the issuer will need to issue more stock and will decrease the influence of existing equity holders. This is known as the dilution effect.

Upward resets are also issued and operate in the opposite way to the downward reset feature described above. This type of clause leads to the likelihood that the CB will be redeemed rather than converted but will also mitigate the dilution effect if conversion does take place and new shares are issued.

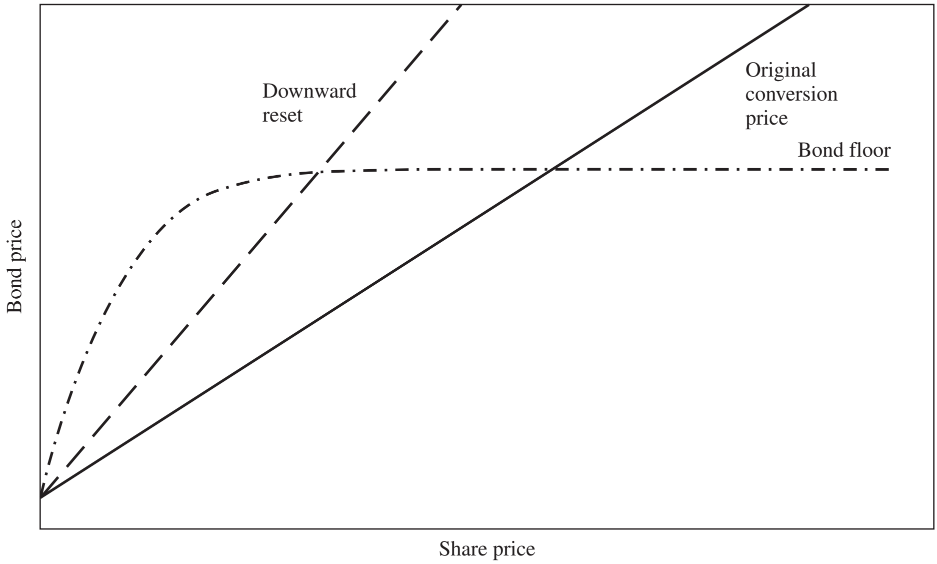

Figure 12.9: Reset convertibles

12.2.5 *Embedded calls*

Convertible bonds are often issued with a call feature that enables the issuer, if they so wish, to force the holder to convert when the company's share price reaches a certain level. Frequently the call provision is activated or triggered when the share is 30% above the conversion price. So with a conversion price of £10 the bond might be deemed callable when the price rises above £13 and stays at above that level for more than a specified period of time. Whether or not the bond is callable will depend upon any conditions which were specified in the bond's original indenture. To protect the holder of a CB from such an eventuality provisions are normally made in the indenture to the effect that the bond cannot be called in the first x years following issue. This stipulation affords some protection to the holder of the bond and is referred to as *hard non-call* or *hard call*. Once the x years have elapsed the bond moves into a *soft non-call* or *soft call* period. The bondholder now faces the risk that the bond may be redeemed once the call level defined in the indenture has been reached. To offer the holder protection against market "spikes" there may be an additional condition to the effect that the bond cannot be called unless it has traded at, or above, the defined trigger call level for a specified number of days. The protection afforded by the hard call period ensures that the bondholder should, in the majority of cases – ignoring the risk of default – receive a coupon advantage over the holder of equity since coupons are usually higher than dividends.

If all issuers decided to call their bonds once the conditions in the indenture had been met the effect would be to form a collar as demonstrated in Figure 12.10. In practice, for a variety of reasons, convertible bonds are not always called once the trigger has been activated. Figure 12.11, which ignores the possibility of a collapsing floor, illustrates a traded bond's potential profile before the bond reaches maturity. At low share prices the bond floor will provide the bond price. At slightly higher prices the CB price rises above the floor but once the share price has exceeded some critical level the CB price will be drawn in

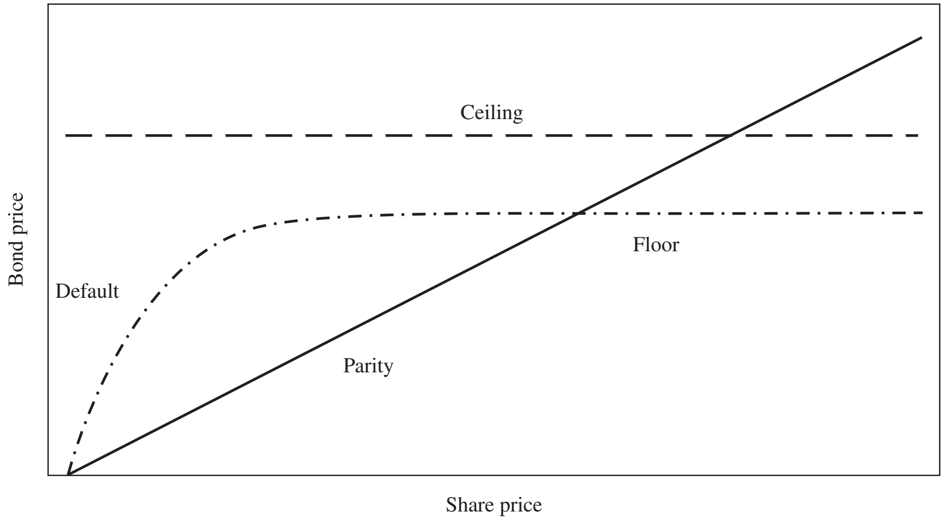

Figure 12.10: CB with embedded call option

towards the parity line. This is, of course, the typical behaviour of a standard call option prior to expiration.

Pricing a CB with an embedded call is relatively straightforward in a binomial framework. Take for example the CB if the bond houses a call feature, say, callable at end of year 3 and the soft call period has been entered. If the bond price at any node is greater than parity then the price of the bond must be reduced to parity. The logic behind this is that since the bond is callable and the call provision in the bond's indenture has been met, the bond may

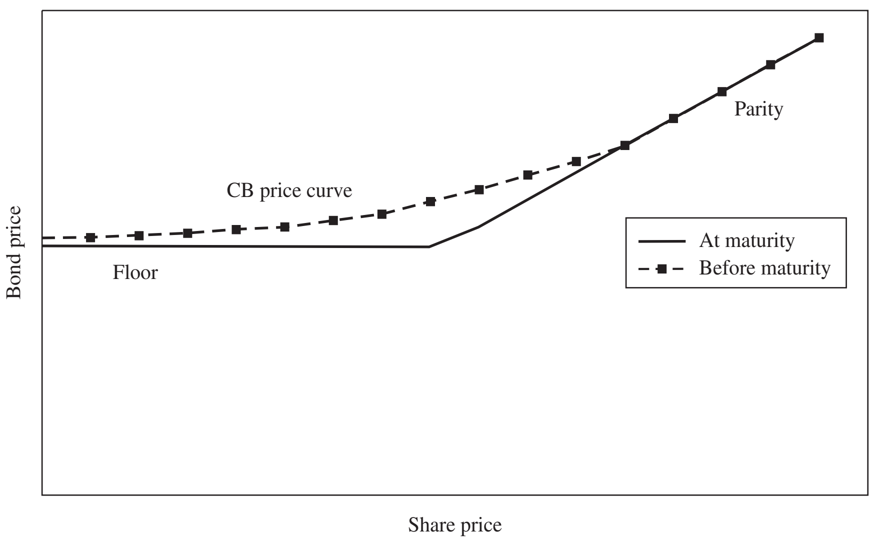

Figure 12.11: Convertible bond price behaviour

be withdrawn from the market by the issuer at any time. To clarify this consider the following example.

If parity is currently £142, say, and the bond price is estimated in the tree as £144, the bond should trade at £142. If it trades at £144 and is called, the holder will convert and the bond will cease to exist. Knowing this the potential purchaser of the bond should not be willing to pay an amount that would be received back in the form of shares if conversion is forced as a consequence of the issue of a call notice.

Applying this idea to the credit adjusted spread tree depicted in Figure 12.7, if the bond is callable at £105 from the end of year 2 when the price of the bond trades at £105 or higher, nodes *G* and *K* will need to be cut from the tree to establish today's fair price. Taking this change into account, suitably modifying the tree, and recalculating the deltas the revised convertible bond price can be calculated. The effect of the call is to reduce the CB's price to £88.03. In effect the bondholder has sold the call option and the bond's price is reduced by the call premium.

As discussed above, although convertible bonds may be callable there is no guarantee or compulsion on the part of the issuer to issue a call notice. Many issues trade well above their call trigger levels. In this event the bond price will continue to rise as the share price rises and it is parity that defines the cap on the price of the CB under those circumstances.

12.2.6 *Embedded puts*

This type of instrument conveys the right but not the obligation to return the CB to the issuer at a price defined in the indenture at a specified point in time. Clearly now the lower branches of the pricing tree become important in making the decision to exercise the put. However, lopping the lower CB values from the tree will have the effect of increasing the bond's market price. In this case the issuer is selling a put option to investors and that privilege will have to be paid for.

12.2.7 *Mandatory conversion structures*

Apart from the problem of pricing CBs the bond/equity mix embodied in a convertible leads to other difficulties. From an external perspective, how, for example, should a convertible be treated by the relevant regulatory authorities: is the instrument a bond generating a series of period cash flows which may be tax deductible or is it equity and part of the company's permanent capital base? How should financial market analysts regard the instrument when advising on investment strategies? From the company's internal perspective the big question concerns whether or not conversion takes place. To some extent the inclusion of reset clauses, discussed above, may go some way to ensuring that a CB is converted at maturity, but conversion is by no means certain; for this reason some convertible bonds have mandatory conversion stipulated in their contract specification.

There are several types of structures that fall into this category:

1. Equity commitment notes – where repayment is effected through a special fund established from moneys obtained through the sale of equity securities that qualify under accounting practices as primary capital.
2. Equity notes – where repayment is undertaken through either an optional fund created in the same way as described in 1 or by physically delivering equity from the stock of primary capital.

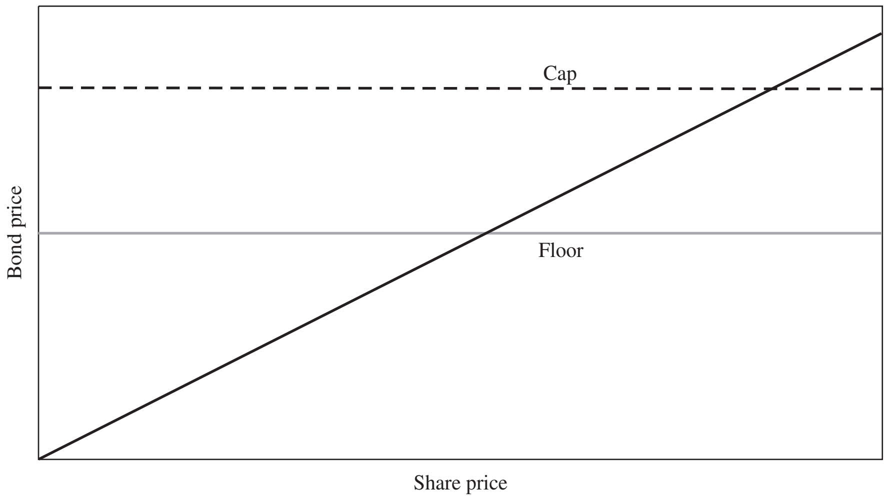

Figure 12.12: Convertible bond with cap and floor

Some mandatory conversion structures also have caps built into them. Typical in this category would be the PERCS (preferred equity redemption cumulative stock), ELKS (equity-linked debt securities) and EYES (equity-linked yield securities). The examples cited do not exhaust the mnemonics created by the financial engineering and structural product creators.

In the case of these mandatory conversion issues, the conversion rights section of the contract specification will be phrased to reflect the issuer's concern with potential dilution problems. The wording may take the form: "The Bond can be converted into the Issuer's shares subject to the defined conversion cap and floor. However, the Issuer reserves the right to cash settle the transaction." An example of a mandatory CB with capped share price of 132 and floor share price of 100 is illustrated in Figure 12.12.

The valuation of CBs with these complex structures can be handled by the tree methods already described. Another possibility, however, is to synthesise the product. In the case of a capped mandatory convertible this would entail the purchase of the underlying shares and, simultaneously, the writing of an OTM call option on the shares with an exercise price set at the level of the specified cap and with the same maturity date as the CB.

When a floor is included in the transaction this would also require the simultaneous purchase of a put option on the shares with an exercise price equal to the specified floor level and, again, with a horizon matching that of the maturity date of the CB.

12.2.8 *Reverse convertibles*

This type of CB is engineered so that at maturity it is either converted into shares if the underlying share price multiplied by the conversion factor is *less than* the redemption (par) value of the bond or redeemed at its par value plus final coupon. This decision rule is clearly the exact opposite of that used in the analysis of the CBs discussed in this chapter. It is thus the issuer of the bond who has the possibility of deciding at maturity which redemption strategy is the more economic.

This instrument creates a position which from an investor's perspective is the same as having purchased a fixed-coupon bond and having written a European-style put option in the issuing company's equity. The instruments do not really offer equity or bond fund managers the features of the normal convertible bond. From the point of view of equity participation the upside potential of a reverse CB is capped and, moreover, it offers no supporting floor on the downside. On the plus side these instruments will offer coupon rates higher than those enjoyed by similar quality paper.

12.3 Currency considerations

This chapter has attempted to demonstrate the basic construction and pricing of convertible bonds. The focal point has been on single currency CBs and some modifying features that have been introduced in recent years. One further aspect, which adds another complicating dimension to the pricing of these instruments, but may also enhance their attractiveness to issuers and/or end users, is currency. Convertible bonds are often issued with features which involve more than one currency. In the indenture, for example, the shares may be denominated in one currency but the bond redeemable in another. Under such a regime the bond's price will reflect not only movements in yields but also movements in the foreign exchange markets. It is clear that this type of instrument brings with it additional risks as well as potential gains. A good exposition of the way in which such convertibles can be priced appears in Connolly (1998).

12.3.1 *Convertible bond credit spreads*

The main concentration of this chapter has been on the equity and bond characteristics that the instrument offers. The time has now come to examine credit aspects. This can be achieved by breaking down the bond into its component parts and structuring an asset swap to handle each identified exposure.

Scenario 1

A 5.00% convertible bond with two years to maturity is purchased and held in a portfolio by a fixed income fund manager attracted by the possibility of gaining a relatively safe investment but with the good upside potential offered by the equity component of the bond. The fund manager realises at this point that the equity component has performed well since its purchase and he/she wishes to lock into the equity gains without converting.

This "locking-in" could be achieved by writing a call option and buying a put option with a strike price that reflects the current price of the underlying equity. If exchange-based options are available they could be used but in their absence OTC options would be a substitute. Alternatives to this would be to short a USF – if available – or use CFDs. Up to a point this effectively covers movements up or down in the value of the share price leaving the overall position flat. Figure 12.13 illustrates how this position is hedged. It should be noted that should the underlying share decline below the floor of 97.5 the short position will continue to generate profits that will augment those

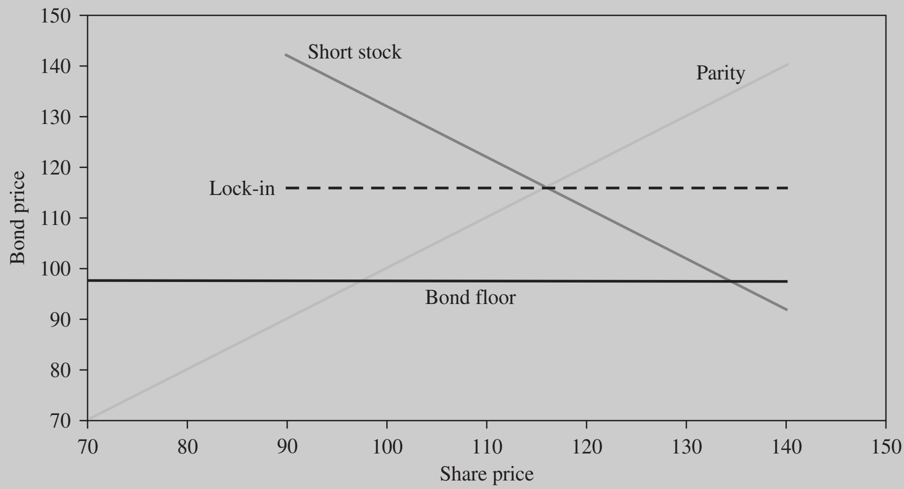

Figure 12.13: Equity lock-in

Outcome	Action
CB reaches maturity and is in-the-money	Closes the "short" position (if the market price of the equity is below the locked-in price the "short" position will deliver a profit. If the market price of the equity is above the locked-in price the "short" position delivers a loss). Converts CB at the conversion price and sells equity at market.
CB reaches maturity out-of-the-money	Closes the "short" position. Since the price of equity has fallen the "short" position will deliver a profit to ensure the locked-in value. Receives the bond's par redemption value.
CB is called in-the-money	Closes the "short" position at parity. Receives the call price from the issuer.
CB is called out-of-the-money	Closes the "short" position and receives a positive payoff to ensure locked-in price is achieved. Receives the call price from the issuer.
Issuer defaults on CB	Closes the "short" position. Receives a positive payoff to ensure locked-in price is achieved. Receives the recovery value from the bond issuer.

Table 12.12

obtained from the CB's floor. This, of course, raises the possibility of reducing the cost of the lock-in strategy by writing a put option with a strike in the region of 97.5. Ignoring this last possibility there are a number of possible outcomes to this strategy. These are summarised in Table 12.12.

Scenario 2

An investor has purchased a CB issued by a company with a sound trading background and whose bonds are regarded as investment quality. The credit spread on the CB appears wider than what would normally be reasonable for an issuer of this standing.

The investor writes a call option on the equity component of CB which will be activated when the company's share price reaches a certain level. At this point in the strategy the investor holds a synthetic callable bond. There are several outcomes possible and these are summarised in Table 12.13.

There are many other strategies available using CBs but those in Tables 12.12 and 12.13 provide a flavour of what is possible.

Outcome	**Action**
CB reaches maturity and is in-the-money	Call option is exercised by the holder. Investor converts into equity at the conversion rate and sells equity at market. The investor receives par and passes equity gain to the option holder.
CB reaches maturity out-of-the-money	Call option expires worthless. Receives indenture-specified redemption value from the issuer.
CB is called in-the-money	Call option is exercised by the holder. The investor receives indenture-specified call value and passes equity gain to the option holder.
CB is called out-of-the-money	Call option expires worthless. Investor receives indenture-specified call value from the issuer.
Issuer defaults on CB	Receives the recovery value from the bond issuer.

Table 12.13

12.4 Guaranteed equity products

Recent years have seen a variety of equity index-linked bonds arriving on the market. Typically these products offer an investor the opportunity to benefit from the receipt of the growth, average growth, or some proportion of that growth, on a specified equity index or set of equity indices over some defined time period or, in the event of the growth in the index being flat or negative, return of the investor's original capital investment plus a small consolation payment.

For such products the methodology used to determine the index's start and final levels is important. Determination of the final index level may take the form of an average value of the index based on daily closing prices of the index over a specified period of time. The start level may be captured by a single index figure on the date when the product comes into existence or by a series of index levels over several days at start-up.

These products clearly offer the individual investor, who may be risk averse and reluctant to enter the stock market, the opportunity of participating in a buoyant stock market without the normal downside risk associated with positions in shares. The position enjoyed by the individual investor is that of a long call option as depicted in Figure 12.14.

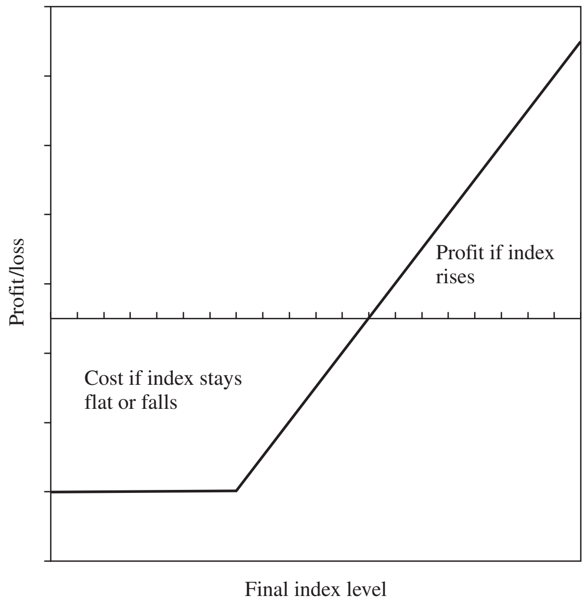

Figure 12.14: Investor's perspective

Needless to say the position taken by the individual investor is not totally risk free, there will be costs associated with holding this type of product. The most obvious cost is that of the interest forgone by purchasing a product which only pays out after a defined fixed period of time has elapsed.

To illustrate this consider an investor who ignores the index-linked bond and invests £10,000 in a fixed-term building society/bank deposit account for a five-year period at a fixed rate of 2% p.a. compounded annually. At the end of the five-year period the deposit account will have grown to approximately £11,040.

Now consider the investor's payout if, having chosen to invest in the guaranteed product, the benchmark index has not grown from its starting value or, indeed, has fallen in value when calculation of the final level takes place. Under this scenario if the consolation payment is 2% the total payout will be only £10,200, that is, £10,000 plus £200 interest over the entire investment period. So compared to the payout from a standard building society/bank deposit, a loss of £940 arises as a result of choosing the guaranteed product over the fixed-rate, fixed-term investment. This loss can be interpreted as the limited potential downside loss that results from holding the guaranteed product in a falling market as identified in Figure 12.14.

The index level breakeven point for this product can be found by taking the start-up and maturity index values, and comparing their ratio to the growth that would be obtained if the risk-free deposit account were to be used.

$$\left[\frac{FI}{SI}\right] = \frac{\text{Final risk-free deposit}}{\text{Original deposit}} \tag{12.10}$$

where

FI denotes the reference equity index value on maturity of the bond;

SI denotes the start-up index value.

If it is assumed that the index level at start-up is 5200 and that £10,000 grows to £11,040 after a five-year fixed-rate deposit then using equation (12.10):

$$\left[\frac{FI}{5200}\right] = \frac{11{,}040}{10{,}000} \tag{12.10$'$}$$

$$FI = (5200)(1.104) = 5740.$$

This indicates that the reference index will have to rise to around 5740 before the return on the equity index-linked investment outperforms a standard savings deposit. One thing that must be realised is that when full index participation is not offered by the issuing institution then the rise in the equity index will have to be greater. Equation (12.10) would need to be modified in order to reflect any reduced level of reference index participation. Equation (12.11) illustrates how this can be achieved.

$$\left[\frac{FI}{SI}\right] pt = \frac{\text{Final risk-free deposit}}{\text{Original deposit}} \tag{12.11}$$

where pt represents the reference index participation rate expressed as a decimal.

As an example, take the case where the issuer is offering 80.00% participation in the reference index growth ($pt = 0.8$). In this case the index will have to rise to 7176 before the guaranteed bond matches the bank deposit investment as shown in equation (12.11′).

$$FI = (5200)(1.104)/(0.8) = 7176. \tag{12.11$'$}$$

Other risks may also exist but they may well be less apparent than that described above. For instance, consider the situation where the equity index rises from year 1 to year 4 with only minor downturns. However, in year 5 the index falls dramatically from its peak in year 4. The index still finishes above it start-up value but substantially below the heights it had reached one year earlier. In this case the investor, unable to take advantage of the high index level in year 4, because of the fixed-term nature of the investment, has missed out on a potentially large payout. This type of situation is depicted graphically in Figure 12.15.

This situation might have been anticipated by the issuing institution, who could offer the investor the opportunity to use part of their capital investment to buy an option that would allow them to take advantage of high index levels which might prevail during the instrument's lifetime. This might be achieved by the institution offering to sell a cliquet or (ratchet) option to the investor. This would give the investor the opportunity to lock-in any reference index gains and reset the exercise price on pre-specified days during the bond's life. Alternatively the investor could be offered the opportunity of locking in gains when the reference index reaches pre-specified levels during the bond's lifetime – a ladder option. An example of how a cliquet option works is presented in Table 12.14 and is based on the following scenario.

An investor purchases a guaranteed bond which carries with it the opportunity to lock-in and reset the instrument's exercise price at the end of the second and fourth years of the bond's life.

Column 1 of Table 12.14 indicates the dates of particular interest over the bond's life. The second column defines the strike for measuring payouts at the bond's maturity. When

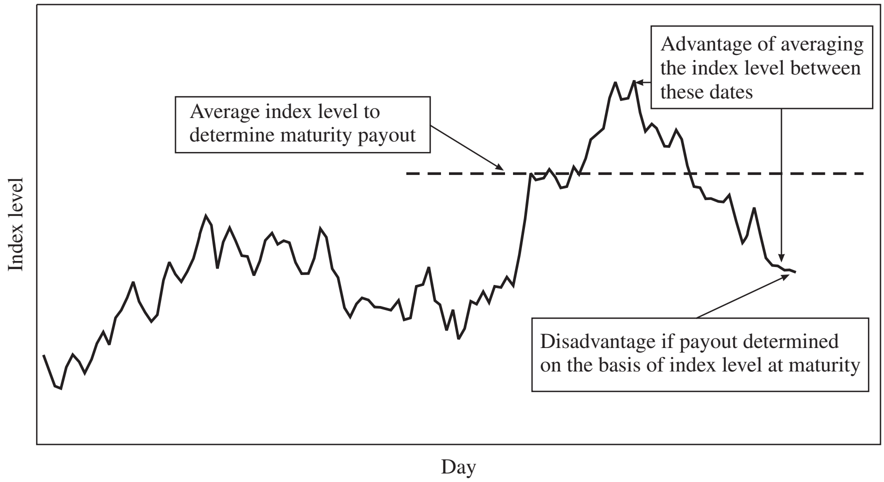

Figure 12.15: Index over time

Date	Index level	Strike set at	Cliquet payout	Regular option payout
Now	5200	5200	–	–
End Year 2	5750	5750	550	–
End Year 4	5500	5500	0	–
Maturity	5700	–	200	500

Table 12.14

the bond is issued the strike price of the reference index is set at 5200. At the end of year 2 the cliquet uses the standard payout rule to determine the option's value: max $(5750 - 5200, 0)$, the option is worth 550 index points and the strike is reset to the current index level of 5750. At the end of year 4 the reference index stands at 5500. Applying the standard decision rule again yields: $(5500 - 5750, 0)$ clearly the option in this period has no value but the strike is now reset to the index level of 5500. When the bond reaches maturity the option generates a further 200 index points for the bondholder. Over the five-year life of the bond the holder of the bond with an embedded cliquet option will have gained 750 index points. This compares to 500 points when no lock-in opportunities are provided.

It must be emphasised, however, that this attractive feature will have to be paid for. Recall that a fraction of the investor's capital will be used at start-up to buy the cliquet option. The eventual return will no longer be based on the £10,000 deposit but on (£10,000 − x) where x will be the front-end cost of the embedded cliquet option.

One more way in which the investor could be protected against over-reliance on the performance of the reference index on one particular day's value, would be to offer an average payout. Reverting to Figure 12.15, if the reference index were to be averaged over a period towards the end of the instrument's life the payout to the investor would have been more attractive.

Example 12.2

A company has just issued a special, limited issue bond. The bond offers the holder the opportunity of enjoying an equity index payoff over a five-year period with limited downside risk on the part of the investor.

The total amount of the bond issue involved is £20,000,000. Specifically the company has offered holders of the bond a five-year investment period at the end of which the investor will receive the greater of:

(a) return of original capital plus 70% of the growth in the index; or

(b) return of 100% of capital invested.

There are several ways in which such a guaranteed product might be hedged. The most popular way, and probably the most effective, is to use a zero-coupon bond plus ATM call options.

The data in Table 12.15 sets the scene. The reference equity index specified in the contract currently stands at 5200. It provides information on the index level at start-up, at the horizon, and rates of interest that need to be used in the analysis. The first scenario considered assumes that by the horizon the underlying index will have fallen by 3.0% to 5044.

Scenario 1: Index falls to 5044 at the investment horizon

Table 12.16 summarises the outcome if a zero-coupon bond is used to replicate the guaranteed product's potential payout. Assuming a discount rate of 3.0% an investment floor can be established that meets the return of capital by purchasing a zero-coupon bond. The cost involved in the establishment of the floor through the purchase of such a five-year zero-coupon bond for an investment of £20,000,000 is:

$$£20{,}000{,}000/(1.03)^5 = £17{,}252{,}176.$$

This will leave £2,747,824 available for the purchase of call options on the future level of the index. Using B&S **Basic.XLS** on the accompanying CD to price the option side of this transaction with input values set at volatility 14.0% p.a., exercise price and current index level 5200, risk-free rate of interest 4.25% p.a., dividend yield 1.1% p.a., and 1826 days to expiration, the option premium is found to be 9961 index points. If this

Principal	£20,000,000	
Horizon	5 year	
Discount rate	3%	
Interest rate for deposits	4.25%	
Index now	5200	
Index horizon	5044	Index growth = −3%
Participation	70.00%	
or	0.00%	

Table 12.15

Payout required at end of period		Principal +	£0		
Buy zero-coupon bond @		£17,252,176	Contracts required	269
Required to purchase options		£2,679,509	Strike	5200
Available to purchase options		£2,747,824		
Payout options	£0			
Zero	£20,000,000			
Interest on deposit	£15,805			
Excess funds	£68,315			
Total	£20,084,120			
Required	£20,000,000			
Profit/loss	£84,120			

Table 12.16

transaction is performed through an exchange where one index point is valued at £10, approximately 385 five-year ATM call contracts will be needed at a cost of approximately £3,834,985, if the full growth in the index is to be returned to the investor. However, since the contract only offers 70% of the reference index growth only 269 contracts will be needed and the cost will be reduced to £2,679,509. This leaves some £15,805 available to meet any contingencies which may arise.

The results for this scenario are reported in Table 12.16. Looking at the table, it is clear that, in the event of a market decline, the option component has zero payout and the zero-coupon bond pays out £20,000,000. These sums together with excess funds and accrued interest result in a profit of £84,120 in the replicating strategy. Thus the issuer has successfully hedged the position.

Scenario 2: Index rises by 156 points to 5356 by the horizon

The background data and assumptions for this case are presented in Table 12.17. In this scenario the rise in the index is compensated for by the option payouts at the end of the five-year period. The payout required to meet the guarantee is £420,000. The option payout alone is £419,640. When the excess funds and interest received are taken into account the profit figure rises to over £83,760, as shown in Table 12.18.

Principal	£20,000,000	
Horizon	5 year	
Discount rate	3%	
Interest rate for deposits	4.25%	
Index now	5200	
Index horizon	5356	Index growth = 3%
Participation	70.00%	
or	0.00%	

Table 12.17

Payout required at end of period		Principal +	£420,000		
Buy zero-coupon bond @		£17,252,176	Contracts required	269	
Required to purchase options		£2,679,509	Strike	5200	
Available to purchase options		£2,747,824			
Payout options	£419,640				
Zero	£20,000,000				
Interest on deposit	£15,805				
Excess funds	£68,315				
Total	£20,503,760				
Required	£20,420,000				
Profit/loss	£83,760				

Table 12.18

It must be emphasised that the zero-coupon bond has been priced off a low discount rate of 3%. Had the rate been higher the price of the zero-coupon bond, which creates the floor of the strategy, would have been cheaper and more funds would have been available to purchase call options. This would have given the issuer the chance of offering the investor more favourable index growth terms. Moreover the volatility off which the call options will be priced is also crucial. If interest rates are generally low and/or the volatility of call options high, strategies such as less than 100% index participation will be the order of the day. Alternatively the institution may offer a "blended" index growth bond where the return to the investor will be determined by the growth or decline across several indexes. An example of this might be: the investor will have a 100% principal guarantee plus the average growth from the FTSE 100®, the Nikkei 225, and the Dow Jones EuroSTOXX 50. If over the bond's life the reference indexes deliver 38%, 140% and −36% growth the payout to the investor will be approximately 47%. The payout is damped by the decline in the value of the EuroSTOXX index. There are, of course, many variations on this type of structure too.

This chapter has shown how equity-linked structures can be created, the uses to which they may be put and how, in some cases, they might be hedged. For a comprehensive description of these instruments the reader is referred to Satyajit Das's *Structured Products and Hybrid Securities*.

Selected bibliography and references

Connolly, K., *Priang Convertible Bonds*, Wiley 1998
Das, S., *Structured Products and Hybrid Securities*, 2nd edition, Wiley 2001
Eales, B. A., *Financial Engineering*, Palgrave 2000
Philips, G. A., *Convertible Bond Markets*, Macmillan 1997

Index

This Page Intentionally Left Blank