

WILEY FINANCE

Exotic options trading

FRANS DE WEERT

Exotic Options Trading

Frans de Weert



John Wiley & Sons, Ltd

Exotic Options Trading

For other titles in the Wiley Finance Series
please see www.wiley.com/finance

Exotic Options Trading

Frans de Weert



John Wiley & Sons, Ltd

Copyright © 2008

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

Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): cs-books@wiley.co.uk

Visit our Home Page on www.wiley.com

All Rights Reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP, UK, without the permission in writing of the Publisher.

Requests to the Publisher should be addressed to the Permissions Department, John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England, or emailed to permreq@wiley.co.uk, or faxed to (+44) 1243 770620.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The Publisher is not associated with any product or vendor mentioned in this book.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the Publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Other Wiley Editorial Offices

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA

Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany

John Wiley & Sons Australia Ltd, 42 McDougall Street, Milton, Queensland 4064, Australia

John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark, Singapore 129809

John Wiley & Sons Canada Ltd, 6045 Freemont Blvd, Mississauga, ONT, L5R 4J3, Canada

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

British Library Cataloguing in Publication Data

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

ISBN 978-0-470-51790-1 (HB)

Typeset in 11/13pt Times by Aptara Inc., New Delhi, India

Printed and bound in Great Britain by TJ International Ltd, Padstow, Cornwall, UK

This book is printed on acid-free paper responsibly manufactured from sustainable forestry in which at least two trees are planted for each one used for paper production.

Contents

Preface	xi
Acknowledgements	xiii
1 Introduction	1
2 Conventional Options, Forwards and Greeks	3
2.1 Call and Put Options and Forwards	3
2.2 Pricing Calls and Puts	6
2.3 Implied Volatility	8
2.4 Determining the Strike of the Forward	8
2.5 Pricing of Stock Options Including Dividends	9
2.6 Pricing Options in Terms of the Forward	10
2.7 Put-Call Parity	11
2.8 Delta	12
2.9 Dynamic Hedging	14
2.10 Gamma	14
2.11 Vega	16
2.12 Theta	18
2.13 Higher Order Derivatives Like Vanna and Vomma	19
2.14 Options' Interest Rate Exposure in Terms of Financing the Delta Hedge	21
3 Profit on Gamma and Relation to Theta	23
4 Delta Cash and Gamma Cash	25
4.1 Example: Delta and Gamma Cash	26

5	Skew	27
5.1	Reasons for Higher Realised Volatility in Falling Markets	27
5.2	Skew Through Time: ‘The Term Structure of Skew’	28
5.3	Skew and its Effect on Delta	29
5.4	Skew in FX versus Skew in Equity: ‘Smile versus Downward Sloping’	32
5.5	Pricing Options Using the Skew Curve	34
6	Simple Option Strategies	35
6.1	Call Spread	35
6.2	Put Spread	37
6.3	Collar	39
6.4	Straddle	40
6.5	Strangle	42
7	Monte Carlo Processes	45
7.1	Monte Carlo Process Principle	45
7.2	Binomial Tree versus Monte Carlo Process	46
7.3	Binomial Tree Example	46
7.4	The Workings of the Monte Carlo Process	48
8	Chooser Option	49
8.1	Pricing Example: Simple Chooser Option	49
8.2	Rationale Behind Chooser Option Strategies	51
9	Digital Options	53
9.1	Choosing the Strikes	55
9.2	The Call Spread as Proxy for the Digital	55
9.3	Width of the Call Spread versus Gearing	55
10	Barrier Options	57
10.1	Down-and-In Put Option	58
10.2	Delta Change over the Barrier for a Down-and-In Put Option	58
10.3	Factors Influencing the Magnitude of the Barrier Shift	60
10.4	Delta Impact of a Barrier Shift	63
10.5	Situations to Buy Shares in Case of a Barrier Breach of a Long Down-and-In Put	63

10.6	Up-and-Out Call	63
10.7	Up-and-Out Call Option with Rebate	64
10.8	Vega Exposure Up-and-Out Call Option	64
10.9	Up-and-Out Put	65
10.10	Barrier Parity	65
10.11	Barrier at Maturity Only	65
10.12	Skew and Barrier Options	66
10.13	Double Barriers	68
11	Forward Starting Options	71
11.1	Forward Starting and Regular Options Compared	71
11.2	Hedging the Skew Delta of the Forward Start Option	72
11.3	The Forward Start Option and the Skew Term Structure	73
11.4	Analytically Short Skew but Dynamically No Skew Exposure	74
11.5	Forward Starting Greeks	75
12	Ladder Options	77
12.1	Example: Ladder Option	77
12.2	Pricing the Ladder Option	78
13	Lookback Options	79
13.1	Pricing and Gamma Profile of Fixed Strike Lookback Options	79
13.2	Pricing and Risk of a Floating Strike Lookback Option	80
14	Cliquets	83
14.1	The Ratchet Option	83
14.2	Risks of a Ratchet Option	85
15	Reverse Convertibles	87
15.1	Example: Knock-in Reverse Convertible	87
15.2	Pricing the Knock-in Reverse Convertible	89
15.3	Market Conditions for Most Attractive Coupon	89
15.4	Hedging the Reverse Convertible	90
16	Autocallables	93
16.1	Example: Autocallable Reverse Convertible	93

16.2	Pricing the Autocallable	95
16.3	Autocallable Pricing without Conditional Coupon	97
16.4	Interest/Equity Correlation within the Autocallable	98
17	Callable and Puttable Reverse Convertibles	99
17.1	Pricing the Callable Reverse Convertible	99
17.2	Pricing the Puttable Reverse Convertible	102
18	Asian Options	105
18.1	Pricing the Geometric Asian Out Option	105
18.2	Pricing the Arithmetic Asian Out Option	107
18.3	Delta Hedging the Arithmetic Asian Out Option	109
18.4	Vega, Gamma and Theta of the Arithmetic Asian Out Option	110
18.5	Delta Hedging the Asian in Option	110
18.6	Asian in Forward	112
18.7	Pricing the Asian in Forward	114
18.8	Asian in Forward with Optional Early Termination	116
19	Quanto Options	119
19.1	Pricing and Correlation Risk of the Option	119
19.2	Hedging FX Exposure on the Quanto Option	122
20	Composite Options	125
20.1	An Example of the Composite Option	125
20.2	Hedging FX Exposure on the Composite Option	126
21	Outperformance Options	129
21.1	Example of an Outperformance Option	129
21.2	Outperformance Option Described as a Composite Option	130
21.3	Correlation Position of the Outperformance Option	131
21.4	Hedging of Outperformance Options	132
22	Best of and Worst of Options	135
22.1	Correlation Risk for the Best of Option	135
22.2	Correlation Risk for the Worst of Option	137
22.3	Hybrids	138

23	Variance Swaps	139
23.1	Variance Swap Payoff Example	140
23.2	Replicating the Variance Swap with Options	140
23.3	Greeks of the Variance Swap	142
23.4	Mystery of Gamma Without Delta	144
23.5	Realised Variance Volatility versus Standard Deviation	145
23.6	Event Risk of a Variance Swap versus a Single Option	146
23.7	Relation Between Vega Exposure and Variance Notional	147
23.8	Skew Delta	147
23.9	Vega Convexity	148
24	Dispersion	151
24.1	Pricing Basket Options	151
24.2	Basket Volatility Derived From its Constituents	152
24.3	Trading Dispersion	153
24.4	Quoting Dispersion in Terms of Correlation	153
24.5	Dispersion Means Trading a Combination of Volatility and Correlation	153
24.6	Ratio'd Vega Dispersion	155
24.7	Skew Delta Position Embedded in Dispersion	156
25	Engineering Financial Structures	157
25.1	Capital Guaranteed Products	157
25.2	Attractive Market Conditions for Capital Guaranteed Products	158
25.3	Exposure Products for the Cautious Equity Investor	160
25.4	Leveraged Products for the Risk Seeking Investor	163
Appendix A Variance of a Composite Option and Outperformance Option		167
Appendix B Replicating the Variance Swap		169
Bibliography		175
Index		177

Preface

This book is appropriate for people who want to get a good overview of exotic options in practice and are interested in the actual pricing of them. When dealing with exotic options it is very important to quantify the risks associated with them and at which stock or interest rate levels the Greeks change sign. Namely, it is usually the case that the Greeks of exotic options show much more erratic behaviour than the Greeks of regular options. This instability of the Greeks forces the trader to choose different hedging strategies than the standard option model would prescribe. Therefore the risk management of exotic options entails much more than just obeying the model, which in turn has an impact on the price. The non-standard risk management of exotic options means that when pricing an exotic option, one first needs to understand where the risks lie that affect the hedging strategy and hence the pricing of the particular exotic option. Once the risks have been mapped and the hedging strategy has been determined, the actual pricing is often nothing more than a Monte Carlo process. Moreover, when knowing the risks, the actual pricing of an exotic option can in some cases even be replicated by a set of standard options. In other words, the starting point for pricing exotic options is to have a full awareness of the risks, which in turn has an impact on how one needs to accurately price an exotic option.

The aim of this book is to give both option practitioners and economics students and interested individuals the necessary tools to understand exotic options and a manual that equips the reader to price and risk manage the most common and complicated exotic options. To achieve this it is imperative to understand the interaction between the different Greeks and how this, in combination with any hedging scheme, translates into a real tangible profit on an exotic option. For that reason, this book

is written such that for every exotic option the practical implications are explained and how these affect the price. Knowing this, the necessary mathematical derivations and tools are explained to give the reader a full understanding of every aspect of each exotic option. This balance is incredibly powerful and takes away a lot of the mystique surrounding exotic options, turning it into useable tools for dealing with exotic options in practice.

This book discusses each exotic option from four different angles. First, it makes clear why there is investor demand for a specific exotic option. Secondly, it explains where the risks lie for each exotic option and how this affects the actual pricing of the exotic option. Thirdly, it shows how to best hedge any vega or gamma exposure embedded in the exotic option. Lastly, for each exotic option the skew exposure is discussed separately. This is because any skew exposure is typically harder to quantify, but it has a tremendous impact on almost every exotic option. For that reason, this book devotes a separate section to skew, Chapter 5, which explains skew and the reasons for it in depth.

Acknowledgements

This book is based on knowledge acquired during my work as a trader at Barclays Capital. Therefore I would like to thank my colleagues at Barclays Capital who have been very helpful in teaching me the theory and practice of options. I would like to thank Faisal Khan and Thierry Lucas for giving me all the opportunities to succeed in mastering and practicing option trading; Stuart Barton, Karan Sood and Arturo Bignardi who have all made an indispensable contribution to my knowledge of options; Arturo Bignardi for his many suggestions and corrections when reviewing my work and Alex Boer for his mathematical insights; and special thanks go out to Neil Schofield for his many suggestions and for reviewing the work very thoroughly. I would also like to thank my great friends for all their support and motivation. Special thanks go out to my parents and my two sisters for having been a motivating force throughout my life and for a fantastic upbringing. I also want to thank my father for reviewing this book and being able to get an understanding of this complicated subject and subsequently being able to correct mistakes.

Petra, the light of my life. Finally, we can shine together!

Introduction

Exotic options are options for which payoffs at maturity cannot be replicated by a set of standard options. This is obviously a very broad definition and does not do justice to the full spectrum and complexity of exotic options. Typically, exotic options have a correlation component. Which means that their price depends on the correlation between two or more assets. To understand an exotic option one needs to know above all where the risks of this particular exotic option lie. In other words, for which spot price are the gamma and vega largest and at which point during the term of this option does it have the largest Greeks. Secondly, one needs to understand the dynamics of the risks. This means that one needs to know how the risks evolve over time and how these risks behave for a changing stock or basket price. The reason that one needs to understand the risks of an exotic option before actually pricing it is because the risks determine how an exotic option should be priced. Once it is known where the risks lie and the method for pricing it is determined, one finds that the actual pricing is typically nothing more than a Monte Carlo method. In other words, the price of an exotic option is generally based on simulating a large set of paths and subsequently dividing the sum of the payoffs by the total number of paths generated. The method for pricing an exotic option is very important as most exotic options can be priced by using a set of different exotic options and therefore saving a considerable amount of time. Also, sometimes one needs to conclude that the best way to price a specific exotic option is by estimating the price with a series of standard options, as this method better captures the risk involved with this exotic option. The digital option is a good example of that and will be discussed in Chapter 9.

Before any exotic option is discussed it is important to fully understand the interaction between gamma and theta. Although this book assumes an understanding of all the Greeks and how they interact, the following two sections give a brief summary of the Greeks and how the profit of an option depends on one of the Greeks, namely the gamma. A more detailed discussion of the Greeks and the profit related to them can be found in *An Introduction to Options Trading*, F. de Weert.

Conventional Options, Forwards and Greeks¹

This section is meant to give a quick run through of all the important aspects of options and to provide a sufficient theoretical grounding in regular options. This grounding enables the reader to enter into the more complex world of exotic options. Readers who already have a good working knowledge of conventional options, Greeks and forwards can skip this chapter. Nonetheless, even for more experienced option practitioners, this section can serve as a useful look-up guide for formulae of the different Greeks and more basic option characteristics.

2.1 CALL AND PUT OPTIONS AND FORWARDS

Call and put options on stocks have been traded on organised exchanges since 1973. However, options have been traded in one form or another for many more years. The most common types of options are the call option and the put option. A call option on a stock gives the buyer the right, **but not the obligation**, to buy a stock at a pre-specified price and at or before a pre-specified date. A put option gives the buyer the right, **but not the obligation**, to sell the stock at a pre-specified price and at or before a pre-specified date. The pre-specified price at which the option holder can buy in the case of a call and sell in the case of a put is called the *strike* price. The buyer is said to *exercise* his option when he uses his right to buy the underlying share in case of a call option and when he sells the underlying share in case of a put option. The date at or up to which the buyer is allowed to exercise his option is called the maturity date or expiration date. There are two different terms regarding the timing of the right to exercise an option. They are identified by a naming convention difference. The first type is the *European* option where the option can **only** be exercised at maturity. The second type of

¹ Parts of this chapter have been previously published in de Weert, F. (2006) *An Introduction to Options Trading*, John Wiley & Sons Ltd, Chichester. Reproduced with permission.

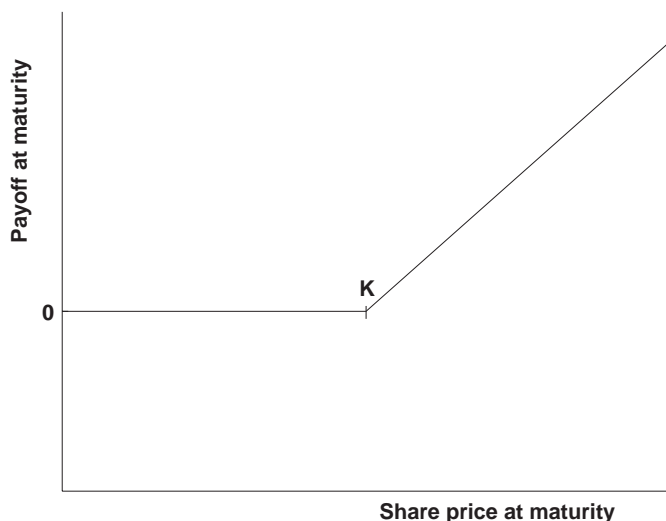


Figure 2.1 Payoff profile at maturity for a call option with strike price K

option is the *American* option where the option can be exercised at **any time up to and including** the expiry date.

Obviously, the buyer of a European call option would only exercise his right to buy the underlying stock if the share price was higher than the strike price. In this case, the buyer can buy the share for the pre-determined strike price by exercising the call and subsequently sell it in the market at the prevailing price in the market, which is higher than the strike price and therefore making a profit. The payoff profile of the call option is shown in Figure 2.1. The buyer of a European put option acts opposite to the buyer of the call option in the sense that the buyer of a put option would only exercise his option right, at maturity, if the share price was below the strike price. In this case the option buyer can first buy the share in the market at the prevailing market price and subsequently sell it at the strike price by exercising his put option, earning a profit as a result. The payoff profile at maturity of a put option is shown in Figure 2.2.

A forward is different to an option in the sense that the buyer of the forward is **obliged** to buy the stock at a pre-specified price and at a pre-specified date in the future. The pre-specified price of a forward is chosen in such a way that the price of the forward is zero at inception of the contract. Therefore, the *expected* fair value of the stock at a certain




Figure 2.2 Payoff profile at maturity for a put option with strike price K

maturity date is often referred to as the forward value of a stock or simply the forward associated with the specific maturity. The payoff profile at maturity of a forward contract is shown in Figure 2.3. Figure 2.3 makes clear that there is a downside in owning a forward. Whereas the owner of an option always has a payout at maturity which is larger




Figure 2.3 Payoff profile at maturity for a forward with strike price K which is equal to the fair forward value

than zero and therefore the maximum loss is equal to the premium paid for the option, the maximum loss on one forward is equal to the strike price of the forward, which occurs if the share price goes to zero. Since the definition of a forward prescribes that the contract is worth zero at inception, the strike price of the forward is equal to the forward value, which is discussed more elaborately in sub-section 2.4.

2.2 PRICING CALLS AND PUTS

In 1973 Black and Scholes introduced their famous Black–Scholes formula. The Black–Scholes formula makes it possible to price a call or a put option in terms of the following inputs:

- The underlying share price, S_t ;
- The strike price, K ;
- The time to maturity, $T - t$;
- The risk free interest rate associated with the specific term of the option, r ;
- The dividend yield during the term of the option, d ;
- The volatility of the underlying during the term of the option, σ .

The thought Black and Scholes had behind getting to a specific formula to price options is both genius and simple. The basic methodology was to create a risk neutral portfolio consisting of the option one wants to price and, because of its risk neutrality, the value of this portfolio should be yielding the risk free interest rate. Establishing the risk neutral portfolio containing the option one wants to price was again genius but simple, namely for a call option with a price c_t and underlying share price S_t

$$c_t - \frac{\partial c_t}{\partial S_t} \cdot S_t \quad (2.1)$$

and for a put option with price p_t

$$p_t - \frac{\partial p_t}{\partial S_t} \cdot S_t. \quad (2.2)$$

Respectively, $\frac{\partial c_t}{\partial S_t}$ and $\frac{\partial p_t}{\partial S_t}$ are nothing more than the derivatives of the call option price with respect to the underlying share price and the put option price with respect to the underlying share price. In other words, if, **at any time**, one holds $\frac{\partial c_t}{\partial S_t}$ number of shares against one call option, this portfolio is immune to share price movements as the speed at which the price of the call option changes with any given share price movement

is exactly $\frac{\partial c_t}{\partial S_t}$ times this share price movement. The same holds for the put option portfolio. Since both portfolios are immune to share price movements, if one assumes that all other variables remain unchanged, both portfolios should exactly yield the risk free interest rate. With this risk neutral portfolio as a starting point Black and Scholes were able to derive a pricing formula for both the call and the put option. It has to be said that the analysis and probability theory used to get to these pricing formulae are quite heavy. A separate book can be written on the derivations used to determine the actual pricing formulae. Two famous theorems in mathematics are of crucial importance to these derivations, namely *Girsanov's theorem* and *Ito's lemma*. It is far beyond the scope of this book to get into the mathematical details, but the interested reader could use Lamberton and Lapeyre, 1996 as a reference. Although the derivations of the call and the put price are of no use to working with options in practice, it is very useful to know the actual pricing formulae and be able to look them up when necessary. The prices, at time t , of a European call and put option with strike price K , time to maturity $T - t$, stock price S_t , interest rate r and volatility σ are given by the following formulae²:

$$c_t = S_t N(d_1) - K e^{-r(T-t)} N(d_2), \quad (2.3)$$

$$p_t = K e^{-r(T-t)} N(-d_2) - S_t N(-d_1). \quad (2.4)$$

In these formulae, $N(x)$ is the standard Normal distribution, and d_1, d_2 are defined as

$$d_1 = \frac{\ln\left(\frac{S_t}{K}\right) + \left(r + \frac{1}{2}\sigma^2\right)(T-t)}{\sigma\sqrt{T-t}}, \quad (2.5)$$

$$d_2 = d_1 - \sigma\sqrt{T-t}. \quad (2.6)$$

Equations 2.3 and 2.4 are incredibly powerful as all the variables that make up the formulae are known or can be treated as such, except for the volatility. Although it is not known what the interest rate will be over the term of the option, it can be estimated quite easily and on top of that there is a very liquid market for interest rates. Hence, the interest rate can be treated as known. Therefore, the only uncertainty left in the pricing of options is the volatility. The fact that the volatility over the term of the option is not known up front might make it impossible to

² The dividend yield is assumed to be zero.

price the option exactly, but it does give the opportunity to start betting or trading, especially on this variable.

2.3 IMPLIED VOLATILITY

Implied volatility is one of the main concepts in options trading. At first, the principle of implied volatility might seem quite difficult. However, it is in fact quite intuitive and simple. Implied volatility is the volatility implied by the market place. Since options have a value in the market place, one can derive the volatility implied by this price in the market by equating the Black–Scholes formula to the price in the market and subsequently solving for the volatility in this equation with one unknown. Except for the volatility, all variables in the Black–Scholes formula are known or can be treated as known and hence equating the Black–Scholes formula for an option to the specific price in the market of that option results in an equation in one unknown, namely the volatility. Solving for the volatility value in this equation is called the implied volatility. Although this is an equation in one unknown, one cannot solve this equation analytically but it can easily be solved numerically.

2.4 DETERMINING THE STRIKE OF THE FORWARD

Forwards are agreements to buy or sell shares at a future point in time without having to make a payment up front. Unlike an option, the buyer of a forward does not have an option at expiry. For example, the buyer of a 6 months forward in Royal Dutch/Shell commits himself to buying shares in Royal Dutch/Shell at a pre-agreed price determined by the forward contract. The natural question is, of course, what should this pre-agreed price be? Just like Black and Scholes did for the pricing of an option, the price is determined by how much it will cost to hedge the forward position. To show this, consider the following example. An investment bank sells a 2 year forward on Royal Dutch/Shell to an investor. Suppose that the stock is trading at €40, the interest rate is 5 % per year and after 1 year Royal Dutch/Shell will pay a dividend of €1. Because the bank sells the forward it commits itself to selling a Royal Dutch/Shell share in 2 years' time. The bank will hedge itself by buying a Royal Dutch/Shell share today. By buying a Royal Dutch/Shell share the bank pays €40, over which it will pay interest for the next 2 years. However, since the bank is long a Royal Dutch/Shell share, it will receive a dividend of €1 in 1 year's time. So over the first year the bank will pay interest over €40 and over the second year interest over

€ 39. This means that the price of the forward should be

$$F = 40 + 40 \times 0.05 + 39 \times 0.05 - 1 = 42.95. \quad (2.7)$$

A more general formula for the forward price is

$$F = \text{Price of underlying} + \text{Cost of carry}. \quad (2.8)$$

In the previous example the cost of carry is the interest the bank has to pay to hold the stock minus the dividend it receives for holding the stock.

2.5 PRICING OF STOCK OPTIONS INCLUDING DIVIDENDS

When dividends are known to be paid at specific points in time it is easy to adjust the Black–Scholes formula such that it gives the right option price. The only change one needs to make is to adjust the stock price. The reason for this is that a dividend payment will cause the stock price to go down by exactly the amount of the dividend. So, in order to get the right option price, one needs to subtract the present value of the dividends paid during the term of the option from the current stock price, which can then be plugged into the Black–Scholes formula (see equations 2.3 and 2.4). As an example, consider a 1 year call option on BMW with a strike price of € 40. Suppose BMW is currently trading at € 40, the interest rate is 5 %, the stock price volatility is 20 % per annum and there are two dividends in the next year, one of € 1 after 2 months and another of € 0.5 after 8 months. It is now possible to calculate the present value of the dividends and subtract it from the current stock level.

$$\begin{aligned} PV \text{ of Dividends} &= e^{\{-\frac{2}{12} \times 0.05\}} \times 1 + e^{\{-\frac{8}{12} \times 0.05\}} \times 0.5 \\ &= 1.4753. \end{aligned} \quad (2.9)$$

Now the option price can be calculated by plugging into the Black–Scholes formula (see equation 2.3) a stock price of $S_t = 40 - 1.4753 = 38.5247$ and using $K = 40$, $r = 0.05$, $\sigma = 0.2$ and $T - t = 1$.

$$d_1 = \frac{\ln\left(\frac{38.5247}{40}\right) + \left(0.05 + \frac{1}{2} \times 0.2^2\right) \times 1}{0.2 \times \sqrt{1}} = 0.1621 \quad (2.10)$$

$$d_2 = d_1 - 0.2 \times \sqrt{1} = -0.0379. \quad (2.11)$$

So, the price of the call option will be

$$\begin{aligned} c_t &= 38.5247 \times N(0.1621) - 40 \times e^{-0.05} N(-0.0379) \\ &= 3.2934. \end{aligned} \quad (2.12)$$

2.6 PRICING OPTIONS IN TERMS OF THE FORWARD

Instead of expressing the option price in terms of the current stock price, interest rate and expected dividend, it is more intuitive to price an option in terms of the forward, which comprises all these three components. The easiest way to rewrite the Black–Scholes formula in terms of the forward is to assume a dividend yield rather than dividends paid out at discrete points in time. This means that a continuous dividend payout is assumed. Although this is not what happens in practice, one can calculate the dividend yield in such a way that the present value of the dividend payments is equal to $S_t \times (e^{d(T-t)} - 1)$, where d is the dividend yield. So, if the dividend yield is assumed to be d and the interest rate is r , the forward at time t can be expressed as

$$\begin{aligned} F_t &= S_t \times e^{\{r(T-t)\}} \times e^{\{-d(T-t)\}} \\ &= S_t \times e^{\{(r-d)(T-t)\}}. \end{aligned} \quad (2.13)$$

From the above equation it is clear that dividends lower the price of the forward and interest rates increase it. As shown in the previous subsection, one can calculate the price of an option by substituting a stock price equal to $S_t \times e^{\{-d(T-t)\}}$ into the Black–Scholes formula. By doing this one can rearrange the Black–Scholes formula to express the price of an option in terms of the forward. The price of the call can then be expressed as

$$\begin{aligned} c_t &= S_t e^{-d(T-t)} N(d_1) - K e^{-r(T-t)} N(d_2) \\ &= e^{-r(T-t)} F_t N(d_1) - K e^{-r(T-t)} N(d_2). \end{aligned} \quad (2.14)$$

In the same way the price of the put can be expressed as

$$p_t = K e^{-r(T-t)} N(-d_2) - e^{-r(T-t)} F_t N(-d_1), \quad (2.15)$$

where d_1 and d_2 are

$$\begin{aligned} d_1 &= \frac{\ln\left(\frac{S_t e^{-d(T-t)}}{K}\right) + \left(r + \frac{1}{2}\sigma^2\right)(T-t)}{\sigma\sqrt{T-t}} \\ &= \frac{\ln\left(\frac{e^{-r(T-t)} F_t}{K}\right) + \left(\frac{1}{2}\sigma^2\right)(T-t)}{\sigma\sqrt{T-t}} \end{aligned} \quad (2.16)$$

$$d_2 = d_1 - \sigma\sqrt{T-t}. \quad (2.17)$$

2.7 PUT-CALL PARITY

The put–call parity is a very important formula and gives the relation between the **European** call and the put price, where the call and the put have the same strike and maturity, in terms of the share price and the strike price. The formula is as follows:

$$c_t - p_t = S_t e^{-d(T-t)} - K e^{-r(T-t)}, \quad (2.18)$$

where d is the dividend yield of the share and r is the risk free interest rate. One can prove this put–call parity by assuming it does not hold and show that one can then build a portfolio that leads to a riskless profit at maturity. For example, when

$$c_t - p_t > S_t e^{-d(T-t)} - K e^{-r(T-t)}, \quad (2.19)$$

one can prove that the portfolio

- sell a call option with strike price K and maturity $T - t$ (income is c_t);
- buy a put option with strike price K and maturity $T - t$ (income is $-p_t$);
- buy a share (income is $-S_t e^{-d(T-t)}$)

leads to a riskless profit at maturity for any level of share price at maturity. Indeed, if $S_T > K$ the call will be exercised and therefore the share within the portfolio is sold at $\text{€ } K$ and the put expires worthless. The net position at maturity is therefore zero, but since the generated income including the financing plus $\text{€ } K$ is larger than zero, the strategy makes a riskless profit. Indeed, rewriting equation 2.19 gives

$$[c_t - p_t - S_t e^{-d(T-t)}] e^{r(T-t)} + K > 0. \quad (2.20)$$

In the same way one can prove that the above portfolio leads to a riskless profit if $S_T < K$.

If one assumes that

$$c_t - p_t < S_t e^{-d(T-t)} - K e^{-r(T-t)}, \quad (2.21)$$

one can show that the following portfolio leads to a riskless profit for any level of share price at maturity:

- buy a call option with strike price K and maturity $T - t$ (income is $-c_t$);
- sell a put option with strike price K and maturity $T - t$ (income is p_t);
- sell a share (income is $S_t e^{-d(T-t)}$).

The put–call parity also signals for which strike the call and the put are worth the same. As one would expect, the strike for which the put and the call are worth the same is equal to the forward value of the share. Equation 2.13 established that the forward is equal to

$$F_t = S_t \times e^{\{r(T-t)\}} \times e^{\{-d(T-t)\}}. \quad (2.22)$$

Therefore, if $K = F_t$, equation 2.18 shows that the call is worth as much as the put.

2.8 DELTA

Delta is one of the most important Greeks and instrumental to the Black–Scholes derivation of the price of an option (see Section 2.2). Delta measures the sensitivity of an option price to the stock price. Mathematically, delta, δ , is the derivative of the option price with respect to the stock price. By taking the actual derivative of equations 2.3 and 2.4, the delta of European call and put options on a non-dividend paying stock are as given below. This assumes a call option price of c_t , a put option price of p_t and a stock price of S_t .

$$\delta_{\text{call, European}} = \frac{\partial c_t}{\partial S_t} = N(d_1) > 0, \quad (2.23)$$

$$\delta_{\text{put, European}} = \frac{\partial p_t}{\partial S_t} = -N(-d_1) < 0. \quad (2.24)$$

Equation 2.23 shows that the delta of a call option is between 0 and 1 and of a put option between -1 and 0. Table 2.1 shows for which stock price the delta reaches its extremes. The deltas of call and put options with strike K versus the stock price are shown graphically in Figures 2.4 and 2.5 respectively.

Table 2.1 Extreme delta values

Type of Option	Delta (δ)
Far in the money call option	1
Far out of the money call option	0
Far in the money put option	-1
Far out of the money put option	0




Figure 2.4 Variation of delta with stock price for a call option




Figure 2.5 Variation of delta with stock price for a put option

2.9 DYNAMIC HEDGING

Black and Scholes showed that a portfolio consisting, **at any time**, of one option and **minus** delta shares is risk neutral and should therefore yield the risk free interest rate. For a long call option, delta hedging means selling shares and for a long put option, **minus** delta shares means effectively buying shares as the delta of a put option is negative. To buy **minus** delta shares against one option is called delta hedging. In practice it is not possible to have minus delta shares against one option **at any time**. This would mean that one would have to continuously adjust the number of shares, which is not possible. Therefore, in practice an option is delta hedged at discrete points in time and is called dynamic hedging. Although this might seem inconvenient, it is exactly the reason that traders make money on volatility. This is discussed in Chapter 3.

2.10 GAMMA

Since delta changes whenever the stock price changes, it is useful to have a measure that captures this relationship. This measure is called gamma and gives the sensitivity of delta to a small change in stock price. Mathematically, gamma is the derivative of delta with respect to the stock price. Since delta is the derivative of the option price with respect to the stock price, gamma is the second order derivative of the option price with respect to the stock price. For European call and put options, gamma is given by the following formula.

$$\gamma_{\text{call, European}} = \frac{\partial \delta}{\partial S_t} = \frac{\partial^2 c_t}{\partial (S_t)^2} = \frac{N'(d_1)}{S_t \sigma \sqrt{T-t}} > 0, \quad (2.25)$$

$$\gamma_{\text{put, European}} = \frac{\partial \delta}{\partial S_t} = \frac{\partial^2 p_t}{\partial (S_t)^2} = \frac{N'(d_1)}{S_t \sigma \sqrt{T-t}} > 0. \quad (2.26)$$

where d_1 is defined as in equation 2.5 and

$$N'(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}. \quad (2.27)$$

The above formulae show that the gamma of a European call option is equal to the gamma of a European put option. Respectively, Figures 2.6 and 2.7 indicate the way in which gamma varies with the stock price and the time to maturity. It is helpful to interpret gamma in terms of how the delta hedge of an option changes for a change in stock price.




Figure 2.6 Variation of gamma with stock price

If gamma is positive, one needs to sell shares if the stock price goes up in order to be delta hedged and buy shares if the stock price goes down. If gamma is negative the reverse holds. Namely, one needs to buy shares if the stock price goes up in order to be delta hedged and sell shares if the stock price goes down.

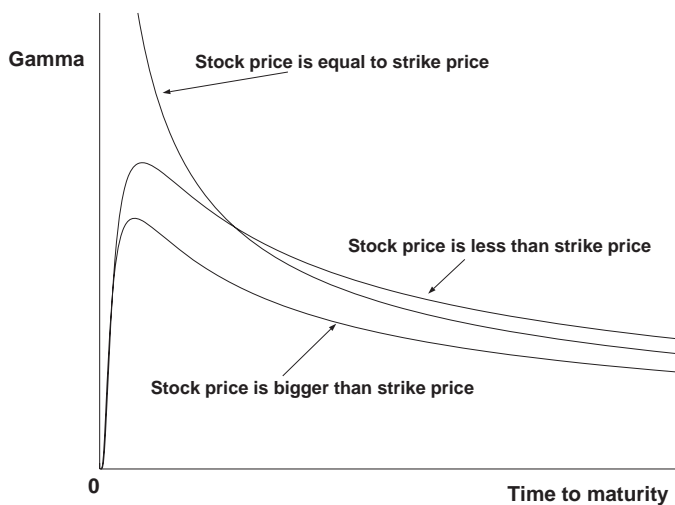


Figure 2.7 Variation of gamma with time to maturity

It is worth noticing that gamma becomes very large if an at the money option is close to expiring. This is caused by the fact that small stock price changes heavily affect the probability that this option will expire in the money. Since, for options close to expiration, $|\delta|$ is approximately equal to this probability,³ small stock price changes heavily affect δ . This explains why ‘at the money’ options close to maturity have large gammas.

Gamma is a very important Greek as it enables traders to derive the profit on an option for any given stock move. The precise way this is done is discussed in Chapter 3. For exotic options, gamma is usually one of the distinguishing features as it typically changes sign for different levels of stock price. Since gamma prescribes how to adjust one’s delta hedge, it is closely related to the actual profit of an option (see Chapter 3). Hence, for any exotic option, it is extremely important to identify the levels of stock price where the gamma changes sign.

2.11 VEGA

Vega measures the option price’s sensitivity to changes in volatility, or rather implied volatility. Mathematically, it is therefore the derivative of the option price with respect to the implied volatility, σ . The vegas, v , of a European call and put option are the same. Therefore, when trading the volatility of an underlying stock by trading options, it does not matter whether one trades a call or a put option. In formula form, the vegas, v , of a European call and put are expressed as:

$$v_{\text{call, European}} = \frac{\partial c_t}{\partial \sigma} = \frac{1}{\sqrt{2\pi}} e^{-\frac{d_1^2}{2}} S_t \sqrt{T-t} > 0, \quad (2.28)$$

$$v_{\text{put, European}} = \frac{\partial p_t}{\partial \sigma} = \frac{1}{\sqrt{2\pi}} e^{-\frac{d_1^2}{2}} S_t \sqrt{T-t} > 0. \quad (2.29)$$

The vega of a regular option is always positive, meaning that, if the implied volatility goes up, the option becomes more valuable. For exotic options, it is not always the case that the price goes up whenever the implied volatility does. An example is an up-and-out call, which is discussed in Chapter 10 in Section 10.8.

³ The probability that a call option expires in the money is equal to $N(d_2)$ and the probability that a put option expires in the money is equal to $N(-d_2)$. This explains why, for call and put options close to maturity, this probability is very close to $|\delta|$.




Figure 2.8 Variation of vega with stock price for a European call and put option

Respectively, Figures 2.8 and 2.9 give the variation of vega with respect to the stock price and the variation with respect to the time to maturity for a European option on a non-dividend paying stock. From these figures it is clear that vega is largest when the share is trading close to the strike price and the longer the term of the option the larger the




Figure 2.9 Variation of vega with time to maturity for a European call and put option

vega. So, in comparing gamma and vega it is good to remember that both these variables are largest when the share is trading close to the strike price, but when considering the time to maturity, gamma is large when the term of the option is short whereas vega is largest when the term of the option is long.

2.12 THETA

Theta measures the option price's sensitivity to the passage of time while all other variables remain unchanged. So, it is the rate of change of the option price with respect to time, and is usually indicated with the Greek letter θ . It is good to be aware of the impact of theta. Even if variables like stock price, interest rate and volatility remain unchanged, the option price will still change. Mathematically, theta is the derivative of the option price with respect to time. The theta of a European call option is always negative, which means that as time passes the option price decreases. The variation of theta with stock price for a European call option is plotted in Figure 2.10. The variation of theta with time to maturity for a European call option is plotted in Figure 2.11. The theta of a European put option is almost always negative. An example of a

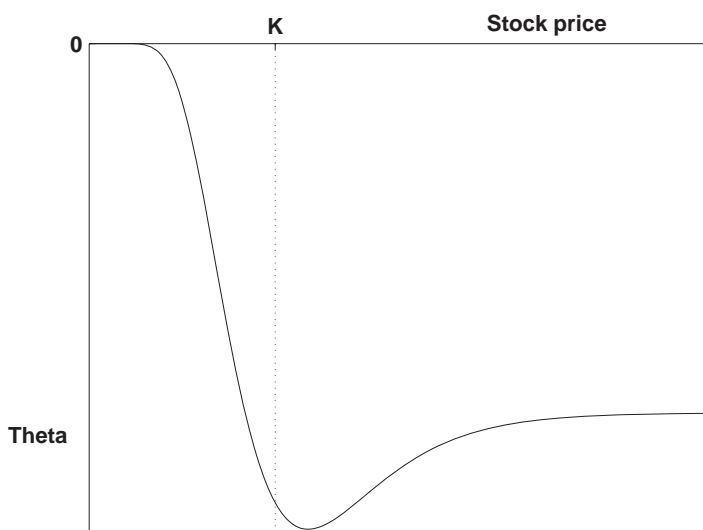


Figure 2.10 Variation of theta with stock price for a European call option, when the interest rate is strictly positive




Figure 2.11 Variation of theta with time to maturity for a European call option

European put option with a positive theta could be an ‘in the money’ European put option on a non-dividend paying stock, provided the interest rate is strictly positive. For a far ‘in the money’ put option this can be seen intuitively. After all, the fact that the upward risk is bounded is not worth anything for a far in the money put option. This together with an interest rate advantage ensures that as time passes, a far in the money European put option becomes more valuable. In Figure 2.12, the variation of theta with stock price is plotted for a European put option when the interest rate is strictly positive, and in Figure 2.13 when the interest rate is zero, which is the same as for the European call option.

2.13 HIGHER ORDER DERIVATIVES LIKE VANNA AND VOMMA

Vanna is the sensitivity of vega with respect to changes in the underlying and vomma is the sensitivity of vega with respect to changes in implied volatility. However, this book will limit the use of these two terms and other higher order derivatives but will rather specifically describe the feature at hand. In fact, higher order derivatives become quite meaningless for options where the Greeks themselves are already not smooth. In particular, exotic options have certain trigger points where the Greeks

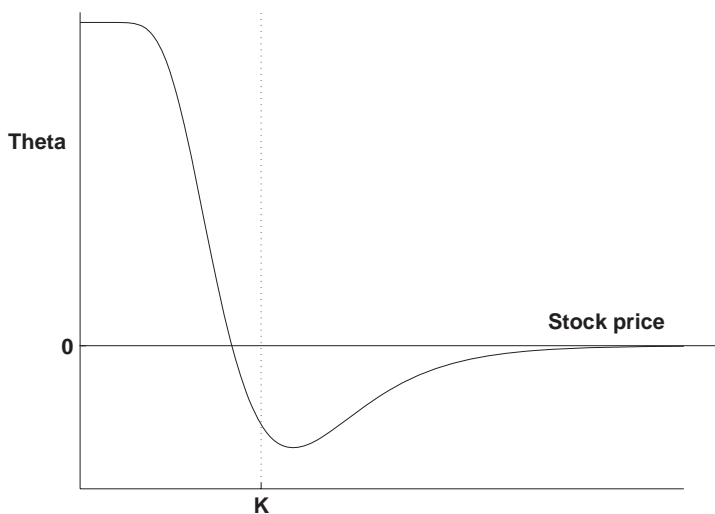


Figure 2.12 Variation of theta with stock price for a European put option, when the interest rate is strictly positive




Figure 2.13 Variation of theta with stock price for a European call and put option, when the rate is zero

'blow up' and therefore are not smooth at all. Only for variance swaps (Chapter 23) is it worth mentioning vanna and vomma, but for most exotic options these higher order derivatives show such erratic behaviour that they will only blur the picture.

2.14 OPTIONS' INTEREST RATE EXPOSURE IN TERMS OF FINANCING THE DELTA HEDGE

It is important to remember that an option is exposed to changes in interest rate. Mathematically this is nothing more than the derivative of the option price with respect to the interest rate. However, it is much more important to understand why an option is exposed to the interest rate. As usual in option theory, the answer can be found in how an option is hedged. For example, consider a call option. A call option is delta hedged by selling $|\delta|$ shares. By selling short these shares, the trader receives money which he can put in a savings account and subsequently earn interest on it. This means that the higher the interest rate the higher the call price. For the put option it is exactly the opposite. Indeed, a put option is delta hedged by buying $|\delta|$ shares. In order to do this the trader needs to borrow the funds which buy him $|\delta|$ shares. On these funds the trader incurs financing and therefore benefits from a lower interest rate. Typically, it is very useful to turn to the financing on the delta hedge in order to derive the interest rate exposure of an option.⁴ Doing this also provides a better understanding of the pricing of options in terms of real cash flows instead of abstract mathematical formulae.

⁴ In this approach the premium of the option is not taken into account. For the put option this premium increases the short interest rate exposure. However, for a call option the premium decreases the long interest rate exposure. Nonetheless, the net effect for the call option is a long interest rate exposure.

Profit on Gamma and Relation to Theta

The easiest way to show how the profit of an option can be determined by the gamma and the movements of the underlying is by means of a simple example. Suppose Philips is worth € 25 and a three month call option with a strike of € 25 has a delta of 0.50. Assuming that a trader buys 1000 of these call options he will need to sell 500 shares to be delta hedged. If the next day Philips' stock price goes to € 27 the delta will increase to, let's say, 0.70. This means that in order to be delta hedged the trader will sell an additional 200 shares at € 27. The next day Philips' stock price goes back down to € 25 and the corresponding delta of the option goes back to 0.50, which means that the trader will buy back 200 shares at € 25. Just from rebalancing the delta, the trader has made $2 \times 200 = € 400$. To put this in a general formula; if a stock moves by y after which the option is delta hedged and the stock then moves back to its original level, the profit on one option is:

$$y \times (\text{change in } \delta).$$

This means that the profit from a move y without moving back to its original level is

$$\frac{1}{2} \cdot y \times (\text{change in } \delta). \quad (3.1)$$

The change in δ can easily be rewritten in terms of gamma and is nothing more than $y \times \gamma$. This means that the profit on a long option after a move in the underlying of y is equal to

$$\frac{1}{2} \cdot \gamma \cdot y^2. \quad (3.2)$$

Obviously, the fact that the holder of an option makes money whenever the stock price moves does not come for free. The way the option holder pays for the right of buying low and selling high is by means of the theta, the time decay of an option. In other words, the holder of an option needs to earn back the daily loss in value of the option by taking advantage of

the moves of the underlying. Naturally, the reverse holds for the seller of an option. In this case the option seller makes money on the theta and loses it by rebalancing the delta, which effectively means buying high and selling low.

The relationship between gamma and theta is easy to understand now. The expected profit on gamma as a result of a move in a small time period should equal the loss in value of the option of that specific time period. The expected profit is easy to calculate as the expected move, y , of the stock in a small period of time, I , equals

$$y = \sigma S_t, \quad (3.3)$$

where σ has I as the unit of time and is the implied volatility of the option. This means that the expected profit over a small period of time, I , is

$$\frac{1}{2} \gamma \sigma^2 S_t^2. \quad (3.4)$$

As this expected profit should be equal to theta over the small time interval I , the above analysis has arrived at the main formula of option theory, namely:

$$\frac{1}{2} \gamma \sigma^2 S_t^2 + \theta = 0. \quad (3.5)$$

When plugging in the formulae for theta and gamma this can also be shown mathematically. A more detailed proof of the above is given in Chapter 7 of *An Introduction to Options Trading*, F. de Weert.

Delta Cash and Gamma Cash

Before discussing any exotic options it is important to understand the terminology of delta cash and gamma cash. Luckily these two definitions are easy to understand and very intuitive. Instead of delta and gamma, traders tend to use delta cash and gamma cash. This is for the simple reason that traders think in percentage movements rather than absolute movements. Indeed, both delta and gamma are absolute units and quantify the change for an absolute movement.

Delta cash, δ_c , is defined as delta, δ , times the share price S_t . Mathematically,

$$\delta_c = \delta \cdot S_t. \quad (4.1)$$

The profit on delta cash for a percentage move, $y\%$, is now nothing more than

$$\delta_c \cdot y\%.$$

Gamma cash, γ_c , is defined as the change in delta cash for a 1 % move in the underlying. Now one can derive that

$$\gamma_c = \frac{\gamma \cdot S_t^2}{100}. \quad (4.2)$$

Indeed, the change in delta for a 1 % move is

$$\gamma \cdot \frac{S_t}{100},$$

which means that the delta cash after a 1 % move is

$$\left[\delta + \gamma \frac{S_t}{100} \right] \cdot S_t.$$

Therefore the change in delta cash for a 1 % move, i.e. the gamma cash is

$$\left[\delta + \gamma \frac{S_t}{100} \right] \cdot S_t - \delta S_t = \frac{\gamma \cdot S_t^2}{100}.$$

4.1 EXAMPLE: DELTA AND GAMMA CASH

To make the use of delta and gamma cash more intuitive, consider the following example. A trader buys 100 thousand 1 year € 25 Philips call options. The delta of each € 25 Philips call option is 0.5. This means that the trader delta hedges himself by selling 50 thousand Philips shares. Assume also that the gamma cash of the 100 thousand Philips options amounts to 50 thousand euros per 1 %. This means that, if the trader is fully delta hedged initially and Philips' shares move down by 1 %, the trader needs to buy € 50 thousand worth of Philips shares to be fully delta hedged again. Indeed, when the share price moves down by 1 %, the delta cash of the portfolio comprising of 100 thousand € 25 call options plus a 50 thousand short share position goes from a delta cash equal to zero to a delta cash of minus € 50 thousand because of the gamma cash.

Skew is the principle that lower strike options on the same underlying have higher implied volatilities than options with higher strikes. The reason for skew can be explained by the observation that if markets go down they tend to become more volatile and, on very few occasions, the market actually crashes, in which case it will be incredibly volatile. This alone does not explain skew as this realised volatility is the same regardless of any strike price. The existence of skew is apparently saying that this increase in volatility has a bigger impact on lower strike options than on higher strike options. The reason behind this becomes apparent when thinking in terms of realised gamma losses as a result of rebalancing the delta of the option in order to be delta hedged. In a downward spiraling market the gamma on lower strike options increases, which combined with a higher realised volatility causes the option seller to rebalance his delta more frequently, resulting in higher losses for the option seller. Naturally, the option seller of lower strike options wants to get compensated for this and charges the option buyer by assigning a higher implied volatility to lower strikes. In other words; for any specific maturity, the lower the strike price of an option on any single underlying the higher the implied volatility of this option. This principle applies regardless of the in or out of the moneyness of an option. Whether it is a lower strike 'in the money' call option or a lower strike 'out of the money' put option makes no difference from a skew perspective. Indeed, if there were a difference there would be an arbitrage opportunity.

5.1 REASONS FOR HIGHER REALISED VOLATILITY IN FALLING MARKETS

An easy way to see why falling markets tend to become more volatile is to compare volatility to uncertainty. Markets go down because there is more uncertainty about the future and hence investors are more uncertain about how to value stocks. A price earnings ratio of 12 can be considered cheap one day and expensive the next day, just because of a different assumption on future interest rates. For example, at the end of a rate hike cycle there

is a lot of uncertainty about how to value equity markets as central banks need to control inflation by increasing rates but can easily overshoot, in which case the economy can end up in a recession. The reason a central bank could easily increase interest rates too much is because the effect of any former rate hikes can only be observed several months later and therefore results in uncertainty about the right course of action; keep increasing rates or acknowledge the end of the rate hike cycle.

An added feature of falling markets is that the number of short players will increase to take advantage of the declining market. This will significantly increase volatility as short players tend to work with tight stop losses, resulting in the so-called short squeezes, which cause big moves on the upside, after which the market typically falls again.

A natural question is of course why cheap valuations of stocks do not result in volatile markets on the upside. The answer to this question is that investors are by nature cautious and therefore any adjustment of too cheap a market will come slowly and is therefore unvolatile. An investor first wants to see confirmation of his view before he increases his positions, i.e. buys more stocks. However, for individual stocks where there is very much uncertainty about their future earnings, the stock might be volatile on the downside as well as on the upside. Therefore, individual stocks with a bullish sentiment can have very shallow skews or even reverse skews. However, the market as a whole will always have skew, i.e. options on indices will have skew built into them and will typically have steeper skews than options on stocks.

Yet another argument why stocks are more volatile on the downside than on the upside is that volatility goes hand in hand with fear of the investor. Fear in turn comes from investors being scared of losing money and therefore cutting positions. Since the market as a whole is long stocks, the fear of the total investor community will be largest when stock markets go down as this is when the investor community as a whole loses money and therefore will act in fear and irrationally, which results in higher volatilities.

5.2 SKEW THROUGH TIME: ‘THE TERM STRUCTURE OF SKEW’

A natural question is how skew behaves for different maturities. In other words, is the skew for long term maturities lower or higher than for short term maturities? The answer is that the skew for any specific stock is

higher (steeper) for short term maturities than for long term maturities.¹ To understand this, one first needs to make the observation that skew is mainly there because traders are afraid to lose money on downside strikes in case the market goes down and becomes more volatile. This means that skew is a result of traders wanting to be compensated on short downside options because rebalancing the delta on these options more frequently results in higher losses. Obviously, the larger the gamma the more imminent the problem. Since short term downside options have larger gammas when the share price moves down to the lower strikes, the effect of skew is largest for short term options. Another reason is that for short term maturities the trader exactly knows whether an option is a downside option or not. For long term options the trader cannot qualify whether it is a downside strike, as the trader does not know where the stock will be trading in, for example, two years' time. This means that a two year downside option currently, might be an upside option in two years' time if the markets go down. Two years from now is exactly when the effects of skew for this particular two year option are most pronounced, hence reducing the effect of skew for longer term options. Graphically, the term structure of skew looks like Figure 5.1, where the skew is quantified as the volatility difference per 10 % strike differential. For example, Figure 5.1 shows that the one year maturity has a skew of 2 volatility points differential between strikes that are 10 % apart.

5.3 SKEW AND ITS EFFECT ON DELTA

The skew curve of a particular stock can have a big impact on a trader's delta hedge against any option positions he has. Consider a trader who owns a 1 year 120 % call on BMW. Suppose the skew curve of BMW for the 1 year maturity indicates that every 10 % decrease in strike translates into a 1.5 % increase in implied volatility. This means that, if the 1 year ATM implied volatility of BMW is 20 %, the 1 year 120 % strike has an implied volatility of 17 %. Suppose that the trader decides to mark this 120 % call on an implied volatility of 21 %, i.e. higher than the ATM implied volatility. As a result the implied volatility curve plotted against strike takes the shape of a smile. Because the trader marks the 120 % call on a 21 % implied volatility instead of a 17 % implied volatility, his

¹ Only in very special situations is the skew for long term maturities higher than for short term maturities. For example, when a stock is expected to go bankrupt but people know it will not be in the next few months but any time from 6 months to 2 years.

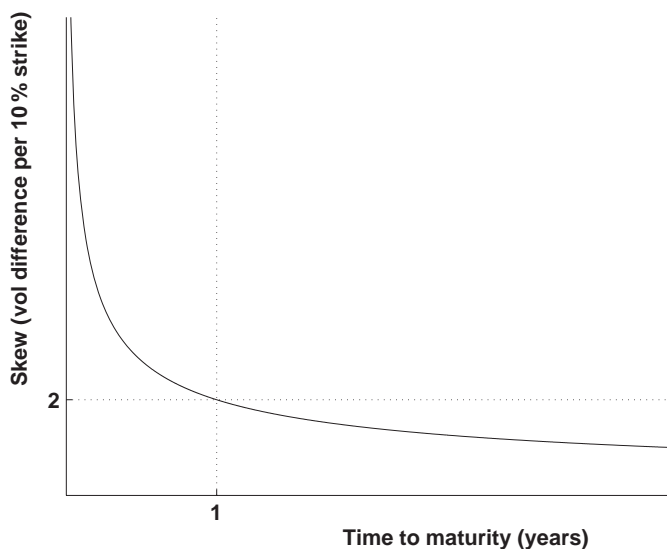


Figure 5.1 Skew plotted per time to maturity, clearly showing that the shorter the time to maturity the steeper the skew. Skew is expressed as the volatility differential per 10 % difference in strike

delta is larger than it should be and therefore sells more BMW shares against the 120 % call. This can be seen intuitively by comparing delta to the probability that an option expires in the money.² By increasing the implied volatility the probability that the 120 % call expires in the money increases and hence the delta increases. This shows that skew and the way a trader marks his volatility surface have a large impact on a trader's delta hedge. A good question is whether one would expect a downward sloping implied volatility curve per maturity or a smile curve, i.e. both higher and lower strikes than ATM have a larger implied volatility than the ATM implied volatility. The answer to this question is that it depends on the asset class. For options on equities one can say that the implied volatility curve per maturity is always downward sloping, except for very special situations. This can be seen by looking at the way a trader delta hedges his options for upside strikes and by comparing the delta to the probability that the option expires in the money.

² The probability that a call option expires in the money is $N(d_2)$ and the delta of a call option is $N(d_1)$, where $d_2 = d_1 - \sigma\sqrt{T-t}$. Since d_1 is also a function of σ it is the case that for upside strikes a higher σ results in a greater d_2 and therefore a higher probability that the call expires in the money. This is also what one would expect intuitively.

Assuming that a trader marks his upside strikes on a higher implied volatility than the ATM implied volatility means that he sells more shares against these long upside calls than when the skew curve has a downward sloping shape. This also means that, considering that the ATM implied volatility for this particular maturity is a trader's best estimate for what the stock will realise (volatility wise) during the term of the option, the trader assigns too high a probability to the option expiring in the money and therefore hedges the upside call on too high a delta. To make matters worse, the trader will continue to sell too many shares with the share price going up. However, with the share price going up the stock is likely to realise even less, making the probability and therefore the proper delta of this option even lower. In other words, by marking upside strikes on a higher implied volatility than the ATM implied volatility, the trader sells too many shares and, especially when the share price moves up, the trader not only loses money on the extra shares he sold, he also continues to sell too many shares. It is these dynamics that force a trader to mark his volatility surface per maturity with a skew that has a downward sloping shape rather than a smile.

There is one more parameter that is used to mark an implied volatility surface, namely the curvature. Indeed, it is the case that for very high strikes the implied volatility does not decrease any longer but flattens out, which can be marked in by using a curvature measure. At the same time this curvature parameter ensures that very low strikes have an even higher implied volatility than the skew parameter indicates. This is exactly what a trader would want, as very low strikes have very little premium and therefore sellers want to get properly compensated, which means that the skew parameter alone will not be enough and the curvature parameter will ensure these low strikes are marked on a larger implied volatility. In Figure 5.2 the implied volatility is plotted against strike, clearly showing the effects of both the skew parameter and the curvature parameter on the shape of the implied volatility surface for a specific maturity against strike price.

Without confusing matters too much, even in equity there is a small element of a smile shape for the implied volatility per strike. However, this is almost negligible as it is only concerned with upside strikes for which the options have a very small premium. Obviously, nobody will sell an option for nothing and therefore a very high strike might get an incredibly large implied volatility for the simple reason that a small premium already translates into a large implied volatility.




Figure 5.2 Implied volatility plotted against strike price for a specific maturity, pointing out how the skew and the curvature parameters affect the implied volatility surface to best fit the real world

5.4 SKEW IN FX VERSUS SKEW IN EQUITY: ‘SMILE VERSUS DOWNWARD SLOPING’

In sub-section 5.3 it was shown that the implied volatility plotted against strike has a downward sloping shape per maturity in the equity asset class. In Foreign exchange (FX), the implied volatility plotted against strike takes the shape of a smile for any specific maturity. In other words, the implied volatility for options on an exchange rate is higher for both lower strikes and higher strikes than ‘at the money’. This is because in FX either currency of the exchange rate can collapse. For example, taking the exchange rate of EUR versus USD, which is currently trading at € 0.8 per \$ 1, it is clear that if the exchange rate goes to € 0.4 per \$ 1, the dollar is crashing against the euro and an exchange rate of € 1.2 per \$ 1 means that the euro is crashing against the dollar. In both scenarios one can expect more uncertainty and therefore more volatility. For this reason a trader would charge a higher implied volatility for both the 0.4 strike and the 1.2 strike as it is the risk of crashing that causes traders to charge higher implied volatilities. Figure 5.3 plots the implied volatility against the strike for the EUR/USD exchange rate, clearly showing the ‘smile’ shape of the implied volatility surface in FX. Obviously,

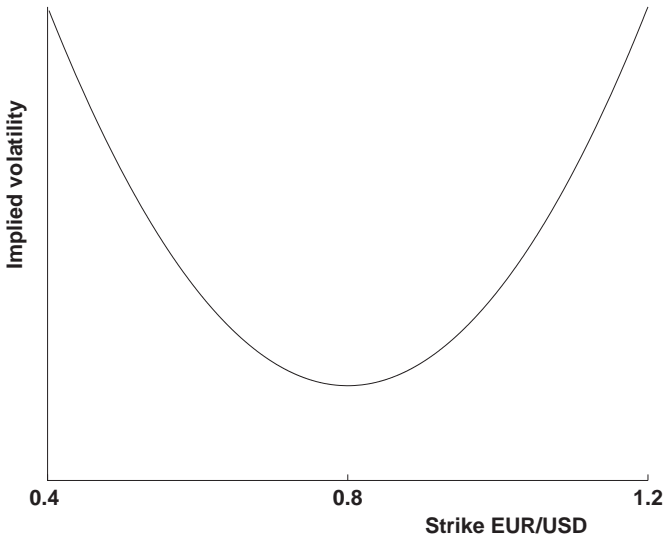


Figure 5.3 Implied volatility plotted against strike price for the EUR/USD exchange rate, pointing out the ‘smile’ shape of the volatility surface in FX

the extent of ‘skew’ in either direction depends on which currency is considered to be more stable. For example, both the euro and the dollar are considered stable currencies and therefore the ‘skew’ for both lower and higher strikes is very similar. However, the magnitude of ‘skew’ on an exchange rate between a stable currency and an unstable currency can be very different for downside strikes than for upside strikes. For example, consider the exchange rate between the euro and the Brazilian real (BRL), which is currently trading at BRL 2.6 per € 1. In this case it is very unlikely to see a crash of the euro versus the BRL, whereas a crash of the BRL versus the euro is highly possible. In other words, a trader would charge a much higher implied volatility for an FX option with a strike of BRL 3.6 per € 1 than for an FX option of BRL 1.6 per € 1. The reason being that an exchange rate of BRL 3.6 per € 1 means a large devaluation of the Brazilian real, which will be accompanied with large volatility, whereas an exchange rate of BRL 1.6 per € 1 is a large devaluation of the euro and will most likely be accompanied with a decrease in volatility as it means that the BRL is becoming more stable. Figure 5.4 plots the implied volatility against strike for the BRL/EUR exchange rate, and clearly shows that the skew is much more prominent for higher strikes, i.e. a devaluation of the BRL, than for lower strikes. It

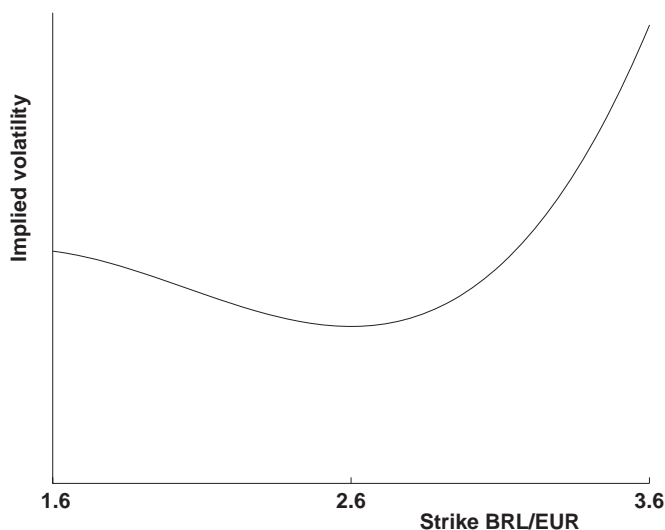


Figure 5.4 Implied volatility plotted against strike price for the BRL/EUR exchange rate, pointing out the biased ‘smile’ shape of the implied volatility surface between a stable and an unstable currency

could even be argued that the implied volatility should initially go down for slightly lower strikes than ‘at the money’, after which it picks up for very low strikes.

5.5 PRICING OPTIONS USING THE SKEW CURVE

When pricing options it is important to take the skew curve into account. That is to say, pricing an option means pricing it off the full volatility surface. Therefore, if one prices, for example, a downside option, one would look at the volatility curve and price it on a higher implied volatility than an ‘at the money’ option. For a plain vanilla downside option the effect of skew on the price is straightforward. However, for a large class of exotic options the effect of skew on the price is not that obvious. Hence, each section discusses separately the effect of skew on the price of the exotic option in question.

Simple Option Strategies¹

Before really delving into everything that has to do with exotic options, it is good to be familiar with some simple but commonly used option strategies. It is also a prerequisite to know the jargon of the different option strategies when dealing with exotic options. Especially because simple option strategies are often used to hedge the risks of an exotic option.

The most popular option strategies amongst investors usually involve more than just buying or selling an option outright. In this chapter several different option strategies are explained. Why investors would execute these option strategies is discussed from a break-even point of view. As a bonus, some of the following option strategies have an embedded skew position and can therefore serve as a simple introduction to more complex skew positions.

6.1 CALL SPREAD

One of the most popular strategies is the call spread. A call spread involves nothing more than two calls, one with a low strike and another one with a higher strike. An investor is said to be buying the call spread if he buys the lower strike and sells the higher strike. Selling the call spread means selling the lower strike and buying the higher one. Buying the call spread is called a bullish strategy because the investor benefits if the underlying increases in value. However, the investor's profits are capped because he has sold another call with a higher strike to fund his bullish view on the stock. As an example consider an investor who buys an at the money call on BMW and partially funds this by selling a 120 % call on BMW. In this case the investor profits from an increase in BMW's share price up to the point where the share price has reached

¹ Parts of this chapter have been previously published in de Weert, F. (2006) *An Introduction to Options Trading*, John Wiley & Sons Ltd, Chichester. Reproduced with permission.

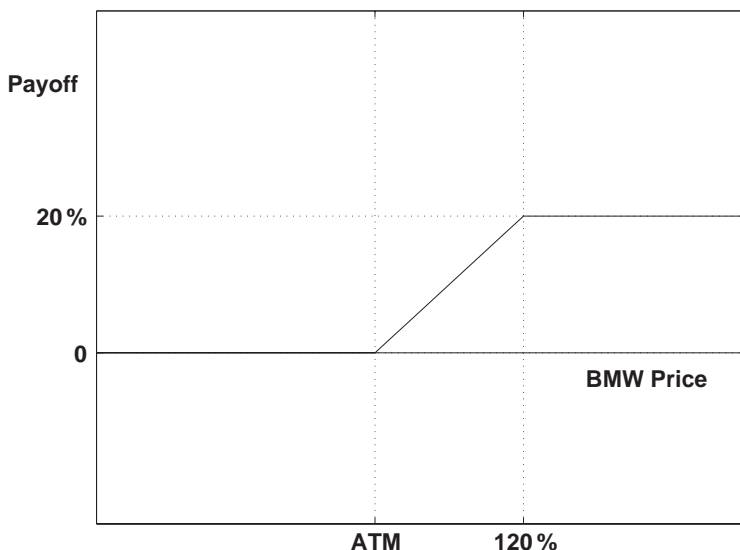


Figure 6.1 Payoff profile at maturity for an ATM/120% call spread on BMW

120% of its initial value. Every percentage gain over 20% does not make the investor any money, hence his profits are capped. When the two payoff profiles of a long ATM call option and a short 120% call option are combined, one gets the payoff graph in Figure 6.1.

From a skew perspective it is interesting to look at the call spread position from the trader's point of view. The trader, who is selling the 100% call and buying the 120% call, is short skew as he sells a lower strike option with a higher implied volatility than the higher strike option which he buys. The trader would therefore benefit if the skew goes down. To put it differently, the trader would make money if the skew becomes more shallow. This can happen when

- The implied volatility of the 100% call goes down and the implied volatility of the 120% call stays the same. This results in a smaller difference in implied volatility between the 100% call and the 120% call and hence the skew is more shallow.
- The implied volatility of the 120% goes up and the implied volatility of the 100% call stays the same. This means that the implied volatility of the 120%, which the trader is long, moves closer to the implied

volatility of the 100 % call. The balance is that the skew becomes more shallow and the trader makes money as a result.

- The implied volatility of the 100 % call goes down and the implied volatility of the 120 % call goes up. Again, the result is that the skew is less steep, which results in a profit to the trader.

The above shows the possibility for the trader to make money as a result of skew becoming less steep. Obviously, one could think of even more scenarios in which the trader benefits from a more shallow skew. For example, when both the implied volatility of the 100 % call and the implied volatility of the 120 % go down, but the implied volatility of the 100 % call goes down more than the implied volatility of the 120 % call. However, this scenario is not as clear cut as the above three scenarios. For example, if the vega of the 120 % call is larger than the vega of the 100 % call, a parallel downward shift in implied volatility for all strikes has a larger negative effect on the trader's long 120 % call position than the positive counter effect on the trader's short 100 % call position. This means that, even if the implied volatility of the 100 % call goes down more than the implied volatility of the 120 % call it might not be enough to offset the negative effect because of the long vega loss.

6.2 PUT SPREAD

A put spread is opposite to a call spread in the sense that, when an investor buys a put spread he is executing a bearish strategy, whereas buying a call spread is a bullish strategy. Buying a put spread means buying a higher strike put and selling a lower strike put. Since a put with a higher strike is always more expensive than a put with a lower strike and the same maturity, the investor has to pay a premium to put this strategy on. When an investor buys a put spread he partially funds his bearish view on the underlying stock by selling a lower strike put. However, by partially funding his bearish view he is also limiting the profits of his strategy, because the investor only benefits from the stock going down, up to the strike of the downside put. As an example consider an investor who buys an at the money put on Volkswagen (VW) and partially funds this by selling an 80 % put on VW. In this case the investor does profit from a decline in VW's share price up to the point where it has reached 80 % of its initial value. If the share price drops below 80 % of its initial value,

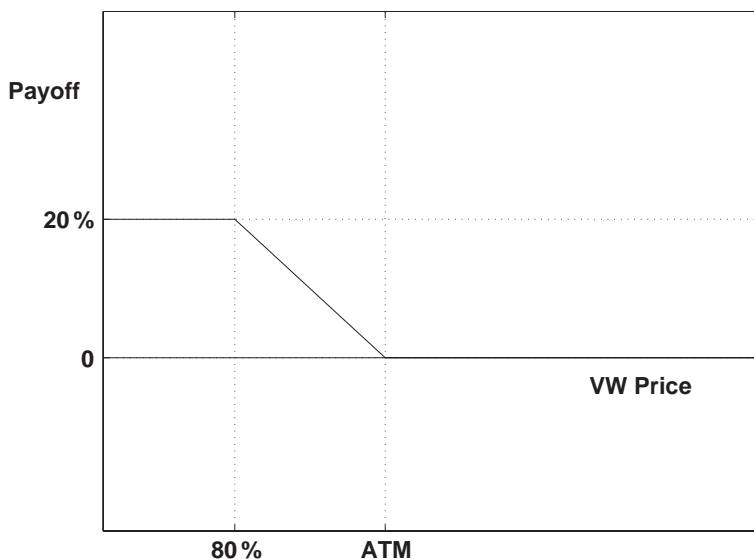


Figure 6.2 Payoff profile at maturity for an ATM/80 % put spread on Volkswagen

the investor does not benefit any longer. Again, when the two payoff profiles of a long at the money put and a short 80 % put are combined, one gets the payoff graph in Figure 6.2. It is important to understand that the payoff in Figure 6.2 is not the investor's profit because he would have paid an initial premium to put this strategy on. So, if this strategy costs him 2 % of the initial share price, the investor only starts to make money if the share price drops below 98 % of the initial share price. However, had he just bought an at the money put it would have cost him considerably more, for example 5 %, and therefore the share price would have to drop below 95 % of the initial share price before the investor would start to make money. Obviously, the put outright would not limit his upside potential to 18 % but in theory to 95 % if the share price went to zero.

Consider the trader who is selling the 100 %/80 % put spread. In this case the trader is long the 80 % put and short the 100 % put. Since he is long the lower strike option and short the higher strike option, the trader is long skew. This means that the trader would benefit from a steepening skew. In other words, if the difference in implied volatility

for a 10 % change in strike goes up, the trader, who is long skew, makes a profit.²

6.3 COLLAR

The collar or the risk reversal is used to give gearing to a certain view on the underlying rather than limiting the upside potential of that strategy, as for the call or put spread. A collar consists of a call with a higher strike and a put with a lower strike. If an investor is very bullish on a stock and he thinks the downside for the stock is very limited, he could buy an upside call and sell a downside put to partially fund this strategy. However, the investor might not be limiting his upside potential when the stock goes up, but by selling the put he does expose himself to the stock going down. So the investor is said to be gearing his view because he can actually lose more than his investment when the stock ends up lower than the strike of the put. If an investor is bearish on a stock he could buy a downside put and sell an upside call to fund his view partially or in full, but he therefore exposes himself to the stock price going up. Whether an investor is paying or receiving money to put on a risk reversal strategy depends on the forward and the skew. It is therefore ambiguous to talk about buying or selling a risk reversal.

Consider an investor who is very bullish on BMW's share price and therefore wants to buy a 1 year 110 % call. He also thinks that if the share price goes down it will not go down more than 10 % in the coming year. Therefore he wants to fund his 110 % call by selling a 90 % put with the same maturity. By combining the payoff profile of a long 110 % call with a short 90 % put one would get the graph in Figure 6.3. Again, it is important to remember that Figure 6.3 does not show the profit profile of the investor. The profit depends on how much the investor received or paid for putting this 90 %/110 % collar on in the first place. If BMW's dividend yield is higher than the risk free interest rate, the investor will most definitely have received money for this strategy. This is because, if the dividend yield is higher than the risk free interest rate, the 1 year forward is less than 100 % of the initial share price, which makes the

² Depending on the net vega position of the put spread, a steepening in skew does not always result in a profit to the trader. This was discussed in sub-section 6.1. A steepening of the skew as a result of the 80 % implied volatility going up and the 100 % staying the same or of the 100 % going down and the 80 % implied volatility remaining equal or lastly of the 80 % implied volatility going up and the 100 % volatility going down always makes the trader a profit.

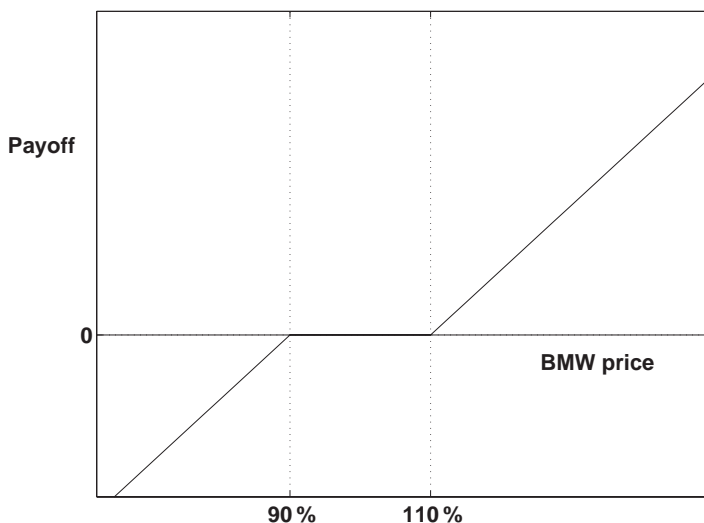


Figure 6.3 Payoff profile at maturity for a 90 %/110 % collar on BMW

1 year 90 % put more valuable than the 1 year 110 % call. On top of that the skew makes the 90 % put even more valuable and the 110 % call less valuable.

Another feature of the collar is that it has a relatively high delta, because, if a trader buys a put, he needs to buy shares to hedge himself and if the trader also sells a call, he needs to buy even more shares in order to be delta hedged. Because of this high delta the collar can also be used to express a dividend view. If a trader thinks BMW is likely to increase its dividends he quite happily takes the other side of the investor's strategy to buy the put and sell the call. Because, if BMW does increase its dividends, the forward should have been lower than the forward with the current dividends and the put will therefore increase in value and the call will decrease in value. The trader is also exposed to changes in skew as he is long the lower strike option and short the higher strike option. He is therefore long skew and benefits if the skew goes up.

6.4 STRADDLE

The straddle is a very popular strategy for investors who think the underlying stock will move away from its current level but do not know whether it will be up or down. A straddle consists of a call option and

a put option with the same strike and same maturity. Since buying a straddle means buying both a call and a put option, the underlying needs to move away from the strike price considerably for the strategy to make money. For that very same reason selling a one year ATM straddle is a very attractive strategy for investors who think the underlying will move sideways for the coming year. The investor will take in the premium of the straddle and will most likely have to pay very little at maturity because the stock will not have moved away much from current levels. So, on either the put or the call, the payout will be nothing and on the other one it will be very little.

As an example, consider an investor who knows that TomTom, the maker of car navigation systems, is due to come out with a statement which will be a real share mover. However, the investor does not know whether it will affect the share price positively or negatively. For that reason he decides to buy a 6 months at the money straddle on TomTom. Combining the two payoff profiles of a long at the money call and a long at the money put, one will get Figure 6.4. To show the difference between the payoff of buying an ATM straddle on TomTom and the profit of buying this straddle, suppose the investor paid €3 for the ATM straddle on TomTom and the share price is currently trading at €20. The




Figure 6.4 Payoff profile at maturity for a long ATM straddle on TomTom




Figure 6.5 Payoff profile at maturity for a long ATM straddle on TomTom

profit profile at maturity of buying an ATM straddle on TomTom would look like Figure 6.5.

The straddle is the perfect tool to play the volatility. Since a straddle consists of both a call and a put, it has a very large vega exposure. On top of that it has very little dividend exposure as the deltas of the call and the put offset each other. Therefore, a trader who believes that the volatility will go up in the near future should buy a straddle and subsequently delta hedge it. A trader who believes that the volatility will go down should sell a straddle and delta hedge it rigorously.

6.5 STRANGLE

The strangle is very similar to the straddle with the difference that the call and the put do not have the same strike. So, a strangle consists of a call with a higher strike and a put with a lower strike. Therefore, like the straddle, the strangle is also for investors who think that the underlying share will move away from current levels. However, an investor might prefer a 90%/110% strangle to an ATM straddle because he thinks that the share price will move much more than 10% from current levels but he wants to minimize his loss in case the share price does not move at

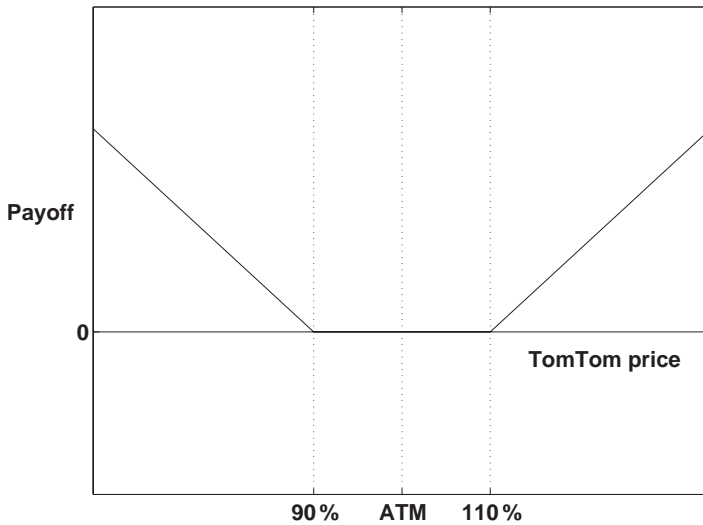


Figure 6.6 Profit profile at maturity for a long 90%/110% strangle on TomTom

all. By using the strangle to support his view, his initial investment is much less than when buying the ATM straddle. If the share does move much more than 10 % from current levels he profits almost as much as with an ATM straddle, but he has paid much less than in the case of the straddle, which makes his strategy less risky. As an example, consider an investor who thinks that TomTom's share price will either move away from its current €20 level by at least 20 % in the next six months or not do anything at all. Instead of buying an at the money straddle where the initial investment is large and in case of a big move the profit at maturity is large as well, the investor decides to buy a 6 months 90%/110% strangle where the initial investment is small and in case of a big move the profit is almost as large as for the straddle. The payoff profile at maturity for this strategy is shown in Figure 6.6.

To see that this strategy is less risky than buying a six months ATM straddle, suppose the investor paid €1.50 for the strangle and as per the example in section 6.4 he would have to pay €3 for the ATM straddle. If the share price were to move down by 20 %, TomTom's share price would be €16. The strangle would therefore make a profit of €0.50 (the strike of the 90 % put is €18 and therefore makes €2 on this put) and the straddle would make a profit of €1. The two different profit

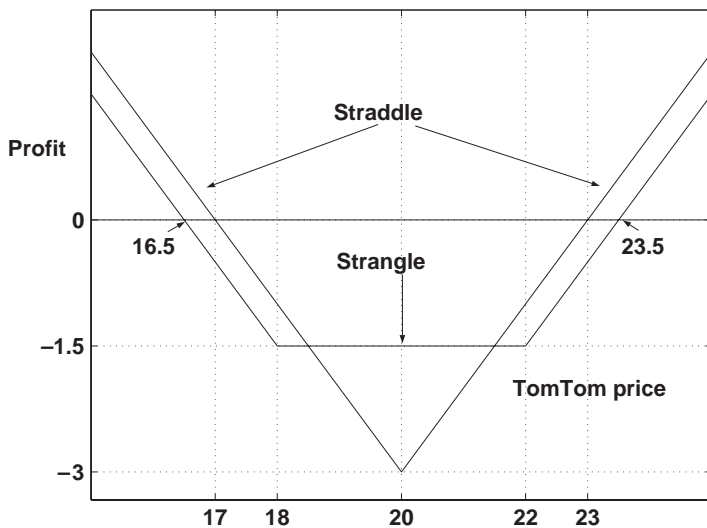


Figure 6.7 Profit profiles at maturity for a long 90%/110% strangle and an ATM straddle on TomTom

profiles of the straddle and strangle are shown in Figure 6.7. Figure 6.7 clearly shows that the strangle is a less risky strategy but the potential profit is less as well. However, for very big moves the difference in profit between the straddle and the strangle is very small in percentage terms.

Monte Carlo Processes

Most exotic options are ultimately priced with the aid of a Monte Carlo process. However, before using a Monte Carlo process one first needs to establish the inputs for this Monte Carlo process. These inputs can only be determined when understanding the exact risks and behaviour of the Greeks. Therefore, when pricing exotic options it is much more important to be aware of the risks than to fully understand the actual pricing of it through a Monte Carlo process. Nonetheless, it is important to have an understanding and intuition of the workings of a Monte Carlo process. This section is looking to provide exactly this understanding without discussing the different mathematical models that are concerned with generating the paths of a Monte Carlo process.

7.1 MONTE CARLO PROCESS PRINCIPLE

The principle of a Monte Carlo process is to generate a large but finite number of paths, aggregate the option payoffs associated with each path and subsequently divide the sum by the number of simulated paths. These paths are generated according to the process by which the stock or other underlying is assumed to move. For example, under Black–Scholes, stock price movements are assumed to behave like a Brownian motion. This means that the paths would be generated by a computer program using a standard normal distribution, where at any point in time the next move is distributed as a standard normal distribution with the implied volatility as standard deviation of this normal distribution. Armed with a large set of paths one can estimate the price of an option by adding the option payoffs corresponding to each respective path and dividing it by the number of simulated paths. In other words, let ω_i be the path corresponding to the i -th simulation and let $f(\omega_i)$ be the corresponding option payoff, the option price obtained through Monte Carlo simulation is

$$\frac{1}{N} \sum_{i=1}^N f(\omega_i). \quad (7.1)$$

Although formula 7.1 comes across as a very rough estimate of the specific option price, the Central Limit Theorem says that the error in this estimation converges to zero with a speed of

$$\frac{1}{\sqrt{N}}. \quad (7.2)$$

This means that a large enough number of paths ensures that an option price obtained with Monte Carlo simulation is very close to the fair value of this option.

7.2 BINOMIAL TREE VERSUS MONTE CARLO PROCESS

A binomial tree is another popular iterative method to calculate the price of an option. The difference between a binomial tree and a Monte Carlo process is that the binomial tree requires a discrete (not continuous in time) probability measure and it is therefore a less accurate approximation of the real movement of the underlying. On the other hand, because it has a finite number of observations, a real expectation can be calculated as the probabilities do add up to 1, whereas a Monte Carlo process only converges for a large enough number of paths to the expected value and therefore the price of the option. Since, for a continuous process, the probability of one specific path occurring is zero, it is impossible for a Monte Carlo process to calculate an expectation directly. However, for a large enough number of paths, one can expect the proportion of paths within a certain range to be equal to the probability of that stock price falling in that range prescribed by the process according to which the stock is supposed to move. Although the binomial tree and a Monte Carlo process bear some significant differences, the reason that the binomial tree is discussed in this section is because it is the easiest way to show how a Monte Carlo process derives the value of, for example, an American option where the possibility of early exercising should also be taken into account in the price.

7.3 BINOMIAL TREE EXAMPLE

Consider the binomial tree in Figure 7.1 on BMW's share price. Figure 7.1 depicts a 2 year binomial tree where with a 50 % chance the stock price is multiplied by 1.1 the following year and with another




Figure 7.1 Binomial tree with 50 % chance of stock price being multiplied by 1.1 and 50 % chance of being multiplied by $\frac{1}{1.1}$ and a dividend of € 2 in the second year

50 % chance the stock price is divided by 1.1 the following year. There is one catch in this binomial tree because there is a € 2 dividend in the second year and therefore the stock prices in the second year are € 2 lower than expected. Knowing the 2 year stock movement according to the binomial tree in Figure 7.1 it is easy to price a 2 year € 97 European call. Indeed, there are only two paths where the € 97 European call has a payout. Namely when BMW's stock price is € 98 after two years and when it is € 119, both of which have a 25 % chance of occurring. Assuming that the interest rate is zero, the price of the 2 year € 97 European call is

$$0.25 \times (98 - 97) + 0.25 \times (119 - 97) = \text{€ } 5.75$$

When pricing a two year € 97 American call, one also has to check at every point in the tree whether it is optimal to early exercise the American call option. For the tree in Figure 7.1, the only point that needs to be checked is the € 110 stock price after 1 year. In other words, one has to look at the part of the tree in Figure 7.2 to determine whether the option of early exercising has an impact on the price of the American option. The payout on the € 97 American call is € 13 if one were to exercise after 1 year. Assuming that the interest rate is zero and that BMW's share price is trading at € 110 after one year, the present value of the expected value of the € 97 American call in the second year is

$$0.5 \times (119 - 97) + 0.5 \times (100 - 97) = \text{€ } 12.5$$




Figure 7.2 Second year of the binomial tree from Figure 7.1 zoomed in to the ‘up’ path

This means that it is optimal to exercise the 2 year €97 American call on BMW after the first year if BMW is then trading at €110. The price of the American call can therefore be established as

$$0.5 \times (110 - 97) + 0.5 \times 0 = \text{€}6.5.$$

7.4 THE WORKINGS OF THE MONTE CARLO PROCESS

Knowing that, for a large enough number of paths, the percentage of paths falling within a certain range converges to the probability of that range occurring according to the process by which the stock is supposed to move, a Monte Carlo simulation can effectively be seen as an expectation. Therefore, determining the fair option value through a Monte Carlo process is ultimately the same idea as calculating the expected value of an option from a binomial tree. For a Monte Carlo process, a computer program randomly generates paths for which the associated option payoffs are aggregated and subsequently divide by the simulated number of paths to get the fair option value. On the other hand, a binomial tree reduces the stock price movements to a finite number of paths and hence compromises the accuracy of the stock price movements, but the expected value is real as the probabilities do add up to 1. In other words, increasing the number of paths of a Monte Carlo process means that the outcome resembles an expected value more closely and therefore the Monte Carlo estimation is closer to the fair value of the option, whereas increasing the number of paths of a binomial tree means that the binomial tree more closely resembles the actual stock price movements and hence the expected value of the option is closer to the fair value of the option.

Chooser Option

Before proceeding with discussing real exotic options, the chooser option is an interesting one to discuss. The chooser option is often characterised as an exotic option because it comes across as an option that is difficult to price. However, closer investigation of the chooser option shows that it can actually be priced as a series of regular options.

The buyer of a chooser option has the right to decide, up to a certain date, whether it should be a call option or a put option. The strikes of either the put or the call can be the same but need not necessarily be the same. When the chooser option specifies that the strikes are the same, the chooser option is referred to as a simple chooser. When the strikes or even the expiries are not the same the chooser is referred to as a complex chooser. The following sub-section discusses an example of how a simple chooser can be priced.

8.1 PRICING EXAMPLE: SIMPLE CHOOSER OPTION

Consider a European chooser option on BMW with a maturity of 2 years, a strike price of €40 and the buyer has 2 months to choose whether he wants the chooser to be a €40 call or a €40 put, see Table 8.1. In other words, the fact that the chooser option is European means that, once the buyer has chosen whether the chooser should be a put or a call, the chosen option is a European style option and can therefore only be exercised at maturity. The European property has therefore nothing to do with the timing of the buyer's choose date. Namely, the buyer can pick any day up to 2 months to decide whether the chooser should be a call or a put. However, the buyer would always pick the last possible date to decide whether the chooser should be a put or a call as there is no benefit in choosing earlier. Suppose that the 2 year forward is €40 when BMW is worth €38 after 2 months, i.e. the 22 months forward in 2 months is €40 if BMW's share price is equal to €38 in two months' time. Armed with this assumption the 2 year chooser option on BMW can be priced

Table 8.1 Terms of the chooser option

Underlying	BMW
Option style	European
Option type	Chooser
Maturity	2 years
Choose date	Up to two months
Strike	€ 40

as follows by the trader. In other words, if the trader buys the following position he is fully hedged against the short chooser option.

- Buy a 2 year € 40 strike European call **plus**
- Buy a 2 months € 38 strike European put.

To see that the sum of the above two options prices the chooser option on BMW accurately, consider the following two scenarios

1. After two months BMW's share price is above € 38.

Since the 2 months put expires worthless, the trader is just left with a € 40 call option expiring in 22 months. This mirrors exactly the position as a result of the choice of the buyer of the chooser option. Indeed, the buyer of the chooser option decides to turn the chooser option into a € 40 call option instead of a put option since the 22 months forward is higher than € 40 and therefore the call is worth more than the put.

2. After two months BMW's share price is below € 38.

In this case the put expires 'in the money' and therefore the trader exercises the € 38 put within his hedging portfolio, which results in a short share position for the trader. This means that the trader is left with a short share position combined with a long call position. Put-call parity means that a short share position combined with a long call position is equal to a long put position. Again, this mirrors exactly the short chooser position. Indeed, the buyer of the chooser position decides to turn the chooser into a 22 months € 40 put as the 22 months forward is below € 40 and therefore the 22 months € 40 put is worth more than the 22 months € 40 call.

The above shows that a simple chooser is easy to price. Consider a chooser option with a strike price K_c , a maturity T_1 and a maturity up to the choose date of T_2 . To price this chooser option, one just has to

determine the stock level, S_f , such that the forward on the choose date with an expiry equal to the maturity date of the chooser is equal to the strike, K_c , of the chooser, i.e. the forward value of a stock price S_f and a time to maturity of $T_1 - T_2$ is equal to K_c . A put and a call option with a strike equal to the forward and the same maturity are worth the same. Hence, on the choose date, S_f determines the tipping point for the chooser call option to be worth more than the chooser put option. The chooser option can now be priced as

- A call option with a time to maturity of T_1 and a strike K_c **plus**
- A put option with a time to maturity T_2 and a strike price S_f .

Alternatively, the chooser option can also be priced as

- A put option with a time to maturity of T_1 and a strike K_c **plus**
- A call option with a time to maturity T_2 and a strike price S_f .

This is because being long a put option and a share is the same as being long a call option, which mirrors exactly the position of the short chooser option when, on the choose date, the share price is above S_f in which case the owner of the chooser turns the chooser into a call option. Obviously, when, on the choose date, the share price is below S_f the owner of the chooser decides to choose the put option. Again, mirroring the position of the hedged portfolio as the call with strike S_f and maturity T_2 is 'out of the money' and therefore not exercised. This results in a long put position in the hedged portfolio.

8.2 RATIONALE BEHIND CHOOSER OPTION STRATEGIES

The investors in chooser options apply a similar strategy to investors in straddles. They both believe that a stock is about to go through a volatile period but are unsure about the direction of the stock price. The difference between an investor in a chooser option and an investor in straddles is that the first is confident that the direction is revealed within a certain time frame and is therefore willing to choose between a put and a call after this time frame. As a result the investor pays less for the chooser option than he would have paid for the equivalent straddle.

Digital Options

A digital option is one of the most straightforward options. It is nothing more than an option that gives a fixed payout if it is below or above a certain point and does not give a payout at all in all other cases. Despite being so simple, the digital option could still classify as an exotic option as its payout cannot be replicated by a set of standard options. However, rather than viewing this as an exotic option and using a Monte Carlo process to calculate the price of a digital option, one can make a good estimate of the price of a digital option by using option spreads. On top of that one will find that the only real way to risk manage the digital option is with option spreads. The vast majority of digital options are European. Therefore this chapter only considers European digital options.

Consider the following example. An investor buys a 3 month European digital option on BMW which pays € 10 if after 3 months the stock is above € 50 and pays 0 if BMW is below € 50 at maturity. Figure 9.1 shows the payoff profile at maturity for this digital option on BMW. The trader who sells this digital option on BMW can easily replicate the price of this digital by using a geared call spread. The gearing of this call spread depends on the width of the call spread. The wider the call spread the lesser the gearing and the more conservative the price, i.e. the trader would charge the investor more to buy the digital. The trader believes that a € 2.5 wide call spread is enough to be able to risk manage this position. In other words, the trader will price this by selling a € 47.5/50 call spread that is 4 times geared. To see that this replicates the behaviour of the digital, consider three different BMW share prices at maturity. First, suppose BMW's share price at maturity is € 55. Under the terms of the digital the investor is supposed to get € 10. This is exactly what the trader's position prescribes. Indeed, the trader has a 4 times geared € 47.5/50 call spread and each call spread gives a payout of € 2.5 if the stock price ends up at € 55, resulting in a total payout to the investor of € 10. Secondly, If the share price of BMW is € 49 at maturity, the investor is not supposed to receive anything under the terms of the digital. However, the 4 times geared € 47.5/50 call spread prescribes a payout of € 6, resulting in a windfall




Figure 9.1 Payoff at maturity from buying a digital call option on BMW. Strike price = € 50 and digital payout € 10

of € 6 to the trader. This obviously shows that the trader has priced the digital conservatively by choosing this € 47.5/50 call spread and would have been more aggressive by choosing a tighter call spread. However, the width of the call spread is necessary to make up for the large gamma and pin risk around € 50. Indeed, just before expiry with the stock price exactly at € 50 the digital is not worth anything, but if the stock price goes up to € 50.01 the digital is all of a sudden worth € 10. This is obviously extremely hard to risk manage because the trader would not know what delta to put against this option. Therefore the best way to price this digital is by means of a geared call spread. Even with a call spread the trader would still need to manage a large pin and gamma risk, but the fact that he prices the digital as a call spread gives him a provision against this risk. However, since the call spread gives the trader a short position in the € 47.5 call and a long position in the € 50 call, the trader has shifted his pin risk to the 47.5 strike and enjoys the benefit of pin on the 50 strike. In a way, the trader has given himself a cushion to manage the pin risk of the digital by forcing himself to be hedged from € 47.5 onwards because of the call spread. Thirdly, when the share price of BMW is € 46 neither the digital nor the call spread prescribes a payoff to the investor.

9.1 CHOOSING THE STRIKES

A natural question would be ‘why is the digital in the previous section not replicated as a € 50/52.5 call spread’? The answer is that the € 50/52.5 call spread achieves exactly the opposite of what the trader wants to achieve, namely giving himself a cushion. The € 50/52.5 call spread starts to act once the digital event has already occurred. In a way the € 50/52.5 call spread makes the trader feel ‘richer’ than he actually is. For example, if the stock price ends up at € 51 the 4 times geared call spread would specify a payoff of € 4 whereas the digital would demand a payoff of € 10. In other words, the hedging scheme of the € 50/52.5 call spread does not build up to a € 10 payout at a stock price of € 50 whereas a € 47.5/50 call spread does exactly that. In a way, the € 50/52.5 call spread is chasing the tail of the digital whereas the € 47.5/50 is forcing the trader to act before the digital actually kicks in and is therefore well positioned to deal with the ‘explosive’ digital event.

9.2 THE CALL SPREAD AS PROXY FOR THE DIGITAL

A call spread is not only used to price a digital but, from the perspective of the trader, it is also the product that he actually trades. In other words, when a trader sells a digital he books a call spread in his risk management system instead of the exact terms of a digital. In this way, he can accurately risk manage the digital as a call spread. The previous sub-section showed that a call spread is a conservative proxy for the digital and therefore at expiry of the European digital the trader checks the payout of the call spread against the payout according to the terms of the digital. A payout of the call spread versus a non payout of the digital subsequently results in a windfall to the trader.

9.3 WIDTH OF THE CALL SPREAD VERSUS GEARING

The narrower the width of the call spread the higher the gearing necessary to replicate the digital. In turn, the higher the gearing of the call spread the larger the risk that needs to be managed on the short strike, and since the strikes of the call spread are so close together the trader gives himself less of a cushion to prepare for the digital event. In other words, the narrower the width of the replicating call spread the higher the risk and the more the call spread behaves like the digital.

Barrier Options

Barrier options are very popular amongst retail investors as the barrier feature provides the investor with additional protection or leverage. From a risk management perspective, barrier options are interesting because the risks are discontinuous around the barrier and therefore the Greeks become less predictable and very often even change sign around the barrier. Barrier options are priced with Monte Carlo processes, but before one can price barrier options one has to be fully aware of the risks around the barrier. The risks associated with the barrier are typically of such a nature that the terms of the barrier option have to be adjusted slightly to be able to capture all these risks in the price and to also be able to manage the barrier risk properly. This section discusses in depth the risks embedded in barrier options and shows how these risks should be taken into account in the price. This section also emphasises the skew risk associated with a barrier option. Although there are basically 8 types of single barrier options, the down-and-in put is used as a leading example in order to get across all the risks within a barrier option.

The reason behind this approach is that the underlying causes for barrier option risk are generic and once the drivers of barrier option risk are understood for the down-and-in put, one will be able to derive the risks for other types of barrier options. However, in this regard it is important to distinguish between the drivers behind risk and the actual risk. The fact that the drivers behind risk are the same for the different types of barrier options does not mean that the risks are the same. In fact, between different types of barrier options the symptoms of risk can be completely opposite in nature. There are in total 8 different types of single barrier options and the down-and-in put is chosen as a leading example because of its popularity amongst retail investors. The 8 types of single barrier options, the typical retail demand¹ and the barrier risk to the trader for each type of barrier option are summarised in Table 10.1.

¹ Whether it is a barrier option or not, retail investors will typically sell puts and buy calls. However, it does occur that a retail investor effectively buys a down-and-out put as part of a structured product.

Table 10.1 Types of single barrier options

Barrier type	Typical retail demand	Trader's barrier risk
Down-and-in put	Sell	Significant
Down-and-out put	Sell/Buy	Significant
Up-and-in put	Sell	Low
Up-and-out put	Sell	Low
Down-and-in call	Buy	Low
Down-and-out call	Buy	Low
Up-and-in call	Buy	Significant
Up-and-out call	Buy	Significant

Aside from single barrier options another class of barrier options exists, namely the double barrier options. Double barriers are discussed in sub-section 10.13.

10.1 DOWN-AND-IN PUT OPTION

A *down-and-in* put option is an option that only becomes a European put option when the share price breaches a downside barrier. For example, consider a trader who buys a 1 year 100 % European put option on BMW that only becomes a put option if BMW's share price ever drops below 70 % of BMW's initial share price during the one year term of the option. In other words, the trader is long a 100/70 % down-and-in put and its payout at maturity is shown graphically in Figure 10.1. The dotted diagonal line in Figure 10.1 indicates that the payoff in this area only occurs at maturity if the barrier of 70 % has been hit during the life of the option. If this has not been the case, the payoff would still be 0.

Although the pricing of the down-and-in put is nothing more than a Monte Carlo process, there are two phenomena that one needs to understand before one can price, let alone, trade a down-and-in put. The main risk of a down-and-in put lies in the *delta change* over the barrier. Another risk that needs to be understood is the *skew* exposure of a down-and-in put. The following sub-sections will discuss these risks in more detail.

10.2 DELTA CHANGE OVER THE BARRIER FOR A DOWN-AND-IN PUT OPTION

Consider the trader who buys the 100/70 % down-and-in put on BMW from an investor. If the delta at inception is 0.4, the trader hedges himself

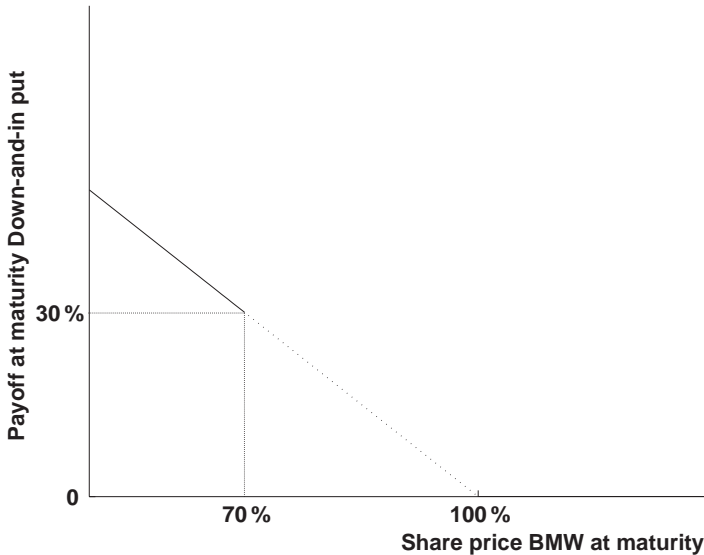


Figure 10.1 Payoff at maturity of a down-and-in put with strike price 100 % and barrier level of 70 %

by buying BMW shares in a ratio of 0.4 shares per option. When BMW's share price breaches the 70 % barrier level, the down-and-in put option goes from not being an option to a put option that is 30 % in the money. This obviously means that the value of the down-and-in put increases significantly over the barrier, even if the share price moves down ever so slightly, i.e. from 70.1 % to 69.9 %. Since delta measures the sensitivity of the option price to share price movements, the absolute delta of the down-and-in put becomes extremely large when the share price approaches the barrier. This absolute delta can even, and typically does, become greater than 1 when the stock price gets close to the barrier. Since the absolute delta of a regular option can never be greater than 1, the trader would inevitably accumulate too many BMW shares when BMW's share price approached the barrier. This means that the trader would need to sell any excess shares when the share price breached the 70 % barrier. The initial price of the down-and-in put, which any Monte Carlo price comes to, assumes that these excess shares can be sold exactly at the barrier level of 70 %. However, in practice this proves to be extremely difficult as the share price is already going down for the barrier to be breached in the first place, and the fact that the trader needs to sell a large quantity of shares will push the share price down even further.

Therefore, the trader almost certainly sells these excess shares below the barrier level and as a result incurs a loss on the sale of the excess shares. To avoid this loss, the trader should give himself a cushion to sell any excess shares over the barrier. In order to do that, the trader prices and risk manages a slightly different option. Namely, an option where the barrier is shifted downward in such a way that the trader has enough room to sell the excess shares without incurring a loss. This means that, if the trader believes that he needs a 3 % cushion to sell the shares over the barrier, the trader prices and risk manages a 100/67 % down-and-in put rather than a 100/70 % one. The following sub-section discusses the factors that influence the magnitude of the barrier shift. However, before that is discussed it is important to recognise that a long down-and-in put position goes from being long gamma to short gamma as the share price approaches the barrier. In other words, *the gamma of a down-and-in put changes sign around the barrier*.

10.3 FACTORS INFLUENCING THE MAGNITUDE OF THE BARRIER SHIFT

There are 5 main influencing factors that impact the magnitude of the barrier shift, which are all discussed separately in this sub-section. This sub-section only discusses the influencing factors for the down-and-in put, but the same analysis can easily be extended to other types of barrier options. Therefore, this sub-section serves as a blueprint for all barrier options.

- ***The size of the down-and-in put transaction.*** The larger the size the more shares have to be sold over the barrier and therefore the trader is more likely to move the share price against him, i.e. downward. In other words, the larger the size the larger the barrier shift.
- ***The difference between strike price and barrier level.*** If this difference is large, the down-and-in put goes from not being a put to a put that is far in the money when the share price breaches the barrier. This means that the greater the difference between strike price and barrier level the greater the barrier shift that a trader would apply. For example, a 130/70 % down-and-in put should have a larger barrier shift than a 100/70 % down-and-in put.
- ***The volatility of the underlying stock.*** The larger the volatility of a stock the larger is the risk to the trader of the stock price approaching

the barrier level. This can be seen by a simple example. Consider the 100/70 % down-and-in put on a highly volatile stock with the stock currently trading at 71 % of its initial level. Suppose that the trader has priced this option as a 100/67 % down-and-in put and is also risk managing it as such. Suppose further that the absolute delta of the 100/67 % is 2.5 and that the trader is long 200 thousand of these 100/70 % down-and-in puts. Since the trader is risk managing the 100/70 % down-and-in put as a 100/67 % one and the delta of the latter is 2.5, the trader is long 500 thousand of the underlying shares. If the stock gaps down 10 %, the down-and-in put obviously knocks in and therefore becomes a regular deep in the money 100 % put. Suppose that the delta of this deep in the money 100 % put is 1 and hence the trader needs to have a long position of 200 thousand shares against this put. However, the trader was long 500 thousand shares and therefore he had to sell 300 thousand shares over the barrier. Because the stock gapped down, the trader can only sell these shares at 61 % of the initial stock price. Since the trader risk managed this as a down-and-in put with a barrier of 67 %, he is basically long 300 thousand shares outright from 67 % of the initial level to 61 % of the initial level. If the initial level was, for example, € 100, this would amount to a loss of € 1.8 million. The above example shows that, if a stock is or can be highly volatile, the buyer of a down-and-in put needs a larger barrier shift in order to be protected against a larger move. In other words, **the larger the volatility the larger the barrier shift.**

- **The barrier level.** Since lower stock prices tend to go hand in hand with higher volatilities and higher volatilities result in larger barrier shifts, traders need to apply larger barrier shifts to lower barrier levels. For example, the barrier shift for a 100/60 % down-and-in put should be larger than for a 100/70 % down-and-in put.

Some cases are ambiguous as to which down-and-in put should have the larger barrier shift. For example, does a 120/70 % down-and-in put need a larger barrier shift than a 100/60 % down-and-in put? On the one hand one can argue that it does as the difference between strike and barrier is larger for the 120/70 % down-and-in put than for the 100/60 % one. On the other hand, the barrier level of the 100/60 % down-and-in put is lower than the 120/70 % one and therefore it can be argued that the 100/60 % down-and-in put should have the larger barrier shift. This example shows that determining the magnitude of a barrier shift is not an exact science. It is up to

the trader's own risk assessment which one of the two down-and-in puts he assigns a larger barrier shift to. However, for sure the trader will assign a larger barrier shift to both the 120/70 % and 100/60 % down-and-in puts than the barrier shift he finds appropriate for the 100/70 % and even the 110/70 % down-and-in put.

- **The time left to maturity.** The closer one gets to maturity the larger the absolute delta will be just before the barrier and therefore the larger the change in delta over the barrier. This obviously translates into larger barrier shifts for shorter maturities and typically traders use a scattered barrier table where the barrier shift increases once the down-and-in put gets closer to maturity. The reason that, if the stock price is just above the barrier, the absolute delta becomes larger when the time left to maturity gets shorter is because there is less time left for the down-and-in put to knock in. This means that the barrier becomes an all or nothing event and starts to mimic the features of a digital option and therefore the change in value of the down-and-in put over the barrier becomes larger when the time left to maturity shortens. Since the change in value increases for a barrier breach when the time left to maturity is shorter, the absolute delta is larger for a shorter time to maturity, see Figure 10.2. This all adds up to the fact that **the shorter the time left to maturity the larger the barrier shift a trader applies to a down-and-in put.**




Figure 10.2 Delta of a down-and-in put plotted against the time left to maturity when the share price is just above the barrier level

10.4 DELTA IMPACT OF A BARRIER SHIFT

Although a barrier shift is mainly there to manage the Greeks close to the barrier (gamma and delta are larger without a barrier shift) one has to realise that a barrier shift also affects the delta further away from the barrier. For example, the larger the barrier shift on a long down-and-in put the smaller the absolute delta and therefore the smaller the amount of long shares to delta hedge the down-and-in put. This means that a barrier shift can also be used to express a delta view. If one is bullish on the underlying stock one should choose to have a smaller barrier shift and therefore effectively run a longer delta (more long shares) against the down-and-in put. If one is bearish on the underlying stock one should choose to increase the barrier shift and therefore run fewer long shares against the long down-and-in put position.

10.5 SITUATIONS TO BUY SHARES IN CASE OF A BARRIER BREACH OF A LONG DOWN-AND-IN PUT

There are specific situations in which the trader, who is long a down-and-in put, does not need to sell shares when the stock price goes through the barrier level, but actually needs to buy shares. These situations occur when the stock price is still far away from the barrier level and therefore has not accumulated any excess delta after which it goes through the barrier in a gap move down. In these cases the absolute delta actually increases because of a barrier breach.

10.6 UP-AND-OUT CALL

An up-and-out call is a call option that ceases to exist if the underlying stock hits a certain upside barrier during the life of the option. An up-and-out call is obviously cheaper than a regular call and can therefore be advantageous to investors who are bullish on a stock but do not believe it will go above a certain level during the life of the option. Consider a trader who sells a 1 year 100/120 % up-and-out call option on BMW. This means that if BMW's share price ever breaches 120 % of the initial level, the call knocks out and therefore ceases to exist. In this case the up-and-out call loses a lot of value over the barrier as it goes from a call that is 20 % in the money to something that is worthless. The trader, who holds a short position in this up-and-out call, is therefore short shares against this option as his delta hedge when BMW's share price gets close to the 120 % barrier. Note that a regular short call position is

hedged by buying shares. This is initially also the case for an up-and-out call, but when the share price gets closer to the upper barrier level there is an inflection point where the trader will need to go short shares to be delta hedged. However, once the up-and-out call knocks out, the trader needs to buy back these shares as it is no longer a hedge against anything (the option has knocked out). Unlike the long down-and-in put position the trader needs to shift the barrier upward for a short up-and-out call position. However, apart from the direction of the shift the factors that influence the magnitude of the barrier shift are the same for both the long down-and-in put position and the short up-and-out call position. These influencing factors are described in sub-section 10.3.

10.7 UP-AND-OUT CALL OPTION WITH REBATE

Consider a trader who sells the 1 year 100/120 % up-and-out call option with the additional feature that when the up-and-out call option knocks out it pays a rebate to the investor of 10 %, i.e. half its intrinsic value. This can obviously still be priced as a Monte Carlo process. A more interesting question is how does this rebate affect the barrier shift? Because of the rebate, the up-and-out call goes from an option with an intrinsic value of 20 % to something that is worth 10 % when the stock breaches the barrier of 120 %. Hence, the change in value over the barrier is far less in the case where the up-and-out call pays a rebate than if it does not pay a rebate. In other words, the discontinuity of the up-and-out call price is less with a rebate than without a rebate. This means that the trader will be short fewer shares when the share price gets close to the barrier and therefore needs to buy back less shares when the barrier is breached. For this reason the barrier shift can be smaller when the up-and-out call pays a rebate than if it does not pay a rebate.

10.8 VEGA EXPOSURE UP-AND-OUT CALL OPTION

A long position in an up-and-out call option is not necessarily long vega. In fact, more often than not a long up-and-out call position results in being short vega. This means that, if the implied volatility of the underlying goes up, the up-and-out call option actually becomes less valuable. The reason being that a higher implied volatility results in a higher probability of the call knocking out and therefore a lower chance of the up-and-out call having a payout at maturity. Naturally, an up-and-out

call is not always short vega. Indeed, if it is a very upside barrier the probability of the up-and-out call knocking out is very low anyway and therefore an increase in implied volatility has a much bigger impact on the option part of the up-and-out call than on the fact that there is a higher chance of the option knocking out. Hence an up-and-out call is long vega for very upside barriers and short vega for lower barriers.

10.9 UP-AND-OUT PUT

There are certain types of barrier options where there is no need to apply a barrier shift. For example, a trader who is long an up-and-out put does not need to apply a barrier shift. The reason being that the delta hedge against a long up-and-out put position is always a long share position. This means that the trader needs to sell shares over the barrier. But for that he does not need to shift the barrier as he can always sell out his shares 1 cent below the barrier in which case he would, at worst, only lose a cent on the number of shares he needs to sell.

10.10 BARRIER PARITY

Now that both knock-in as well as knock-out barriers have been explained, it is time to introduce the barrier parity. Barrier parity states that a knock-out (KO) option plus a knock-in (KI) option with the exact same barrier have the same payoff profile as the regular European option. Mathematically this reads as follows,

$$\text{KO-call} + \text{KI-call} = \text{European call} \quad (10.1)$$

$$\text{KO-put} + \text{KI-put} = \text{European put} \quad (10.2)$$

As an example, consider a 100/70 % down-and-out put plus a 100/70 % down-and-in put. The down-and-out put means that a 100 % put ceases to exist whenever the 70 % barrier is breached. However, the down-and-in put ensures that a new 100 % put comes into existence at the same time and therefore a down-and-in put plus a down-and-out put is effectively a regular European put.

10.11 BARRIER AT MATURITY ONLY

The previous sub-sections discussed barrier options where the barriers were live continuously throughout the life of the option. However, certain

barrier options are only live at maturity or on specific days. These barrier options can still be priced as a Monte Carlo process. However, barrier options that are only live at maturity can also be priced as a combination of European options.

Consider a trader who buys a 1 year 100/70 % down-and-in put option on BMW where the option can only knock in at maturity. This option is worth less than the equivalent down-and-in put option that can knock in throughout the life of the barrier option. The way to price this down-and-in put option with knock in at maturity only is as follows:

- The trader buys 10 times a 1 year 70 % European put option
- The trader sells 10 times a 1 year 67 % European put option
- The trader buys one 1 year 70 % European put option

The fact that the trader buys 10 times the 70 % put and sells 10 times the 67 % put is an overhedge and therefore a conservative way to replicate the payoff of the 100/70 % down-and-in put at maturity only. The 3 % wide put spread can be seen as a barrier shift. The payoff at maturity for the down-and-in put (DIP) with knock in at maturity only and the payoff replicated through a combination of European puts are both shown in Figure 10.3. Figure 10.3 clearly shows that the replication through a geared put spread plus a 70 % put is a lower estimate and therefore results in a more conservative (lower) price for the 100/70 % down-and-in put with knock in at maturity only. Obviously, the tighter the put spread the more the replication converges to the actual price of the down-and-in put with knock in at maturity only. The gearing of the put spread depends on the width of the put spread and the difference between strike and barrier level. The gearing can be calculated by dividing the strike/barrier differential by the width of the put spread. In other words, if one wants to replicate the 100/70 % down-and-in put with knock in at maturity only in terms of a 70 % put plus a 2 % geared put spread, i.e. a 70/68 % put spread, the gearing of the put spread would be 15 ($30/2 = 15$).

10.12 SKEW AND BARRIER OPTIONS

Barrier options tend to have large skew exposures. Luckily it is easy to see whether a barrier option is long or short skew. Risk wise one can compare a knock in barrier to a long option position at that barrier and a knock out barrier to a short option position at the specific barrier. This

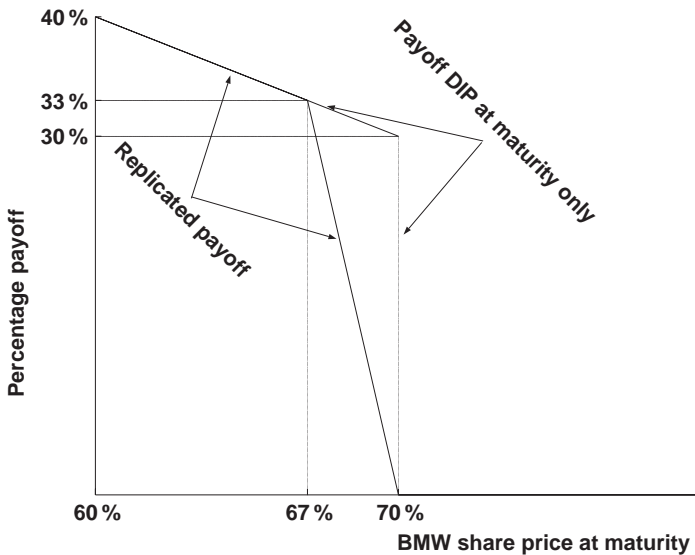


Figure 10.3 Payoffs for the down-and-in put, knock in at maturity only, and its replicated payoff, 10 times geared 70/67 % put spread plus a 70 % put

means that the owner of a down-and-in put has a risk profile that can be compared to being long an option at the knock-in level. Since this is naturally a downside option, the owner of a down-and-in put option is long skew. Another way to see that the owner of a down-and-in put option is long skew is that, if the skew goes up, the stock is expected to become more volatile when it goes down and therefore the probability of the down-and-in put knocking in goes up. As a result the down-and-in put becomes more valuable. A trader who has a short position in an up-and-out call is short skew. This can be seen by comparing his position to a call spread, where he is short a call with a strike price equal to the strike of the up-and-out call and long an upside call with a strike equal to the barrier level of the up-and-out call. Another way to see that he is short skew is by looking at the price change of the up-and-out call if the skew flattens (goes down). If the skew flattens the upside implied volatility goes up and therefore it becomes more likely that the up-and-out call knocks out, which makes the up-and-out call less valuable. The trader, who holds a short position in the up-and-out call, makes money from the skew flattening and is therefore short skew. Table 10.2 gives

Table 10.2 Skew position barrier option

Position barrier option	Skew exposure barrier option	Direction barrier shift
Long down-and-in put	Long skew	Downward
Long up-and-out put	Long skew	No shift needed
Long up-and-in call	Short skew	Upward
Long down-and-out call	Short skew	No shift needed
Short up-and-out call	Short skew	Upward

different barrier option positions with their associated skew exposures. Obviously, if the position is opposite to the one stated in Table 10.2, the skew exposure is the opposite as well.

10.13 DOUBLE BARRIERS

Apart from options with just one single barrier, there is another common class of barrier options, namely the so-called double barriers. However, even with a double barrier option one has to consider each barrier separately to determine the magnitude and direction of the barrier shift. Double barrier options are typically used within structured products to achieve a specific type of payoff. Although specific examples of double barrier options and their applications within structured products are given in Chapter 25, it is important to devote a sub-section to double barrier options that discusses the two main types of double barrier options. Before looking at the different types of barrier options it is worth emphasising once more that the pricing of any barrier option is established by using a Monte Carlo process. Therefore, from a trading perspective, the most interesting element of a barrier option is determining the magnitude and direction of the barrier shift.

The first type of double barrier option features a dependency between the trigger of the first and second barrier. That is to say, the second barrier can only trigger if the first barrier has been triggered. Although the Monte Carlo modeling is slightly more complex because of the built in condition, it is in fact more transparent from a risk perspective than if the barriers were completely independent. Indeed, one specifically knows the conditions and therefore the option characteristics under which the barriers can be triggered. This means that one can more easily establish the required barrier shift for each of the barriers. Consider the following

example. A trader sells a 100 % call that knocks in at 90 % and knocks out at 120 %. However, it can only knock out after it has breached the 90 % level first. For this double barrier it is relatively easy to determine the direction of the respective barrier shifts. Since the 120 % barrier is dependent on the 90 % barrier, one should first focus on determining the shift for the 90 % one. Because the trader is short the double barrier call option, and the double barrier call option increases in value when the 90 % level is breached, the trader will want to shift the 90 % barrier level downward. To determine the required shift for the 120 % barrier, it is important to establish that this barrier can only trigger once the double barrier has already become a call option (the 90 % barrier has been breached). Therefore, to determine the shift for the 120 % barrier one basically has to determine the shift for a regular 100/120 % up-and-out call. The required barrier shift for a short up-and-out call is upwards. Indeed, the option goes from being an 'in the money' call option to not being a call at all and therefore goes down in value significantly. The trader will have hedged himself by selling shares, which he needs to buy back once the barrier level is breached. To give himself a cushion to buy back these shares, the trader wants to shift the barrier upwards at the 120 % level.

For the second type of double barrier option each of the two barriers can be triggered independently. In other words, when the specific barrier level is breached, the barrier triggers regardless of whether the other barrier has been triggered. This type of double barrier option is slightly easier to model as a Monte Carlo process. However, from a risk perspective it is actually less transparent. The reason being that one has an indefinite picture of what the option characteristics look like when either of the barriers breaches. This becomes clear by considering the example above with the difference that the 120 % barrier can trigger regardless of whether the 90 % barrier has been triggered. This double barrier option is worth less than the conditional one in the example above as it can already knock out before it actually has become a call option. This also means that the barrier shift at 120 % can be less if it breaches before the 90 % barrier has been breached. However, if the 90 % barrier has been breached first, one requires the same barrier shift as for a regular up-and-out call. Hence, an unconditional double barrier has a bit more ambiguity regarding the exact barrier shift. Nonetheless, by taking the same barrier shifts as the conditional double barrier the trader just prices the unconditional double barrier slightly conservatively.

Forward Starting Options

A forward starting option is an option which does not start until the so-called forward start date. This also means that the strike price of this option is dependent on the underlying share price at this forward start date. The price of a forward starting option can in most cases be estimated with the equivalent regular option which has the same term. Obviously, to exactly price this forward starting option one would use a Monte Carlo process that simulates all the different paths according to the stochastic variables, calculates the price by adding the payoffs for each path and subsequently divides this sum by the number of generated paths. This would give the exact price of the forward starting option, assuming that the stochastic variable inputs are correct and enough paths are generated. In other words, pricing the forward starting option is easy once one understands how to input the right stochastic variable for the Monte Carlo process. This is exactly what this section will focus on. Moreover, this chapter will show that skew is the variable to be most wary of when pricing a forward starting option.

11.1 FORWARD STARTING AND REGULAR OPTIONS COMPARED

Consider a forward starting 100 % put option on BMW that starts in one month and then has a maturity of 1 year. The fact that the forward starting option has a strike of 100 % means that, in one month, the strike of the forward starting option is set to the prevailing stock price on that date. Suppose that BMW stock is trading at € 40. One can estimate the price of this forward starting option on BMW by pricing a 1 year ‘at the money’ put on BMW. This price will be almost the same as the 100 % one month forward starting option which has a maturity of 1 year from the date of the forward start date. Obviously, there is a slight difference in price because of the fact that the buyer of a forward starting option only holds an option in one month but has to pay for it up-front. This means that one has to discount the price of the regular 1 year option by one month to get the price of the 1 month forward starting option. Apart from this

interest rate differential there are two other potential pitfalls in deriving the price of a forward starting option from the equivalent regular option. The **first** pitfall is that there might be a dividend in the 1 month up to the forward start date and not the year after. This means that the forward starting option has one dividend less than the regular option. Therefore, to derive the price of the forward starting option from the regular option, one has to price the regular option without the dividend that occurs in the month up to the forward start date. Otherwise, the price of the forward starting put option would be overestimated. The **second** pitfall is that, depending on the strike level in one month, one might not be using the correct implied volatility to price the forward starting option. Indeed, it was shown in Chapter 5 that the lower the strike price the higher the implied volatility. This means that if BMW's share price goes down in the month up to the forward start date, the regular option should have been priced at a higher implied volatility in order to derive the price of the forward starting option from the equivalent regular option. However, if the stock price were to go up during that month, the regular option should be priced at a lower implied volatility to derive the correct price for the forward starting option. The problem of skew can be solved by going long delta if one buys the forward starting option and going short delta if one sells the forward starting option. The next sub-section shows how to hedge the skew risk during the forward start period for the forward starting option. The last sub-section deals with another distinct skew feature that strongly affects the price of a forward starting option, namely the term structure of skew.

11.2 HEDGING THE SKEW DELTA OF THE FORWARD START OPTION

The previous sub-section showed that the price of any forward starting option is dependent on the underlying share price at the forward start date, as the strike determines what implied volatility each option should be priced at. However, the strike price only becomes known on the forward start date. Therefore, the implied volatility is always an unknown in pricing the forward starting option. Whether a trader prices the forward start option with a Monte Carlo process or by deriving it from the equivalent regular option, he always uses the current share price as a reference for the implied volatility of the forward starting option. This means that, if the trader buys a forward starting option, he loses money if the share price goes up and, if he sells a forward starting option, he loses

money when the share price goes down. The trader can easily hedge this exposure by buying or selling the underlying shares. The following example shows how this is done in practice.

Consider the 100 % forward starting put option on BMW, where the strike price is set in 1 month to then have a maturity of 1 year. Suppose BMW is trading at €40 and the ATM implied volatility is 20 %. Furthermore the vega of this forward starting option is 0.30 % and the trader buys 1 million of these forward starting puts. Also, because of the effect of skew the ‘at the money’ implied volatility goes down by 0.15 % for a 1 % (€0.40) move up in BMW’s share price. Since the trader’s best guess for the ATM implied volatility in one month’s time is the current implied volatility of 20 %, the trader exposes himself to an increase in BMW’s share price. Luckily the trader can easily hedge himself to this sensitivity in share price by buying shares. The number of shares the trader needs to buy can be calculated as follows. For every €0.40 increase in BMW’s share price the trader loses 1 million multiplied by €40 multiplied by 0.30 % multiplied by 0.15 %, which equals a loss of €180 per €0.4 increase in share price. This can be hedged by buying $180/0.4 = 450$ BMW shares.

11.3 THE FORWARD START OPTION AND THE SKEW TERM STRUCTURE

It is important to recognise that the term structure of skew strongly impacts the price of a forward starting option. This has been overlooked by numerous investment banks for a prolonged period and has therefore resulted in large losses on their side. The following example clearly shows the risks of the skew term structure in pricing a forward starting option.

Consider a trader who is selling a three year 70 % put option on BMW that is forward starting in 2 years. If the trader blindly prices this option off his implied volatility curve he will look at the 3 year point on his volatility surface, which has a much lower skew than the skew for the 1 year maturity. However, a three year option that is forward starting in 2 years is effectively a 1 year option and should therefore be priced with the skew of the 1 year maturity. To make this more intuitive, suppose that the 3 year ‘at the money’ implied volatility is 24, the 1 year skew is 1.5 % per 10 % and the 3 year skew is 0.5 % per 10 % change in strike price. This means that the volatility surface prescribes an implied volatility of 25.5 % for the 70 % 3 year put forward starting in 2 years. However,

since the option is forward starting in 2 years, the forward starting 70 % put should be priced with an implied volatility using the 1 year skew, which takes the implied volatility to 28.5 %. In other words, the implied volatility one should use to price a forward starting option is that derived from the 'at the money'¹ implied volatility point at expiry of the forward starting option and, in case of a downside strike, one should adjust this volatility by the prescribed skew of the maturity of the real term of the forward starting option, i.e. the time between expiration and the forward start date.

11.4 ANALYTICALLY SHORT SKEW BUT DYNAMICALLY NO SKEW EXPOSURE

The paradox with a forward starting option is that a short position in the 70 % put on BMW, expiring in 3 years and forward starting in 2 years, is analytically short skew but has no skew exposure dynamically during the forward starting period. Dynamically short skew means that, if the share price goes down, the position gets shorter vega and if the share price goes up, the position gets longer vega. A short position in a regular 3 year 70 % put option on BMW is dynamically short skew as this position gets shorter vega with the share price moving down. The reason that a forward starting option is not short skew dynamically during the forward starting period is that, if the share price moves down the strike moves down as well. Thus, for a regular short downside option, a downward moving share price means that the share price moves closer to the strike price and therefore the position gets shorter vega. For a short forward starting option, a downward moving share price does not change anything for the vega position as the percentage strike stays the same. However, both a short position in the regular and the forward starting option are short skew analytically. Analytically short skew means that if the skew increases, i.e. the difference in implied volatility increases for a 10 % difference in strike, the position loses money, i.e. the downside option becomes more valuable and therefore being short this downside option costs money. This is obviously the case for a short position in either the regular option or the forward starting option.

The fact that a forward starting option is analytically short skew but not dynamically makes it extremely difficult to hedge one's exposure to

¹ One should really derive it from the 'at the money forward' volatility. This is the implied volatility at the forward level rather than the current 'at the money' level.

forward starting options. This is because one can only hedge with regular options, which are both analytically as well as dynamically short skew. This means that if one hedges a short 70 % forward starting option by buying regular options, one needs to role up the strike if the share price is going up and role down the strike if the share price is going down. To make matters worse, sub-section 11.3 showed that a three year 70 % put forward starting in 2 years is analytically mainly exposed to the one year skew. That would mean that one has to hedge a short position in this option by buying a 1 year 70 % put. However, as time passes one would need to role forward this hedge to continue to be hedged on the 1 year skew. This means that it is extremely difficult to analytically hedge forward starting options against skew. First of all, one would need to hedge with options having a maturity equal to the term of the forward starting period, i.e. the maturity minus the forward start date, and continue to role this maturity forward as time lapses. On top of that, one would need to role the strike up if the share price moved up and role the strike down if the share price moved down.

11.5 FORWARD STARTING GREEKS

The Greeks of a forward starting option are similar to the equivalent option with a time to maturity equal to the period between forward start date and expiration. However, delta,² gamma and theta will only kick in once the forward starting option has actually struck, i.e. on the forward start date. Only then will delta, gamma and theta get a value. For vega this is not the case and it does have a value regardless of whether the forward starting option has struck. In other words, even if the forward starting option has not struck yet it is already exposed to changes in implied volatility. The vega of a forward starting option is similar to an option with a time to maturity equal to the period between forward start date and expiration of the forward starting option. The only main difference is that the level of volatility used to determine this vega should equal the implied volatility of the actual expiration of the forward starting option rather than the level associated with the actual term of the forward starting option.

² This does not take the skew delta into account.

Ladder Options

Ladder options are designed for investors who want to get exposure to the upside of a stock while at the same time locking in the performance of the stock if it ever goes above certain levels. This type of option is particularly popular among retail investors and is typically structured as a capital guaranteed note with unlimited upside participation and the added advantage that a certain performance is guaranteed once the stock goes above a certain level. The following sub-section gives an example of a ladder option.

12.1 EXAMPLE: LADDER OPTION

Consider an investor who wants to go long BMW stock but wants to make sure that if the stock ever goes above 105 % or 110 % he receives at least 5 % or 10 % respectively. The investor decides to buy a 1 year ladder option which has full participation in the upside of BMW stock and assures a payoff of the greater between the performance of BMW and 5 % if the stock ever goes above 105 % and the greater of the performance of BMW and 10 % if the stock ever goes above 110 %. Mathematically the payoff at maturity looks like:

$$C_T^{Ladder} = \max \left[X, \frac{BMW_T - BMW_0}{BMW_0} \right],$$

where

$$X = \begin{cases} 0 \% & \text{If BMW never} > 105 \% \text{ during the one year term of the} \\ & \text{ladder option} \\ 5 \% & \text{If BMW ever} > 105 \% \text{ during the one year term of the} \\ & \text{ladder option} \\ 10 \% & \text{If BMW ever} > 110 \% \text{ during the one year term of the} \\ & \text{ladder option} \end{cases}$$

Although the payoff of the ladder option looks complicated, it is surprisingly easy to price with the knowledge of barrier options with rebates.

12.2 PRICING THE LADDER OPTION

The ladder option can be priced as a series of higher strike knock-out calls with rebates, where the strike of the following knock-out call equals the barrier level of the previous knock-out call and the rebates are equal to the ‘additional’ performance that gets locked in. Obviously, the highest strike call does not have a knock-out level. The ladder option in sub-section 12.1 can be priced by the following set of options:

- Investor is long a 100 % call which knocks out at 105 % to then pay a rebate of 5 %.
- Investor is long a 105 % call which knocks out at 110 % to then pay a rebate of 5 %.
- Investor is long a 110 % call without any knock-out level.

Since, at the barrier level, the rebate of the knock-out call is equal to the intrinsic value,¹ the discontinuity in option price is minimal over the barrier. Because the difference in option price is so small over the barrier the trader barely needs to apply any barrier shift. Moreover, regardless of when the option knocks out during the term of the option, the trader, who is short this knock-out call, would either need to sell shares over the barrier or not do anything at all on the underlying stock.² This means that his barrier risk is basically nonexistent as he can always put an offer in at 1 cent below the barrier for the amount of shares that he would need to sell. This means that the most the trader can lose is 1 cent on the amount of shares he would need to sell.

¹ The time value of the knock-out call will be minimal at the barrier as almost any path of stock price will predict the stock to go through the barrier.

² The ladder structure as a whole might force the trader to buy shares over the barrier because the trader is short gamma from the higher strike calls. However, this short gamma risk is already priced in and therefore does not need to be used as an excuse to apply a barrier shift.

Lookback Options

The principle of a lookback option is to give the investor the maximum payoff based on perfect hindsight. The lookback takes two forms. One where the settlement price of the option is chosen with perfect hindsight of the stock's path during the term of the option and the strike is fixed. And one where the strike is chosen with perfect hindsight and the settlement price is the stock price at maturity of the option. In other words, there are four different types of lookback options:

- **Fixed strike call lookback (Max lookback)**
 - Pays the difference between the strike and the highest stock level during the term of the option.
- **Fixed strike put lookback (Min lookback)**
 - Pays the difference between the strike and the lowest stock level during the term of the option.
- **Floating strike call lookback**
 - Pays the difference between the stock price at maturity and the lowest stock level during the term of the option.
- **Floating strike put lookback**
 - Pays the difference between the stock price at maturity and the highest stock level during the term of the option.

13.1 PRICING AND GAMMA PROFILE OF FIXED STRIKE LOOKBACK OPTIONS

A fixed strike lookback option will typically be priced with a Monte Carlo process. However, the fixed strike lookback can also be priced as a ladder option where the lock in levels are one tick apart. For example, one can price a fixed strike €45 call lookback on BMW, where the tick value is 1 cent, in the following way:

- Investor is long a €45 call which knocks out at €45.01 to then pay a rebate of €0.01.
- Investor is long a €45.01 call which knocks out at €45.02 to then pay a rebate of €0.01.

- Investor is long a € 45.02 call which knocks out at € 45.03 to then pay a rebate of € 0.01.

⋮

- Investor is long a € 50 call which knocks out at € 50.01 to then pay a rebate of € 0.01.

⋮

Obviously a lookback put option can be priced as a ladder option equivalent to a series of down-and-out puts with a cent rebate on each.

The gamma profile of a Max lookback option becomes intuitive when viewing it as a ladder option. Indeed, as long as the stock goes up there will be gamma on the lookback option and the gamma will decrease quickly when the stock goes down, as the options below have already knocked out and therefore have no gamma on them any more. Therefore the gamma of a Max lookback is highest when the stock is at its high, taken from the start of the option term, and increases over time as long as the stock is at or close to its high. For the Min lookback it is exactly the opposite. The gamma is highest when the stock is at its low, taken from the start of the option trade.

Therefore the gamma and vega exposure of a Max lookback option can be hedged by selling a series of plain vanilla upside calls that are, say, 5 % apart. The Min lookback would be hedged by selling a series of downside plain vanilla puts.

13.2 PRICING AND RISK OF A FLOATING STRIKE LOOKBACK OPTION

Like the fixed strike lookback, the floating strike lookback would always be priced as a Monte Carlo process. However, unlike the fixed strike lookback, it is virtually impossible to replicate the floating strike lookback as a set of knock-out options. Since the strike of a floating strike lookback call resets if the share price goes down, one can see that the closest way to replicate the floating strike lookback call option is as a set of forward starting options. However, the situation occurs that these forward starting options are only struck if the share price is lower than the share price of all the previous days since the inception of the lookback option. Unfortunately, a knock-in feature that is dependent on the future share price cannot be captured by the forward starting option.

Nonetheless, to understand the risk of a floating strike lookback option it is very useful to think in terms of forward starting options. It basically means that a floating strike lookback call option is effectively exposed to forward starting risk for as long as the share price goes down. A floating strike lookback put is exposed to forward starting risk for as long as the share price goes up. This means that the floating strike lookback option has a degree of risk associated with forward starting options. To make this more clear consider the following example.

Consider a trader who sells a one year floating strike lookback call option. The trader is obviously short skew, because as long as the share price goes down forward starting options are struck and he therefore gets shorter vega. On top of that, the forward starting option that is struck in this case is on a higher implied volatility than the 'at the money' implied volatility. If the share price goes up he will in turn get longer vega. Unlike a regular forward starting option (see sub-section 11.2), the trader would not go short delta against the lookback option since it is not certain that any forward starting option will get struck. For example, if the share price never goes below the initial share price, no forward starting option gets struck. For a floating strike lookback put option the exact opposite holds. Consider a trader who sells a floating strike lookback put option. In this case the trader is long skew, as he goes shorter vega if the share price goes up because new options get struck.

Cliquet options are options where the strike price of the option potentially resets at predetermined points in time. The most common cliquet option is the *ratchet* option. The ratchet option is an option where the strike price resets at predetermined points in time while at the same time locking in the performance of the previous period. This chapter discusses the pricing and risk of the ratchet option and will show that it is nothing more than a set of forward starting options.

14.1 THE RATCHET OPTION

Consider a 3 year 100 % ratchet call option on BMW that locks in the performance of BMW in each year. In other words, the strike price of the call resets every year to the prevailing stock level of BMW while at the same time locking in the performance of that year. Suppose that the initial stock level of BMW is $S_0 = \text{€}40$, after 1 year it is $S_1 = \text{€}45$, after the second year it is $S_2 = \text{€}50$ and after the third year BMW stock is trading at $S_3 = \text{€}45$. The contribution of each year to the payoff at maturity of this ratchet option can be summarised by

- the contribution of year 1 is $\max [S_1 - S_0, 0] = \text{€}5$;
- the contribution of year 2 is $\max [S_2 - S_1, 0] = \text{€}5$;
- the contribution of year 3 is $\max [S_3 - S_2, 0] = \text{€}0$.

The payoff formula for each year clearly shows that the strike price of the ratchet option resets each year to the prevailing stock price, i.e. the initial strike is S_0 , after one year the strike resets to S_1 and after two years the strike resets to S_2 . The full payoff of an $X\%$ ratchet call option at maturity with $n + 1$ observations and therefore n resets is

$$\sum_{i=0}^n \max \left[S_{i+1} - \frac{X}{100} S_i, 0 \right], \quad (14.1)$$

where S_i is the stock level after the i -th period. Obviously, the payoff at maturity for an X % ratchet put option looks like

$$\sum_{i=0}^n \max \left[\frac{X}{100} S_i - S_{i+1}, 0 \right]. \quad (14.2)$$

Since the strike price resets every year and the performance over the previous period gets locked in, the ratchet option can be priced as a set of forward starting options. Obviously, since the payoff occurs at maturity the value of each forward starting option needs to be discounted by the period from expiry of each forward starting option till the maturity of the ratchet option. In other words, the 3 year ratchet call option on BMW can be priced by adding the following forward starting options:

1. A one year at the money call option with a delayed payout date of two years.
2. A forward starting at the money call option with a forward start date in one year's time and maturing in two years' time with a delayed payout date of one year.
3. A forward starting at the money call option with a forward start date in two years' time and maturing in three years' time.

In a perfectly normally distributed world where there is no skew and BMW is a non-dividend paying stock, the price of the 100 % ratchet call option would be

$$\begin{aligned} & \sum_{i=0}^2 e^{-r \cdot i} [S_0 N(d_1) - S_0 e^{-r} N(d_2)] e^{-(3-(i+1))r} = \\ & \sum_{i=0}^2 e^{-2r} [S_0 N(d_1) - S_0 e^{-r} N(d_2)], \end{aligned} \quad (14.3)$$

where r is expressed on an annualised basis, i.e. the unit of time is one year. Equation 14.3 is derived from the Black–Scholes formula for a call option

$$c_t = S_t N(d_1) - K e^{-r(T-t)} N(d_2), \quad (14.4)$$

where

$$d_1 = \frac{\ln \left(\frac{S_t}{K} \right) + \left(r + \frac{1}{2} \sigma^2 \right) (T - t)}{\sigma \sqrt{T - t}}, \quad (14.5)$$

$$d_2 = d_1 - \sigma \sqrt{T - t}. \quad (14.6)$$

In equation 14.3, $e^{-r \cdot i}$ represents the discounting in order to get the price of the forward starting option and $e^{-(3-(i+1))r}$ takes care of the discounting as a result of the delayed payoff date. Also, in equation 14.3, d_1 and d_2 are calculated with $T - t$ being equal to 1 and $K = S_0$.

In practice skew does exist and therefore a ratchet option would always be priced as a Monte Carlo process.

14.2 RISKS OF A RATCHET OPTION

Since a ratchet option can be priced as a set of forward starting options, the risks involved are very similar to those of a forward starting option. Therefore, the main risk of a ratchet option manifests itself in the skew, see Chapter 11. For example, numerous investment banks have lost money by selling ratchet call spreads. They were, for example, selling 100/130 % ratchet call options. An $n + 1$ year 100/130 % ratchet call spread with yearly resets has the following payoff:

$$\sum_{i=0}^n \max [\min \{S_{i+1} - S_i, 0.3 \cdot S_i\}, 0], \quad (14.7)$$

where S_i is the stock level after the i -th period. However, selling call spreads means selling skew, as the bank sells the lower strike and buys the higher strike. This obviously means that the bank has to take into account all the issues with forward starting options discussed in Chapter 11. For a long period several banks have been pricing these forward starting options using the skew of the specific maturity rather than the skew associated with the maturity equal to the term of the forward starting option, i.e. the maturity minus the forward start date, see Chapter 11, sub-section 11.3. This way of pricing has resulted in large losses for some banks.

Reverse Convertibles

A reverse convertible is a structure where an investor buys a bond and sells an 'at the money' put option to enhance the coupon. Therefore the coupon of a reverse convertible is higher than the coupon on a regular bond. Technically, the reverse convertible in itself is not an exotic option as the structure is nothing more than a bond plus a plain vanilla put option. However, it is imperative to have a good understanding of the reverse convertible in order to understand more exotic structures like the autocallable. Moreover, there are many different iterations to the reverse convertible which can turn a plain vanilla reverse convertible into an exotic reverse convertible. For example, instead of a regular reverse convertible where the investor buys a bond and sells an 'at the money' put, the investor can decide to sell a down-and-in put which gives the investor additional protection. Lastly, the reverse convertible is an extremely popular bullish retail structure and whoever wants to understand exotic options needs to be familiar with such an important structure as the reverse convertible.

15.1 EXAMPLE: KNOCK-IN REVERSE CONVERTIBLE

A very popular iteration to the regular reverse convertible is a structure where the investor buys a bond and sells a down-and-in put to enhance the coupon on the whole structure – a knock-in reverse convertible. Sub-section 10.1 shows that the seller of a down-and-in put enjoys additional protection over the seller of an 'at the money' put. Because the investor receives additional protection, the knock-in reverse convertible pays a slightly lower coupon than the regular reverse convertible. Nonetheless, the knock-in reverse convertible still pays a higher coupon than a regular bond as a compensation for taking the risk of selling a down-and-in put.

Consider a 2 year knock-in reverse convertible on BMW where the investor sells a 100/80 % down-and-in put to enhance the coupon on the structure. In other words, the investor buys a 2 year bond and sells a 2 year 100/80 % down-and-in put and in return the investor receives a semi-annual coupon. The investor buys the full knock-in reverse convertible

structure at 100 %. That is to say, a knock-in reverse convertible works along the same principles as a regular bond. Whoever buys the reverse convertible pays the full notional up-front and throughout the life of the reverse convertible the investor receives semi-annual coupons. In order for the reverse convertible to work it is important that the notional on the down-and-in put is the same as the notional on the bond. This can be established by having the notional divided by the strike number of down-and-in puts in the reverse convertible. This also ensures that if the down-and-in put has knocked in **and** it is in the money at expiry, the issuer of the reverse convertible can simply exercise the down-and-in put for which the issuer should be getting paid an amount equal to the notional that the issuer owes the investor already through the bond part of the reverse convertible. Thus, if the issuer exercises the down-and-in put there is no exchange of cash at expiry. This also shows that a reverse convertible is not a riskless structure to the investor. If, for example, BMW is trading at 70 % of its initial value, the buyer of the reverse convertible loses 30 % of his invested notional as he receives shares, through the assignment of the down-and-in put, that are worth 30 % less. To summarise, at expiry there are three different payout scenarios:

{	100 % in cash	If BMW never < 80 % during the 2 year term of the knock-in reverse convertible
	100 % in cash	If BMW at expiry > 100 % regardless of whether BMW was ever < 80 % during the 2 year term of the knock-in reverse convertible
	$\frac{\text{notional}}{\text{strike}}$ number of shares	If BMW ever < 80 % during the 2 year term of the knock-in reverse convertible and BMW < 100 % at expiry

In the third case, the buyer of the reverse convertible obviously loses money. Instead of seeing his invested notional returned on the bond he receives $\frac{\text{notional}}{\text{strike}}$ number of BMW shares, which he effectively buys at 100 % and are worth less than 100 % at expiry. De facto the investor loses the percentage decline of the BMW share price on the notional invested.

15.2 PRICING THE KNOCK-IN REVERSE CONVERTIBLE

Pricing a reverse convertible typically means that one solves for the coupon that makes the value or re-offer of the reverse convertible 100 %. Consider the two year 100/80 % knock-in reverse convertible on BMW from sub-section 15.1. Suppose that the semi-annual coupon on a regular bond is equal to 1.75 % and also that the interest rate is constant in every maturity. Lastly, suppose that a 2 year 100/77 % down-and-in put on BMW is worth 6.7 %. Note that the issuer would use a 3 % barrier shift to price a 100/80 % down-and-in put, as discussed in Chapter 10. Now it is easy to determine the semi-annual coupon that makes the knock-in reverse convertible worth 100 %. Indeed, since the semi-annual coupon on a regular bond is 1.75 %, the semi-annual discount factors to four decimal places are:

- the 6 month discount factor is $\frac{1}{1.0175} = 0.9828$
- the 1 year discount factor is $\frac{1}{(1.0175)^2} = 0.9659$
- the 18 month discount factor is $\frac{1}{(1.0175)^3} = 0.9493$
- the 2 year discount factor is $\frac{1}{(1.0175)^4} = 0.9330$

the semi-annual coupon, X , for the two year 100/80 % knock-in reverse convertible on BMW can now be calculated as:

$$\begin{aligned}
 0.9828X + 0.9659X + 0.9493X + 0.9330X &= 1 - (0.9330 - 0.067) \\
 &\Downarrow \\
 3.831X &= 0.134 \\
 &\Downarrow \\
 X &= 0.035
 \end{aligned}$$

In other words, a cash semi-annual coupon payment of 3.5 % makes the 2 year 100/80 % knock-in reverse convertible on BMW worth 100 %.

15.3 MARKET CONDITIONS FOR MOST ATTRACTIVE COUPON

It is important to understand under which circumstances the coupon on a reverse convertible looks the most attractive. In other words, during which market conditions one can expect the highest demand for reverse convertibles. There are two inputs that affect the coupon on a reverse

convertible. The first one is the premium for the put or down-and-in put. The higher the premium on the put or down-and-in put the higher the coupon. The two main contributors to a higher put premium are the implied volatility and the dividend yield. The higher the implied volatility or dividend yield the higher the put premium. Secondly, the lower the discount factors the higher the coupon. Discount factors have an inverse relationship to the interest rate. Therefore, the higher the interest rates the higher the coupon.¹ To summarise, there are three factors that impact the magnitude of the coupon:

- The higher the interest rates the higher the coupon on the reverse convertible.
- The higher the implied volatility the higher the coupon on the reverse convertible.
- The higher the dividend yield the higher the coupon on the reverse convertible.

15.4 HEDGING THE REVERSE CONVERTIBLE

Hedging a reverse convertible consists of two hedges. One is the volatility hedge and the second one is the interest rate hedge. There is no exact volatility hedge for the down-and-in put. However, a good proxy hedge does exist. For example, the two year 100/80 % down-and-in put on BMW would typically be hedged with a two year 90 % plain vanilla put. The reason that one would hedge the two year 100/80 % down-and-in put with a downside put, i.e. 90 %, is because a down-and-in put is long skew, which makes the effective implied volatility at which one buys a 100/80 % down-and-in put higher than the 'at the money' implied volatility. In order to hedge this effectively, the buyer of a down-and-in put needs to sell some downside options for which he receives an implied volatility that is slightly higher than the 'at the money' implied volatility. The number of 2 year 90 % put options the buyer of the down-and-in put needs to sell can easily be determined by the vega exposure on the down-and-in put and solving for the number of 2 year 90 % puts that flattens this vega exposure.

¹ Higher interest rates have a dampening effect on the put premium. This is because a long put is hedged by buying shares on which the owner of the put incurs financing. If the interest rates go up, the financing of the hedge becomes more expensive and therefore the buyer is not prepared to pay as much for the put. Since an 'at the money' put has a delta less than 1, higher interest rates have a bigger impact on the bond part than on the put option part of the reverse convertible. Therefore, the result of higher interest rates is that the coupon looks more attractive.

Hedging any interest rate exposure on a reverse convertible is rather simple. Suppose the issuer of a € 10 million 2 year 100/80 % reverse convertible on BMW wants to hedge his interest rate exposure. Suppose also that the delta on the down-and-in put is 0.4. This means that the issuer will buy € 4 million worth of BMW stock on which he incurs financing. On the € 10 million short bond he actually receives interest and therefore the short bond and the delta hedge work in opposite directions in terms of interest rate exposures. In other words, the issuer of the reverse convertible is long floating interest rates on a notional of € 6 million. Since it is a 2 year reverse convertible, the issuer is exposed to the 2 year part of the interest rate curve and will therefore enter into a 2 year swap where he pays floating and receives fixed on a notional of € 6 million in order to hedge his interest rate exposure.

DVO1, the price sensitivity of any structure to 1 basis point,² bp, change in yield curve (the annualised interest rate goes up by 1 bp for every maturity), is typically used as an interest rate risk measure. Since the issuer of the reverse convertible has a 2 year long interest rate exposure and assuming an annual interest coupon payment, a 1 bp increase for the 1 year maturity results in a 0.9659 of a bp profit and a 1 bp increase for the 2 year maturity results in a 0.9330 of a bp profit. This means that the *DVO1* for the issuer of the reverse convertible is equal to

$$€ 6 \text{ million} \times \frac{0.9659 + 0.9330}{10\,000} = + € 1139.34$$

² 1 bp = $\frac{1}{100}$ of 1 %.

An autocallable is a note people can invest in that only pays a coupon if the underlying index is above or below a certain level and the note automatically redeems early if it breaches a second hurdle level, which can be the same or different to the one that determines whether the investor gets paid a coupon. The reason that the note redeems early is that, once the underlying index reaches a certain level, the investor would prefer to invest his money elsewhere rather than to hold his money in a note that pays a fixed coupon, conditional on a certain level of underlying index. To get an enhanced coupon the investor typically also sells an option, which is then embedded in the note. The following sub-section will give an example of an autocallable note.

16.1 EXAMPLE: AUTOCALLABLE REVERSE CONVERTIBLE

Consider an investor who has a range bound view on BMW and believes that in the next 2 years the stock will not drop below 90 % of its value and will not get above 110 % of BMW's current value. He also believes that, if BMW does get above 110 % of its current value, it is a bullish signal and would like to invest money in BMW stock. The autocallable would be the perfect strategy for this investor, where the note autocalls when BMW reaches 110 % of its current value so that he can invest the money of the note in BMW stock directly. The note will pay a semi-annual coupon for as long as BMW stock is not below 90 % of its current value at these semi-annual observation points. Lastly, the investor wants to enhance his coupon by selling a down-and-in put with a strike of 100 % and a continuous barrier of 80 %, which means that his investment is capital guaranteed provided BMW never drops by more than 20 % at any time during the 2 year term of the note. By selling this down-and-in put the investor manages to enhance his coupon to a 4 % semi-annual coupon. Suppose this investor invests 10 million euros in this autocallable reverse convertible. The exact terms of this

note are as follows:

- **Coupons**

- If at any semi-annual observation point, i , during the term of the note

$$BMW_i > 90 \% \cdot BMW_0,$$

the investor gets paid a coupon of 4 % on his investment at this particular semi-annual observation point, i.e. 400 thousand euros.

- If at any semi-annual observation point

$$BMW_i < 90 \% \cdot BMW_0,$$

the investor does not get a coupon at this particular semi-annual observation point.

- **Early redemption**

- If at any of the first three semi-annual observation points

$$BMW_i > 110 \% \cdot BMW_0,$$

the note redeems early and the investor will get his full investment back, i.e. 10 million euros. Because of this redemption the investor obviously loses his right to coupon payments.

- **Redemption at maturity**

- If BMW's share price has never dropped below 80 % of its initial value during the two year term of the note, the note redeems at 100 %, i.e. 10 million euros.
- If BMW's share price has at any point during the 2 year term of the note dipped below 80 % of its initial value, the note redeems at

$$100 \% - \max \left[\frac{BMW_0 - BMW_T}{BMW_0}, 0 \right].$$

It is not straightforward to see whether this conditional autocallable reverse convertible pays a higher or lower coupon than the regular reverse convertible discussed in Chapter 15. On the one hand, the coupon of the above autocallable is conditional and should therefore have a higher coupon than the coupon of the regular reverse convertible, which is unconditional. On the other hand, the autocallable redeems early if, at one of the semi-annual observation dates, BMW's stock price is above 110 %, in which case the trader loses a down-and-in put. This would have a dampening effect on the coupon of the above autocallable. However, taking both aspects into account one can assume that the contingency of the coupon has a larger increasing impact on the coupon than the downward pressure on the coupon from potentially losing the down-and-in put.

16.2 PRICING THE AUTOCALLABLE

The pricing of the autocallable discussed in sub-section 16.1 is relatively easy with the knowledge of digital and barrier options. After all, the conditional coupons are nothing more than a strip of semi-annual 90 % strike digital options which knock out at 110 %, and pay a coupon of 4 %. The *early redemption feature* can be priced by a worthless option that knocks out at 110 % to then pay a rebate of 100 %. This worthless option is typically taken to be a zero strike put. The barriers are only live at the semi-annual observation dates. In other words, the barriers can only be breached at these semi-annual observation dates. The entire structure can therefore be priced as:

- **Coupons**

- The trader sells a 1 time geared 86/90 % European call spread expiring after six months with no knock-out feature, as the investor will still get the 4 % coupon if the note knocks out at the first semi-annual observation date.

Unlike the reverse convertible, this way of pricing the autocallable reverse convertible does not assume that a coupon actually gets paid out in cash at each semi-annual observation date, but is embedded in the pricing of the call spread. The call spread is a conservative way of pricing a digital option that gives a payout of 4 % if BMW's stock price is above 90 % and pays nothing if BMW's stock price is below 90 % at each semi-annual observation date.

- The trader sells a 1 time geared 86/90 % European call spread expiring after one year, with a knock-out barrier of 110 %. The barrier is live only at the first semi-annual date, as the coupon needs to be paid at the second semi-annual observation date, regardless of whether the structure knocks out at this date.
- The trader sells a 1 time geared 86/90 % European call spread expiring after 1.5 years, with a knock-out barrier of 110 %. The barrier is live only at the first and second semi-annual dates.
- The trader sells a 1 time geared 86/90 % European call spread expiring after 2 years, with a knock-out barrier of 110 %. The barrier is only live at the first three semi-annual dates. The barrier is not live at the fourth semi-annual observation date because the note can only redeem at the first three semi-annual observation dates.

- **Early redemption**

- The trader sells a zero strike knock-out put with a barrier of 110 % for the first three semi-annual dates and a 0 % barrier for the fourth

semi-annual observation date. The knock-out put can only knock out at one of the semi-annual observation dates and pays a rebate of 100 %. The barrier for the fourth semi-annual observation date ensures that the client sees his investment returned after the two year period.

- **Financing through down-and-in put**

- The trader buys a 100 % put knocking in at 80 % and knocking out at 110 %. The reason that this option knocks out at 110 % is that the down-and-in put is part of the reverse convertible autocallable note, which knocks out at 110 %. Obviously the trader would price the downside barrier with a barrier shift, i.e. pricing it as a 100 % put with a 77 % knock in. It is important to recognise that the downside barrier at 80 % can knock in continuously, i.e. throughout the full life of the autocallable, whereas the upside barrier at 110 % can only knock out at one of the three semi-annual observation dates.

When determining how to shift the barrier at 110 % for the different semi-annual observation dates, one first has to realise that the different digital options, the zero strike put and the down-and-in put are communicating options when it comes to knocking out. The best way to see this is to imagine BMW's stock price at 109 % just before the third semi-annual observation date. In other words, the autocallable note is close to knocking out. In this case it could well be advantageous to the trader if the structure knocks out. Firstly, this is because the down-and-in put is very unlikely to knock in during the next 6 months, making the down-and-in put almost worthless at that point. Secondly, if the structure knocks out at the third semi-annual observation date, the trader will have to pay 100 % 6 months early, which works to his disadvantage. However, it would also mean that he would not have to pay the fourth coupon, which works to the trader's advantage. Assuming that the rate of interest has not moved, these two effects will work to the trader's advantage, as the coupon plus the present value of 100 % in 6 months is more than 100 % now. This is because the coupon is artificially high as a result of the financing through the down-and-in put within the note. Therefore, the above shows that, if the the knock-in put is close to worthless and the interest rate has been stable or has moved down, it is advantageous to the trader if the note knocks out. In this case the trader obviously shifts the barrier upwards. However, in most of the cases it is disadvantageous to the trader if the note redeems early, as the consequence is that he loses a down-and-in put which he has paid for and therefore shifts the barrier downwards at the 110 % level.

16.3 AUTOCALLABLE PRICING WITHOUT CONDITIONAL COUPON

Another, and probably more common, class of autocallables is an autocallable where the coupon is unconditional but still autocalls at a certain level. Typically, this autocallable with an unconditional coupon is structured such that the coupons are not paid running but only when the structure knocks out, or if it never knocks out the coupons are paid at maturity. Since the coupons are only paid at the observation dates when the autocallable actually knocks out, the coupon that is paid at maturity or when it autocalls is the number of observation dates till knock out multiplied by the agreed coupon. The following example clarifies the autocallable with unconditional coupon.

Consider the autocallable example in sub-section 16.2, only then the coupon is paid regardless of any stock level. One would expect this coupon to be lower than 4 % semi-annually. It would most likely also be lower than the coupon of the regular reverse convertible, as there is always the risk to the trader of losing the down-and-in put if the 110 % barrier is breached. Assuming that the fair semi-annual coupon is 3 % for this unconditional autocallable on BMW, the 2 year unconditional autocallable on BMW would be structured as follows:

- **Early redemption plus coupons**

- The trader sells a zero strike knock-out put with a barrier of 110 % for the first three semi-annual dates and a 0 % barrier for the fourth semi-annual observation date. The knock-out put can only knock out at one of the semi-annual observation dates and pays a rebate of
 1. The rebate at the first semi-annual observation date is 100 % plus one times 3 %, equals a rebate of 103 %
 2. The rebate at the second semi-annual observation date is 100 % plus two times 3 %, equals a rebate of 106 %
 3. The rebate at the third semi-annual observation date is 100 % plus three times 3 %, equals a rebate of 109 %
 4. The rebate at the fourth and last semi-annual observation date is 100 % plus four times 3 %, equals a rebate of 112 %
 A barrier of 0 % for the fourth semi-annual observation date ensures that the client sees his investment returned after the 2 year period.

- **Financing through down-and-in put**

- The trader buys a 100 % put knocking in at 80 % and knocking out at 110 %. The reason this option knocks out at 110 % is that the down-and-in put is part of the reverse convertible autocallable note, which knocks out at 110 %. Obviously the trader would price the

downside barrier with a barrier shift, i.e. pricing it as a 100 % put with a 77 % knock in. It is important to recognize that the downside barrier at 80 % can knock in continuously, i.e. throughout the full life of the autocallable, whereas the upside barrier at 110 % can only knock out at one of the three semi-annual observation dates.

16.4 INTEREST/EQUITY CORRELATION WITHIN THE AUTOCALLABLE

When pricing the autocallable in sub-section 16.2 there was no mention of an interest/equity correlation. However, if one wants to price the autocallable accurately this interest/equity correlation should be taken into account. To see this, consider a purely anti-correlated interest/equity correlation. In other words, there is a lognormal correlation of -1 . This obviously has an impact on the pricing of the autocallable. Indeed, in this case BMW's share price goes up whenever interest rates are going down. This means that, when BMW's share price reaches 110 %, the interest rate is lower than at inception of the trade and it is therefore more advantageous to the trader when the autocallable knocks out than when interest rates have not moved at all.¹ For that reason, the trader can give a higher coupon when he takes this interest/equity correlation of -1 into account. The opposite obviously holds when the interest/equity correlation is perfectly correlated, i.e. 1.

The reason that no interest/equity correlation was taken into account when pricing the autocallable in sub-section 16.2 is that this correlation is very hard to measure and for short term maturities this correlation always hovers around zero.² For that reason traders typically price the autocallable as in sub-section 16.2 and increase or lower the coupon depending on their view on the interest/equity correlation during the term of the note.

¹ The reason that it is advantageous to the trader if the autocallable knocks out if interest rates are lower is that the trader effectively borrows money through the autocallable and pays a fixed coupon for it. When the structure knocks out and interest rates are lower he can refinance himself cheaper.

² Over very long periods of time equity and interest rates have a negative correlation. However, over a 2 year period, equity and interest rates might very well be positively correlated. For example, the start of a rate hike cycle after a recession typically features both equity and interest rates going up.

Callable and Puttable Reverse Convertibles

A callable reverse convertible is a reverse convertible that can be called by the issuer (the seller) at his choice at any of the pre-defined observation dates during the term of the callable reverse convertible. This means that the issuer holds an option as to whether he wants to keep the short reverse convertible position till maturity or unwind it early at the observation date of his choosing. Because the issuer has this additional option the buyer of this callable reverse convertible is compensated by means of a higher coupon than the coupon on a regular reverse convertible. The pricing of such a callable reverse convertible can be seen as an optimisation process of an unconditional autocallable, where the coupon is optimised (maximised) versus the barrier level of this autocallable.

A puttable reverse convertible is the opposite of the callable reverse convertible in the sense that it is not the issuer but the buyer who has the choice of unwinding the reverse convertible early at the observation date of his choosing. Again, the puttable can be priced as an optimisation process of an autocallable. However, in this case the coupon is minimised against the barrier level. The fact that the buyer can redeem his investment plus a part of the coupons early, and therefore has access to his capital earlier, results in a slightly lower coupon.

The next two sub-sections discuss the pricing of the callable and puttable reverse convertibles in more detail.

17.1 PRICING THE CALLABLE REVERSE CONVERTIBLE

Consider a 2 year callable reverse convertible on BMW which can be called (unwound) by the issuer at any of the three semi-annual observation dates. The put within the reverse convertible that is used to enhance the coupon is a 100 % put knocking in at 80 %. Chapter 15 showed that a regular reverse convertible on BMW pays a semi-annual coupon of 3.5 % with a re-offer of 100 %. Therefore, one expects the callable reverse convertible to pay a semi-annual coupon of at least 3.5 %. In order

to solve for the semi-annual coupon of this callable reverse convertible, one prices this structure as an autocallable where the barrier levels at the respective semi-annual observation dates are optimised to give the largest coupon for a re-offer or fair value of 100 %. To begin with one can start with a barrier level of 110 % for all the three semi-annual observation dates which, according to Chapter 16, sub-section 16.3, gives a coupon of 3 %. Secondly, one can try to price an autocallable where the first two observation dates have a barrier of 1000 %¹ and the third semi-annual observation date has a barrier of 104 %. This reverse convertible autocallable appears to give a coupon of 3.65 % for a fair value of the whole reverse convertible structure of 100 %. Proceeding in this way one finds that the most optimal barrier levels at which the structure autocalls are the following for the respective semi-annual observation dates, and with this barrier structure the callable reverse convertible pays a semi-annual coupon of 3.8 %:

- The barrier for the first observation date, above which the structure knocks out, is 120 %. This means that, if BMW's share price is higher than 120 % of its initial value at the first observation date, it is optimal for the issuer of the callable reverse convertible to call the structure and repay the buyer 103.8 %. The reason that it is optimal for the issuer to call the full structure and repay 103.8 % is that the remaining value of the down-and-in put has reduced so much that it is more economical to pay the full 100 % early plus the first coupon of 3.8 % and give up the down-and-in put than to keep the structure with the future liability to pay three artificially high² additional coupons but pay the full 100 % in 1.5 years' time.
- The barrier for the second observation date, above which the structure knocks out, is 112 %. This means that, if BMW's share price is higher than 112 % of its initial value at the second observation date, it is optimal for the issuer of the callable reverse convertible to call the structure and repay the buyer 107.6 %. The reason that it is optimal for the issuer to call the full structure and repay 107.6 % is that the remaining value of the down-and-in put has reduced so much that it is more economical to pay the full 100 % early plus the first two coupons of 3.8 % and give up the down-and-in put than to keep the structure with the future liability to pay two artificially high additional coupons but pay the full 100 % 1 year later.

¹ Such a high barrier makes sure that it does not autocall on either of the first two observation dates. Therefore it allows for the optimisation of the barrier on the third semi-annual observation date.

² The coupons are artificially high as they are increased because of the initial value of the down-and-in put.

- The barrier for the third observation date, above which the structure knocks out, is 104 %. This means that, if BMW's share price is higher than 104 % of its initial value at the third observation date, it is optimal for the issuer of the callable reverse convertible to call the structure and repay the buyer 111.4 %. The reason that it is optimal for the issuer to call the full structure and repay 111.4 % is that the remaining value of the down-and-in put has reduced so much that it is more economical to pay the full 100 % early plus the first three coupons of 3.8 % and give up the down-and-in put than to keep the structure with the future liability to pay one additional coupon but pay the full 100 % 6 months later.

This means that the full callable reverse convertible can be priced as:

- **Early redemption plus coupons**

- The trader sells a zero strike knock-out put with a barrier of 120 %, 112 % and 104 % for the first three semi-annual dates respectively and a 0 % barrier for the fourth semi-annual observation date. The knock-out put can only knock out at one of the semi-annual observation dates and pays a rebate of

1. The rebate at the first semi-annual observation date is 100 % plus one times 3.8 %, equals a rebate of 103.8 %
2. The rebate at the second semi-annual observation date is 100 % plus two times 3.8 %, equals a rebate of 107.6 %
3. The rebate at the third semi-annual observation date is 100 % plus three times 3.8 %, equals a rebate of 111.4 %
4. The rebate at the fourth and last semi-annual observation date is 100 % plus four times 3.8 %, equals a rebate of 115.2 %

A barrier of 0 % for the fourth semi-annual observation date ensures that the client sees his investment returned after the 2 year period.

- **Financing through down-and-in put**

- The trader buys a 100 % put knocking in at 80 % and knocking out at respectively 120 %, 112 % and 104 % for the first three semi-annual observation dates. Obviously the trader would price the downside barrier with a barrier shift, i.e. pricing it as a 100 % put with a 77 % knock in. It is important to recognize that the downside barrier at 80 % can knock in continuously, i.e. throughout the full life of the callable reverse convertible, whereas the upside barriers can only knock out at one of the three semi-annual observation dates.

In pricing the callable reverse convertible it is important to recognise two things. First of all the structure only knocks out if the share price

is above a certain level. In other words, it is only optimal for the issuer to call the callable reverse convertible if the share price is above a certain level. This is obvious, since the issuer of the callable reverse convertible owns a down-and-in put that becomes more valuable when the underlying share price is lower and it is therefore less economical for the issuer to give up this down-and-in put. Secondly, one has to observe that the process of finding the most optimal barriers at which the issuer should call the callable reverse convertible is not a random process based on trial and error but can be done in a systematic way with the Newton–Raphson³ method, provided one starts to solve for the optimal barrier with the penultimate observation date and works backwards to the first observation date. In the above example one would use the Newton–Raphson to first find the optimal barrier for the third semi-annual observation date, while keeping the barriers of the first two semi-annual observation dates so high that the structure would never get called on either of these first two observation dates. Knowing the optimal barrier for the third semi-annual observation date, one then goes on to find the optimal barrier for the second semi-annual observation date, again using a Newton–Raphson process. To be able to optimise the barrier for the second semi-annual observation date one has to use an unrealistically high barrier for the first semi-annual observation date. Lastly, knowing the optimal barriers for both the third and second observation dates, one can finish by solving the optimal barrier for the first semi-annual observation date using the same Newton–Raphson process.

17.2 PRICING THE PUTTABLE REVERSE CONVERTIBLE

With the puttable reverse convertible it is the buyer who has the option to unwind the reverse convertible early, in which case the buyer sees his full investment plus a part of his coupons returned early.

Consider a 2 year puttable reverse convertible on BMW which can be unwound by the buyer at any of the three semi-annual observation dates. The put within the reverse convertible that is used to enhance the coupon is a 100 % put knocking in at 80 %. Chapter 15 showed that a regular reverse convertible on BMW pays a semi-annual coupon of 3.5 % with a re-offer of 100 %. Therefore, one expects the puttable

³ The Newton–Raphson process solves for the barrier by decreasing the barrier until it gives a worse coupon than the previous barrier and then subsequently halving the spread between these last two barriers. With this process one can solve for the optimal barrier to any required accuracy.

reverse convertible to pay a semi-annual coupon of less than 3.5 %. In order to solve for the semi-annual coupon of this puttable reverse convertible, one prices this structure as an autocallable where the barrier levels at the respective semi-annual observation dates are optimised to give the smallest coupon for a re-offer or fair value of 100 %. **However, in contrast to the callable reverse convertible, the barriers are not breached if the stock price is above the barrier but when the stock price is below the barrier. In other words, the barriers are down-and-out rather than up-and-out.** To begin with one can start with a barrier level of 100 % for all the three semi-annual observation dates, which appears to result in a coupon of 2.4 %. Secondly, one can try to price an autocallable where the first two observation dates have a barrier of 0 %⁴ and the third semi-annual observation date has a barrier of 90 %. This reverse convertible autocallable appears to give a coupon of 2.2 % for a fair value of the whole reverse convertible structure of 100 %. Proceeding in this way one finds that the most optimal (lowest coupon) barrier levels at which the structure autocalls are the following for the respective semi-annual observation dates, and with this barrier structure the puttable reverse convertible pays a semi-annual coupon of 2 %:

- The barrier for the first observation date, below which the structure knocks out, is 85 %. This means that, if BMW's share price is lower than 85 % of its initial value at the first observation date, it is optimal for the buyer of the puttable reverse convertible to unwind the structure and receive 102 % from the issuer. The reason that it is optimal for the buyer to unwind the full structure and receive 102 % is that the value of the down-and-in put is sufficiently large that it is more economical to get paid the full 100 % early plus the first coupon of 2 % and lose the short down-and-in put position than to keep the structure with the prospect of receiving three additional coupons but receiving the full 100 % 1.5 years later and having a potential liability on the short down-and-in put.
- The barrier for the second observation date, below which the structure knocks out, is 87 %. This means that, if BMW's share price is lower than 87 % of its initial value at the second observation date, it is optimal for the buyer of the puttable reverse convertible to unwind the structure and receive 104 % from the issuer.
- The barrier for the third observation date, below which the structure knocks out, is 90 %. This means that, if BMW's share price is lower

⁴ Such a low barrier allows for the optimisation of the third semi-annual observation date.

than 90 % of its initial value at the third observation date, it is optimal for the buyer of the puttable reverse convertible to unwind the structure and receive 106 % from the issuer.

In other words, the full puttable reverse convertible can be priced as:

- **Early redemption plus coupons**

- The trader sells a zero strike knock-out put with barriers of 85 %, 87 % and 90 % for the first three semi-annual dates respectively and a 0 % barrier for the fourth semi-annual observation date. The knock-out put can only knock out if BMW's share price is **below** the barrier level at one of the first three semi-annual observation dates and if BMW's share price is above the barrier at the fourth (maturity) semi-annual observation date and pays a rebate of

1. The rebate at the first semi-annual observation date is 100 % plus one times 2 %, equals a rebate of 102 %
2. The rebate at the second semi-annual observation date is 100 % plus two times 2 %, equals a rebate of 104 %
3. The rebate at the third semi-annual observation date is 100 % plus three times 2 %, equals a rebate of 106 %
4. The rebate at the fourth and last semi-annual observation date is 100 % plus four times 2 %, equals a rebate of 108 %

A barrier of 0 % for the fourth semi-annual observation date ensures the client sees his investment returned after the 2 year period.

- **Financing through down-and-in put**

- The trader buys a 100 % put knocking in at 80 % and knocking out (down-and-out) at respectively 85 %, 87 % and 90 % if BMW's share price is below the barrier level at one of the first three semi-annual observation dates. Obviously the trader would price the 80 % knock-in barrier with a barrier shift, i.e. pricing it as a 100 % put with a 77 % knock in. It is important to recognize that the downside barrier at 80 % can knock in continuously, i.e. throughout the full life of the puttable reverse convertible, whereas the other three solved for optimal barriers can only knock out at one of the three semi-annual observation dates.

Asian options are options whose strike price or settlement price¹ does not depend on one observation but on an average of observations. Therefore there are two types of Asian options. The *Asian out* option is a European option where the settlement price is determined by an average of the underlying on a set of predetermined observation dates. This means that, at expiry, the *Asian out* option pays the difference between the average of the underlying (floating settlement price) on these predetermined observation dates and the fixed strike. The *Asian in* option is a European option where the strike price is determined by an average of the underlying on a set of predetermined observation dates. This means that, at expiry, the *Asian in* option pays the difference between the underlying price at expiry and the average of the underlying at these predetermined observation dates (floating strike). The next sub-sections discuss the pricing and risk management of the *Asian out* and *Asian in* options respectively.

18.1 PRICING THE GEOMETRIC ASIAN OUT OPTION

Asian out options can be subdivided into a class where the average part of the Asian options is geometric and one where the average part is arithmetic. The price of an Asian out option where the average part is geometric is easy to price, as the geometric average of a lognormally distributed underlying has a lognormal distribution. Although Asian out options with an arithmetic average part are more common, this sub-section will first discuss geometric Asian out options. Understanding the geometric Asian out option makes it easier to understand the workings of an arithmetic Asian out option.

Consider a 1 year geometric average out Asian call option with monthly observations and a strike price K on a stock S_t . This means

¹ The settlement price is the share price against which an option settles, which, for a regular option, is the share price at the close of business on the expiry day.

that the payout at maturity of this Asian out call option, C_T^g , is:

$$C_T^g = \max \left[\left(\prod_{i=1}^{12} S_i^{\frac{1}{12}} \right) - K, 0 \right], \quad (18.1)$$

where S_i is the stock price of S_t at the end of the i -th month. Suppose σ_i are the implied volatilities of S_i with corresponding time units of $i/12$ -th of a year. And r_i are the interest rates for the i -th month. Since the correlation between σ_i and σ_{i+1} is zero, one derives that

$$\begin{aligned} \text{Var} \left[\ln \left(\prod_{i=1}^{12} S_i^{\frac{1}{12}} \right) \right] &= \text{Var} \left[\sum_{i=1}^{12} \frac{1}{12} \ln(S_i) \right] \\ &= \frac{1}{12} \sum_{i=1}^{12} \text{Var} [\ln(S_i)] \\ &= \frac{1}{12} \sum_{i=1}^{12} \sigma_i^2. \end{aligned} \quad (18.2)$$

and

$$\begin{aligned} E \left[\ln \left(\prod_{i=1}^{12} S_i^{\frac{1}{12}} \right) \right] &= E \left[\sum_{i=1}^{12} \frac{1}{12} \ln(S_i) \right] \\ &= \frac{1}{12} \sum_{i=1}^{12} E [\ln(S_i)] \\ &= \frac{1}{12} \sum_{i=1}^{12} r_i. \end{aligned} \quad (18.3)$$

This means that the geometric Asian out call option can be priced as a call option on the process

$$\prod_{i=1}^{12} S_i^{\frac{1}{12}}, \quad (18.4)$$

which has a volatility of σ^g

$$\sigma^g = \sqrt{\frac{1}{12} \sum_{i=1}^{12} \sigma_i^2}. \quad (18.5)$$

The interest rate r^g , which is effectively the financing of the hedge, for pricing the geometric Asian out option is

$$r^g = \frac{1}{12} \sum_{i=1}^{12} r_i. \quad (18.6)$$

This means that the geometric Asian out option can be priced with the Black–Scholes model using a volatility as in equation 18.5 and an interest rate as in equation 18.6.

It is important to recognise that σ_i is always smaller than σ_{i+1} and r_i is always smaller than r_{i+1} . The longer the period over which the volatility is measured the larger it gets, since any deviation in period i is automatically in period $i + 1$. However, the annualised volatility of σ_i can still be larger than the annualised volatility of σ_{i+1} . In other words, **the volatility of an average of observations is always smaller than the volatility of the share price itself. Obviously, the more observations the Asian option has, the smaller the implied volatility.**² This basically means that the price of a geometric average out option is always smaller than the price of a regular option with the same maturity.³

18.2 PRICING THE ARITHMETIC ASIAN OUT OPTION

Unlike the geometric Asian out option, the arithmetic mean of a log-normally distributed underlying does not have a lognormal distribution. However, one can use the geometric average as an estimate for the arithmetic mean and hence get a good understanding of the Greeks and other features of the arithmetic Asian out option. The actual pricing of the arithmetic Asian out option would always be done with a Monte Carlo process. For the sake of completeness, the payout at maturity for both the arithmetic Asian out call and put are given in formulae 18.7 and 18.8

² It is very important to distinguish between annualised implied volatility and the implied volatility associated with the actual term of the option in this regard. The volatility or implied volatility as referred to in this section is always the volatility associated with the actual term of the option and is therefore not expressed as an annualised volatility.

³ There is a very special situation where the price of a geometric Asian out put is larger than the regular put with the same maturity. This is when the interest rate is so large that the regular put is worth very little. Since equation 18.6 shows that the interest rate for the geometric Asian out option is substantially smaller than the interest rate of the regular put, this geometric Asian out put can be worth more, as the decrease in interest rate has a bigger impact on the price than the decrease in volatility. A geometric average out call option is always worth less than the regular call option with the same maturity.

respectively.

$$\max \left[\frac{1}{12} \left(\sum_{i=1}^{12} S_i \right) - K, 0 \right], \quad (18.7)$$

$$\max \left[K - \frac{1}{12} \left(\sum_{i=1}^{12} S_i \right), 0 \right]. \quad (18.8)$$

The reason one can use the geometric average as an estimate for the arithmetic average is the Jensen inequality. Jensen's inequality states

$$\sqrt[n]{\prod_{i=1}^n (S_i)^{a_i}} \leq \frac{1}{n} \sum_{i=1}^n a_i S_i \quad \sum_{i=1}^n a_i = 1. \quad (18.9)$$

This means that $\prod_{i=1}^n (S_i)^{a_i}$ is a lower bound for $\sum_{i=1}^n a_i S_i$. Although using the geometric average as an estimate for the arithmetic average will not be helpful in accurately pricing the arithmetic Asian out option, one can use this estimate to derive the dynamics of the Greeks and other distinguishing features of the arithmetic Asian out option.

The following sub-section will discuss the delta hedging of the arithmetic Asian out option from an economical perspective rather than a mathematical perspective. The sub-section thereafter will discuss the Greeks of the arithmetic Asian out option using the geometric average as an estimate for the arithmetic average.

Before moving on to the Greeks it is important to observe that, just like the geometric Asian option, the implied volatility of the arithmetic Asian out option is smaller than the implied volatility of the regular option with the same maturity. This can be seen by using the geometric average as an estimate for the arithmetic average and hence deriving formula 18.5. However, a more economic way to understand this is by recognising that the duration of any arithmetic Asian option is smaller than the duration of the regular option with the same maturity. The following sub-section gives a more elaborate discussion on the duration of an arithmetic Asian out option. Obviously, since the effective implied volatility of an arithmetic Asian out option is smaller than the implied volatility of the regular option with the same maturity, the Asian out option is always worth less⁴ than the equivalent regular option.

⁴ Apart from the Asian out put option in a very high interest rate environment.

18.3 DELTA HEDGING THE ARITHMETIC ASIAN OUT OPTION

The best way to show the delta hedging of an arithmetic Asian out option is by means of an example.

Consider a 3 month arithmetic Asian out call option on BMW with a strike of € 40 and monthly observations. Suppose a trader sells this Asian call on BMW to an investor. At inception the delta of the arithmetic Asian out call option is very similar to the equivalent regular three month call option. Assume that BMW's stock price at inception is € 40 and hence the delta for both the arithmetic Asian and the regular call option can be assumed to be $1/2$. This means that if the trader sells 90 thousand of these Asian out call options, he will delta hedge himself by buying 45 thousand BMW shares. Now suppose that, on the day before the observation date of the first Asian setting, BMW stock is trading at € 48. In this case the delta of the arithmetic Asian out call will be larger than the equivalent regular call. This is for the simple reason that the first Asian setting will almost certainly be in the money, i.e. higher than € 40. Suppose that the Asian out call has a delta of $5/6$, i.e. 75 thousand shares long, and the equivalent regular call option has a delta of $3/4$, i.e. 67.5 thousand shares long. However, since the first Asian setting is assured to be in the money, which accounts for $1/3$ of the weight in the arithmetic average, 30 thousand shares of the 75 thousand shares serve as a delta hedge for the first Asian setting. Once this Asian setting is taken the trader does not need to hold these 30 thousand shares long any more and should therefore sell them at the close of the first Asian setting. This means that his delta hedge against the arithmetic Asian out call option after the first Asian setting has reduced to 45 thousand shares. However, since the first Asian setting has been taken, which accounts for 30 thousand shares, the remaining arithmetic average is only on 60 thousand shares and therefore the remaining Asian out call is effectively on 60 thousand shares only. After the first Asian setting, BMW's share price goes down and, on the day before the second Asian setting, BMW stock is trading at € 32. Therefore the delta the trader holds against the Asian call option has decreased significantly to 10 thousand shares long. Since the second Asian setting is pretty much assured to be out of the money, the delta the trader holds is purely against the third Asian setting and nothing against the second one. For that reason the trader does not need to do anything on the close of the second Asian setting. On the day of the third Asian setting, BMW stock is trading at € 42 and therefore in the money. This

means that the trader holds a long position of 30 thousand shares as a delta hedge. Obviously, the trader has to sell these shares at the close on the day of the third Asian setting.

The above example clearly shows the procedure for hedging an arithmetic Asian out option. **On the day of an Asian setting, the trader needs to unwind part of his delta hedge if the Asian setting is in the money and does not need to do anything if the Asian setting is out of the money. If the Asian setting is in the money, the trader needs to unwind as a share position the number of Asian options multiplied by the weight of the setting.** This also shows that the duration of an Asian option is shorter than the equivalent regular option. Namely, on each of the Asian settings the weight multiplied by the full size on the Asian out position rolls off. Effectively, an Asian option is spread out over a set of regular options with maturities equal to the Asian setting dates. Therefore the term of an Asian option cannot be compared to a regular option with the same term but is in fact equivalent to a regular option with a shorter term.⁵ For any option it holds that the longer the time to maturity the higher the price. This once again makes clear that the price of an Asian out option should be smaller than its equivalent regular option.

18.4 VEGA, GAMMA AND THETA OF THE ARITHMETIC ASIAN OUT OPTION

Since sub-section 18.3 clearly shows that the duration of an Asian out option is shorter than the maturity of the equivalent regular option, it is easy to see that the Greeks of the Asian out option change accordingly. This basically means that, compared to a regular option, an Asian out option has a higher gamma and theta but a smaller vega. For example, the risks and Greeks of a 1 year 'at the money' Asian out option with monthly observations are very similar to the risks and Greeks of a 6 months 'at the money' option. Obviously, a 6 months 'at the money' option has in turn a higher gamma and theta than a 1 year 'at the money' option, but a smaller vega.

18.5 DELTA HEDGING THE ASIAN IN OPTION

It is not so much the pricing that makes the Asian in option interesting but the delta hedging of it. Apart from some financing on the delta hedge,

⁵ The term of the equivalent regular option can be estimated to be the weighted average of the different regular option with maturities equal to the Asian setting dates.

the pricing of an Asian in option is really nothing more than the pricing of a forward starting option as discussed in Chapter 11. Also the Greeks of the Asian in option are the same as the Greeks of a forward starting option.⁶ What is not obvious for an Asian in option is the delta hedging up until the date on which the strike is fixed, according to the average of the Asian in settings. That is exactly what will be discussed in this sub-section for the arithmetic Asian in option.

Consider a trader who buys a 3 day Asian in put on BMW with a maturity of 1 year from the date on which the strike price is fixed. The stock has the following price trend over these 3 days:

1. € 30
2. € 20
3. € 10

The trader buys this Asian in option on a notional of € 9 million. On the day of the first Asian in setting, the Asian in option is obviously an ‘at the money’ put as the average over the first 1 day is equal to the share price. Therefore, on the first day the trader assumes that the delta of the Asian in option will be 0.5 (a 1 year ATM put has a delta of 0.5) and hedges this by buying a **third** of € 4.5 million, i.e. 50 thousand shares of BMW. On the second day, the average price over the first two days has fallen to € 25, but the current share price is at € 20 and is therefore already an in the money put, namely the share price is 80 % of the strike. A one year put on BMW that is 20 % in the money appears to have a delta of 0.6. This means that with the current information the trader needs to be long 216 thousand shares on the third day, the day of the last Asian in setting. This 216 thousand is derived from the fact that, given an average share price of € 25, a notional of € 9 million and a delta of 0.6, the delta position should be $0.6 * 9 * 10^6 / 25 = 216$ thousand on the third day. Since the trader has already bought 50 thousand shares on the first day he still has to buy 166 thousand shares over two days. This means that his delta hedge for the second Asian setting is 88 thousand shares. On the third day, the average share price over the full three days has fallen to € 20. Since the share price is at € 10, the put is 50 % in the money. A 1 year 50 % in the money put appears to have a delta of 80 %. This means that the trader’s full delta hedge should be $0.8 * 9 * 10^6 / 20 = 360$ thousand shares. Since the trader has already bought 138 thousand shares in the

⁶ Buying an Asian in option also means buying gamma embedded in the Asian in part, which affects the pricing of an Asian in option. This effect will be elaborated on in sub-section 18.6.

Table 18.1 Hedging the Asian in put

BMW share price for each Asian setting	Delta of 1 year put with strike price equal to the average of the preceding days	Number of shares bought as a hedge
Day 1 setting of € 30	0.5	Buy delta times a third of $9 * 10^6$ divided by the average 1 day share price, $1/2 * 1/3 * 9 * 10^6 / 30 = 50$ thousand shares
Day 2 setting of € 20	0.6	Buy $1/2$ of delta times $9 * 10^6$ divided by the average 2 day share price minus the shares bought on the first day, $1/2 * (0.6 * 9 * 10^6 / 25 - 50 * 10^3) = 88$ thousand shares
Day 3 setting of € 10	0.8	Buy delta times $9 * 10^6$ divided by the average 3 day share price minus the shares bought on the first two days, $(0.8 * 9 * 10^6 / 20 - 116 * 10^3) = 222$ thousand shares

first two days, he has to buy another 222 thousand shares of BMW on the day of the last Asian in setting to be fully hedged. A summary of the hedging process for this Asian in put option can be found in Table 18.1.

18.6 ASIAN IN FORWARD

The Asian in forward is a very popular structure and is often used by corporates to buy back their own stock. An Asian in forward is composed of a long Asian in call and a short Asian in put, both of which have the same observation dates serving as Asian in settings and both expire on the last Asian in setting. In other words, the Asian in forward expires on the very same day that the strike is set according to the Asian in schedule. This means that, if BMW wants to repurchase €900 million worth of its own stock over the next 30 days, it could execute this by buying a 30 day Asian in call and selling a 30 day Asian in put, where each Asian in setting has the same weight, $1/30$.

Executing this stock repurchase through an Asian in forward, rather than buying the stock back in the market place, bears a significant advantage for BMW. This is because there is *hidden gamma* embedded in the Asian in forward. The investment bank that facilitates this Asian in forward and is therefore buying the Asian in put and selling the Asian in call, can easily monetise this gamma. The embedded gamma in this Asian in forward makes it an incredibly powerful structure and a very attractive solution to any corporate that wants to repurchase its own stock. The following example shows that an investment bank does actually materialise a gain on the gamma within the Asian in forward.

Suppose that BMW uses an Asian in forward to repurchase its own stock, where each Asian in setting has the same weight. For the sake of clarity it is assumed that BMW wants to buy back €900 million worth of stock in three days rather than 30 days. BMW stock has the following price trend over the 3 days of the Asian in term:

1. €30
2. €20
3. €10

Since BMW wants to buy €900 million worth of stock at the average 3 day stock price, BMW expects to get delivered 45 million shares of BMW on day 3 and pays €20 for each share. This means that on the third day BMW pays the investment bank €900 million and in return the investment bank delivers 45 million of BMW's own shares. Since the delta of the Asian in forward will be 1 on the day of the last Asian in setting, i.e. the day the strike price is set, the investment bank needs to make sure that it buys €900 million worth of stock over the three day Asian in term. This €900 million worth of BMW stock is determined on the basis of a BMW share price that is equal to the 3 day average BMW price. In other words the investment bank needs to have a hedging scheme that makes sure it has accumulated 45 million stock on the third day. This hedging scheme along with its economic implications can be summarised by Table 18.2. The table makes clear that the investment bank does actually make money on the gamma embedded in the Asian in forward, as the hedging scheme accumulates 45 million shares over the course of three days, for which the investment bank only pays €780 million, after which it sells these 45 million shares for €900 million to BMW itself. Thus, the fact that the share price moves according to the prescribed price trend ensures that the investment bank

Table 18.2 Hedging the Asian in forward

BMW share price for each Asian setting	Number of shares bought as a hedge	Total cost of the hedge on each Asian setting day
Day 1 setting of € 30	Buy a third of 900 divided by the average 1 day share price, $1/3 * 900/30 = 10$ million shares	€ 300 million
Day 2 setting of € 20	Buy $1/2$ of 900 divided by the average 2 day share price minus the shares bought on the first day, $1/2 * (900/25 - 10) = 13$ million shares	€ 260 million
Day 3 setting of € 10	Buy 900 divided by the average 3 day share price minus the shares bought on the first two days, $(900/20 - 23) = 22$ million shares	€ 220 million

makes a profit of € 120 million. Since the Asian in forward fixes the notional rather than the number of shares, the number of shares to be delivered on the day of the last Asian setting is dependent on the average share price over the Asian in period. Since the average goes up if the share price goes up, the investment bank buys fewer shares on a day that the share price goes up and more shares on a day that the share price goes down. This effectively means that the investment bank buys shares low and sells them high, which is exactly the definition of being long gamma. The investment bank will obviously compensate BMW for the fact that the Asian in forward gives the investment bank the opportunity to be long gamma and will therefore pay BMW a better price for the Asian in forward. Obviously, the higher the volatility of BMW the more the investment bank will pay BMW for this Asian in forward.

18.7 PRICING THE ASIAN IN FORWARD

Pricing the Asian in forward is relatively benign and is typically done by a Monte Carlo process where the main input is the volatility. Although a Monte Carlo process can easily provide a price for the Asian in forward,

this sub-section discusses the pricing from an economical perspective and will therefore provide a better insight into the dynamics of the Asian in forward. Sub-section 18.6 showed that the volatility of the underlying strongly affects the price of the Asian in forward. The higher the underlying volatility the more an investment bank is prepared to pay for being long the Asian in put and short the Asian in call. However, to recognise another very important aspect that affects the price of the Asian in forward, it is assumed that the volatility of the underlying is zero. In this case the underlying price stays the same throughout the term of the Asian in period and therefore the delta hedge that is executed on each day is known upfront and the same is true for all Asian in settings. In this case the price of the Asian in forward is very easy to determine. Namely, the investment bank will only charge the corporate for the financing of the delta hedge, because the investment bank will pay money on each of the Asian setting days to purchase the shares but will only receive the money from the corporate on the day of the last Asian in setting. This basically means that the investment bank will lose interest over the term of the Asian in settings and will charge that through to the corporate. To quantify this better, consider the example where BMW wants to repurchase € 900 million of its own stock through a 30 day Asian in forward. Since the volatility is assumed to be zero, the investment bank will hedge itself on each day of an Asian in setting by buying € 30 million worth of stock. This means that the investment bank will lose interest for 29 days on the first € 30 million plus for 28 days on the second € 30 million plus ... plus for 1 day on the penultimate € 30 million. Obviously one can calculate the exact cost of this financing, but an easier method is to say that the investment bank incurs financing for 29 days on half the full notional, i.e. € 450 million. Assuming that the euro interest rate is 0.03 per annum, one calculates that the investment bank will charge BMW $€ 450 * 0.03 * 29/365 = 1.0875$ million, which is equivalent to $1/2 * 0.03 * 29/365 = 0.1192\%$ of the full notional of € 900 million. Obviously, since BMW's share price does move, the investment bank, depending on the volatility, might end up paying BMW for the Asian in forward, as the structure makes the investment bank long gamma. This effect can also be estimated when one knows the gamma on the Asian in forward and by using the profit formula described in equation 2.3. Since the gamma is not constant this estimate will not be accurate and hence a Monte Carlo process is needed to calculate the exact price.

18.8 ASIAN IN FORWARD WITH OPTIONAL EARLY TERMINATION

The last interesting class of Asian option structures is the Asian in forward with a clause that the buyer of the Asian in put, and therefore the seller of the Asian in call, has the option to terminate early the Asian in forward. In other words, the investment bank has the option to deliver the shares early at the average share price up to the early termination day. The reason that this is such a popular structure is that it gives the investment bank additional optionality and therefore the corporate who is looking to repurchase stock takes advantage of yet another gamma feature that enhances the premium the corporate receives for buying back stock. There are two ways of pricing this Asian in forward with optional early termination. One can either price this as a Monte Carlo process or one can estimate the price by pricing the regular Asian in forward and adding the price of an Asian in put to this. Although the latter method will not give the exact price for the Asian in forward with early termination, it is very useful in getting an intuitive feel for the structure and what affects its price. Therefore this sub-section will first discuss how to accurately estimate the price of the Asian in forward with early termination, after which it will discuss the exact pricing of the Asian in forward with optional early termination through the Monte Carlo method.

Consider a 33 day Asian in forward on BMW on a notional of €900 million. The buyer of the Asian in put, and therefore the seller of the Asian in call, has the option to deliver €900 million worth of BMW shares early at any day after day 29, i.e. either day 30, 31 or day 32, at the average share price up to that point. The investment bank would obviously only choose to early terminate if the close on any of the 30th, 31st or 32nd days is lower than the average over respectively the first 30, 31 or 32 days. This means that, if the investment bank wants to early terminate on the 30th day, it will buy all the remaining shares to be able to deliver €900 million worth of stock at the average of the first 30 days. Since the investment bank has the option to early terminate and would only do that if the close on the 30th day is lower than the average of the preceding 30 days, the investment bank is basically long a 30 day Asian in put that expires on day 30. By exercising this Asian in put, the investment bank must buy back the remaining shares that would otherwise have been bought over the next 3 days. This means that the notional on this put should be $\frac{3}{33}$ times the notional on the full

Asian in forward. Since the investment bank can also decide to early terminate on the 31st or 32nd day, a lower estimate for the price of the 33 day Asian in forward with optional early termination from day 30 is to price it as a 30 day Asian in forward on 30/33 times € 900 million plus a 30 day Asian in put on 3/33 times € 900 million. Armed with this intuition one can more easily understand how to price the Asian in forward with optional early termination as a Monte Carlo process.

When pricing the Asian in forward with optional early termination as a Monte Carlo process, it is important to realise that this is a Monte Carlo process with certain conditions. As observed previously, the investment bank which is short the Asian in forward (long Asian in put and short Asian in call) would early terminate on any of the days in the early termination period if the prevailing stock price divided by the average stock price up to that point was less than a certain constant. This constant will always be less than 1, but to get the right price for the Asian in forward with optional early termination, one has to optimise this constant⁷ for each day in the early termination period, i.e. in the example above one would need to optimise these constants for the 30th, 31st and 32nd days. The process of optimising these constants works along the same lines as optimising the barriers for the callable reverse convertible, see sub-section 17.1. One starts backwards and optimises the constant for the last day while keeping the constants for the previous days at zero to make sure that the structure has not been early terminated. Subsequently, one works backwards to optimise the constant for the preceding days. Once these constants on each day in the early termination period have been established by using the Newton–Raphson method, one prices the Asian in forward with optional early termination as the following Monte Carlo process. Suppose that the optimised constants are C_0 , C_1 and C_2 for respectively the 30th, 31st and 32nd day.

- For all paths where $\frac{S_{30}}{\sum_{i=1}^{30} \frac{S_i}{30}} \leq C_0$ the Monte Carlo process prices a 30 day Asian in forward on 30/33 of € 900 million and a 30 day Asian in put on 3/33 of € 900 million.
- For all paths where $\frac{S_{30}}{\sum_{i=1}^{30} \frac{S_i}{30}} > C_0$ **and** $\frac{S_{31}}{\sum_{i=1}^{31} \frac{S_i}{31}} \leq C_1$ the Monte Carlo process prices a 31 day Asian in forward on 31/33 of € 900 million and a 31 day Asian in put on 2/33 of € 900 million.

⁷ Note that this is a conservative estimate as it does not need to be a constant but can be any number associated to any Monte Carlo path which optimises the price of the Asian in forward.

-
- For all paths where $\frac{S_{30}}{\sum_{i=1}^{30} \frac{S_i}{30}} > C_0$ **and** $\frac{S_{31}}{\sum_{i=1}^{31} \frac{S_i}{31}} > C_1$ **and** $\frac{S_{32}}{\sum_{i=1}^{32} \frac{S_i}{32}} \leq C_2$ the Monte Carlo process prices a 32 day Asian in forward on 32/33 of €900 million and a 32 day Asian in put on 1/33 of €900 million.
 - For all paths where $\frac{S_{30}}{\sum_{i=1}^{30} \frac{S_i}{30}} > C_0$ **and** $\frac{S_{31}}{\sum_{i=1}^{31} \frac{S_i}{31}} > C_1$ **and** $\frac{S_{32}}{\sum_{i=1}^{32} \frac{S_i}{32}} > C_2$ the Monte Carlo process prices a 33 day Asian in forward on €900 million.

Quanto Options¹

The quanto option is designed for investors who want to execute an option strategy on a foreign stock but are only interested in the percentage return of that strategy and want to get paid this return in their own currency. The basic principle of a quanto option is that the exchange rate will be fixed to the prevailing exchange rate at inception of the option transaction and the payout of the quanto option will be this exchange rate times the payout of the regular option.

Consider a US investor who is very bullish on the share price of BP (British Petroleum). For that reason this US investor wants to buy an ATM call option on BP expiring in one year. However, he does not want to get his return in British pounds but in US dollars. Assume that BP's share price is £ 5, the exchange rate is currently 2 dollars per pound and in one year's time the share price is worth £ 5.50. The quanto ATM call option will give the US investor a payout at maturity of \$ 1 regardless of the change in exchange rate. In other words, the investor would expect a 10 % return in US dollars on the USD notional² amount he bought calls on, since BP's share price has increased by 10 %.

From the above example it is clear that it is relatively easy to structure a quanto option, however it is much harder to see what effect it has on how this option should be priced and what variables it depends on.

19.1 PRICING AND CORRELATION RISK OF THE QUANTO OPTION

To see how the quanto option should be priced, two new variables have to be introduced. The first variable is the correlation between BP's stock price and the foreign exchange, FX, rate. To get a better understanding of how this correlation affects the price of a quanto call option, an interesting question is whether the US investor, who is buying an ATM USD quanto call, is short or long correlation. Here short correlation

¹ Parts of this chapter have been previously published in de Weert, F. (2006) *An Introduction to Options Trading*, John Wiley & Sons Ltd, Chichester. Reproduced with permission.

² Notional is defined as the number of options times the strike price.

means that the investor benefits (quanto option increases in value) if the correlation goes down and loses if the correlation goes up. The reverse holds for being long correlation. To answer this question, assume the correlation is positive, which means that, if the British pound gets more valuable against the US dollar, BP's share price goes up. Now, one can easily see that the investor is short correlation (selling correlation). Because, if the British pound goes up, the correlation aspect causes the share price to go up and therefore the dollar increase of a plain vanilla call option, which is effectively a quanto option with zero correlation, is more than the dollar increase of the USD quanto call option, which has a fixed exchange rate. A similar analysis shows that, if the British pound goes down, the correlation aspect causes the share price to go down and the combination of a lower share price and a less valuable pound causes the dollar loss of a plain vanilla ATM call option to be less than the dollar loss of an ATM USD quanto call option. To put it differently, if the correlation goes down from for example 1 to -1 , the holder of a USD quanto call option benefits, since he is better off holding a plain vanilla call³ option if the correlation is 1 whereas if the correlation is -1 he is better off holding a USD quanto call option. The above example shows that the holder of a USD quanto call is short correlation. However, when extending this analysis to a USD quanto put one finds that the holder of such an option is long correlation.

The second variable is the FX volatility. For this variable it is less obvious whether the US investor is long or short this FX volatility. At this point it is good to introduce a model that describes the stock price difference in US dollars for a small time interval. This model gives an answer under which circumstances a quanto option is long or short FX volatility and shows once more why the holder of a quanto call is short correlation and the holder of a quanto put is long correlation.

The model most commonly used for modeling a share price on a non-dividend paying stock is the Black–Scholes model which describes the difference in stock price for a small time interval. The formula is as follows.

$$\frac{dS_t}{S_t} = r dt + \sigma_S dW_t, \quad (19.1)$$

³ A plain vanilla option is equivalent to a quanto option where the correlation between the exchange rate and the stock price is zero. The only difference is that the quanto option is priced in the quanto currency and the plain vanilla option in the share's own currency. This means that to have an exact comparison one should have the FX times as many quanto options as regular options.

where

- S_t is the stock price at time t .
- r is the risk free interest rate.
- dS_t is the change in stock price over time interval dt .
- dt is a small time interval.
- σ_S is the volatility of the stock price.
- W_t is Brownian motion, which is a stochastic process characterised by normally distributed intervals dt with a mean of 0 and a variance equal to the length of the interval dt . Mathematically stated the intervals dt have a distribution equal to $N(0, dt)$.
- dW_t is a stochastic process with distribution $N(0, dt)$.

For the purpose of understanding a quanto option, a similar model to equation 19.1 is required which includes two additional variables, namely the correlation, ρ , between the logarithm of the stock price and the logarithm of the exchange rate and the FX volatility, σ_{FX} . However, it is important to realise that this equation defines a new share, F_t , quoted in the currency the option is quanto'd into rather than the share's own currency. F_t is defined in such a way that a regular dollar option on F_t is in fact a quanto option on S_t .⁴

$$\frac{dF_t}{F_t} = (r_{local} - \rho\sigma_S\sigma_{FX})dt + \sigma_S dW_t, \quad (19.2)$$

where

- ρ is the correlation between the logarithm of the stock price and the logarithm of the exchange rate, where the exchange rate is quoted as number of dollars per pound. A positive correlation means that if the pound increases in value the stock price goes up.
- r_{local} is the risk free interest rate of the stock's own currency. In BP's case the r_{local} is the risk free interest rate of the pound.

From this equation it is obvious that anyone who holds a quanto call option benefits if the correlation decreases and the holder of a quanto

⁴ The quanto model described in equation 19.2 is intuitive as the volatility, σ_S , used to price the quanto option should be same as the implied volatility of the underlying stock, since the exchange rate is fixed and therefore the option payout depends solely on the actual movement of the stock. The interest rate or drift part is slightly different as the delta hedge is affected, see sub-section 19.2, by exchange rate movements and therefore affects the financing on the delta hedge. If for example the stock price doubles, this affects positively the financing on a long call option (short shares at a higher stock level therefore receive more interest). However, if the correlation between the stock and the exchange rate is such that a doubling in stock price results in a halving of the dollar value against the pound, the hedge is unaffected and therefore has no impact on the financing of the delta hedge. Therefore it is obvious that the financing part of the quanto stock model is adjusted by $-\rho\sigma_S\sigma_{FX}$.

put benefits if the correlation increases. Depending on the sign of the correlation this formula also makes clear whether the holder of a quanto option is long or short FX volatility. If the correlation is negative the holder of a quanto call is long FX volatility and the holder of a quanto put is short FX volatility. If the correlation is positive the holder of a quanto call is short FX volatility and the holder of a quanto put is long FX volatility. Knowing that the model for the change in stock price for a USD stock with a dividend yield equal to d is

$$\frac{dF_t}{F_t} = (r_{\$} - d) dt + \sigma_S dW_t, \quad (19.3)$$

shows that the price of a quanto option can be derived from a normal option by making an adjustment to the dividend yield of adding $r_{\$} - r_{local} + \rho \sigma_S \sigma_{FX}$ to the dividend yield. In equation 19.3 $r_{\$}$ is the USD risk free interest rate. This also shows that the main theme of a quanto option is the difference in forward compared to a regular option.

19.2 HEDGING FX EXPOSURE ON THE QUANTO OPTION

Hedging the FX exposure on the quanto option might not be very intuitive but is very simple in practice. This is because the FX hedge is captured by the delta hedge, which in turn is because the option notional in the local currency keeps changing with the changing FX. Namely, the notional of the quanto option is agreed in the quanto currency and therefore the notional in the local currency changes whenever the FX changes. This obviously implies that if the quanto currency halves in value with respect to the local currency, the notional of the quanto option in the local currency halves and therefore the trader needs to halve his delta hedge even though the stock price might not have moved. The delta of the quanto option can easily be determined with the help of equation 19.2, the value of the stock in the local currency and the strike in the local currency.

To show how the delta hedging of a quanto option works, which at the same time captures the FX hedging of the quanto option, consider the example in the introduction to this chapter. Suppose a trader sells this ATM quanto call option on BP with a strike of \$ 5 and on a notional of 10 million USD. Assuming that the parameters of equation 19.2 and a stock value of 5 and strike price of 5 give this ATM option a delta of 0.5. As it was assumed that the exchange rate was 2 dollars per pound, the

trader would need to buy 2.5 million pounds worth of BP stock as his delta hedge. Now suppose that USD halves with respect to GBP, which means that the exchange rate goes to 4 dollars per pound. Just because of this change in exchange rate and without the stock price moving, the trader would need to change his delta hedge to being long 1.25 million pounds worth of BP stock. The trader therefore needs to sell 1.25 million pounds worth of BP stock as a result of the exchange rate moving.

The above shows that there is no need for a trader to put an FX hedge in place for a quanto option. However, the trader would need to swap out the FX on the premium received for the quanto option if he sold the quanto option. This is because equation 19.2 prescribes financing in the local currency. If the trader buys a quanto option he effectively ensures financing in the local currency because he would first need to sell the local currency to buy the quanto currency in order to pay for the quanto option. Therefore the trader will only need to do an FX hedge if he sells a quanto option, which is selling the quanto currency and buying the local currency on the premium amount. This ensures financing in the local currency.

Composite Options¹

The composite option is designed for investors who want to execute an option strategy on a foreign stock but want to fix the strike in their own currency and get the payout of this option in their own currency. In contrast to the quanto option, where the holder gets a percentage return regardless of the exchange rate, the holder of a composite option has exposure to the exchange rate. One of the reasons that a composite option is traded is to protect the value in their own currency on a foreign investment. Consider the following example.

20.1 AN EXAMPLE OF THE COMPOSITE OPTION

A US investor owns stock in the British pharmaceutical Glaxo Smithkline (GSK). Assume that the current value of GSK is £ 13.00 and that the exchange rate is 2 dollars per pound. This means that the dollar value of one share is \$ 26. To protect this holding he buys a one year ATM composite put option on GSK. This means that the strike price of this composite put option is equal to \$ 26. Assume that after one year the stock goes down to £ 11.00 and the exchange rate goes from 2 dollars per pound to 1.5 dollars per pound (dollar increases in value). This means that the dollar value of one GSK share has gone down from \$ 26 to \$ 16.5. However, because the strike price of the composite option is fixed in dollars, the dollar loss on the shares is offset by the payout of the composite put option, which is equal to the strike price (\$ 26) minus the new dollar value of one GSK share (\$ 16.5). In summary, the holder of a composite option wants to protect the share value in his own currency to both exchange rate movements and movements in the stock price.

The best way to answer the question whether the holder of a composite option is short or long correlation is to model the change in stock price, for a small time interval, in the currency of the composite. In the example above this would be the dollar. Using the same notation as in

¹ Parts of this chapter have been previously published in de Weert, F. (2006) *An Introduction to Options Trading*, John Wiley & Sons Ltd, Chichester. Reproduced with permission.

Chapter 19, sub-section 19.1 the model looks as follows (see Appendix A):

$$\frac{dF_t}{F_t} = r_{compo}dt + \sigma_{compo}dW_t, \quad (20.1)$$

where

$$\sigma_{compo}^2 = \sigma_S^2 + 2\rho\sigma_S\sigma_{FX} + \sigma_{FX}^2, \quad (20.2)$$

and r_{compo} is the risk free interest rate of the composite currency.

This formula shows that the holder of a composite option is long correlation and the main theme is the difference in implied volatility compared to a regular option. Another way to visualise the USD composite put on GSK is by comparing it to a regular \$ 26 put option on the GSK ADR.²

20.2 HEDGING FX EXPOSURE ON THE COMPOSITE OPTION

If a trader sells the composite put option on GSK in the example of sub-section 20.1, he obviously wants to delta hedge himself. Because equation 20.1 describes a composite option on stock as a normal option on a stock in the composite currency, a delta can easily be determined by setting strikes and share price in the composite currency. However, the trader can only execute his delta on the local stock, which is not quoted in the composite currency. This means that, even if the trader delta hedges himself, he will still have an FX risk, as the payout on the composite option will be in the composite currency and his delta hedge is in the local currency of the underlying stock. To illustrate this FX risk consider the following.

Suppose the trader hedges the short composite put option on GSK by just delta hedging. This means by the example of sub-section 20.1 that, if the trader hedges the composite option on a $\delta = 1$,³ he makes £ 2 (13 – 11) per composite option, which is worth \$ 3 at maturity as it was assumed that the dollar increases from \$ 2 per pound to \$ 1.5 per pound.

² ADR stands for American Depository Receipt and is nothing more than a USD quoted stock on a company that has its main listing in a different country.

³ Using $\delta = 1$ is just for argument's sake. Normally the delta would be much smaller at inception of an option trade and would converge towards one as the option becomes more in the money and gets closer to maturity. Whether the profit on the hedge is equal to the loss on the option ultimately depends on the volatility of both the stock and the FX, which is not taken into account in this example.

As the dollar value of GSK goes from \$26 to \$16.5 during the term of the option, the trader loses \$9.5 on the composite option and only makes \$3 on his delta hedge. This obviously shows that just delta hedging is not enough. The trader would need to buy dollars on the notional of his delta hedge to be fully immune to all the risks of the composite option on GSK. This means that, as an FX hedge, the trader would sell £13 to receive \$26 for every GSK share he shorted as part of his delta at inception of the trade. At maturity the trader can buy back these £13 for \$19.5. The total profit on this FX hedge is therefore \$6.5, making the profit on both the FX hedge as well as the delta hedge equal to \$9.5, which is equal to the loss on the composite option.

To summarise the hedging of the FX exposure on a composite option one just needs to realise that if one sells stocks as a delta hedge, one needs to also sell the currency the stock is quoted in and buy the composite currency in the same notional as the delta hedge. If the delta hedge is to buy stock, the FX hedge is to also buy the currency the stock is quoted in and sell the composite currency in the same notional as the delta hedge. Obviously the FX hedge is not static and should be adjusted along with any delta adjustments. In other words, to be perfectly hedged against both stock movements and FX movements one needs to have, at any time, the same notional of FX hedge as delta hedge.

The above shows how to hedge the FX exposure of a composite option. However, it did not take into account that one also needs to put an FX hedge in place on the premium paid for the composite option at inception of the trade. For the call option the FX hedge on the paid premium works in the opposite direction to the FX hedge on the notional of the delta hedge. For the put option the FX hedge on the paid premium of the composite option works in the same direction as the FX hedge on the notional of the delta hedge. Obviously, the reason that a trader wants to hedge the FX exposure on the premium is that in order to buy the composite option in the composite currency the trader would first need to sell the local currency to buy the composite currency. This effectively gives him an FX position which still needs to be hedged. Counterintuitively, no FX hedge needs to be executed on the premium if the trader sells a composite option, as the trader receives the premium in the composite currency and the model used to determine the price and the delta of the option is equation 20.1, which assumes financing in the composite currency, i.e. USD in the above example. The only reason that

Table 20.1 FX hedge on the composite option

Position composite option	FX hedge on delta notional	FX hedge on premium paid for composite option at inception of trade
Long composite call option	Sell local currency to buy composite currency	Buy local currency to sell composite currency
Short composite call option	Buy local currency to sell composite currency	None
Long composite put option	Buy local currency to sell composite currency	Buy local currency to sell composite currency
Short composite put option	Sell local currency to buy composite currency	None

the trader needs to execute an FX hedge on the premium paid if he buys a composite option is to offset the FX he executes to actually be able to buy the option and therefore ensure financing in the composite currency as the model in equation 20.1 prescribes. Therefore the full FX hedge can be summarised by Table 20.1.

Outperformance Options

The outperformance option measures the outperformance, whether it be on the upside or downside, of one stock against another and pays this difference. Obviously, an outperformance option can also be structured between two different assets, like a stock and the oil price, in which case it measures the outperformance of the stock against the oil price. The beauty of outperformance options is that their pricing and risks look complicated but can easily be understood when comparing them to a composite option. The next sub-section gives an example of an outperformance option.

21.1 EXAMPLE OF AN OUTPERFORMANCE OPTION

Consider an investor who believes that BMW's share price will outperform the share price of Volkswagen (VW) in the next year. Therefore he buys a 1 year outperformance option on BMW versus VW. Suppose that BMW's share price is € 40 and VW's share price is € 70. The investor decides to buy this outperformance of BMW versus VW on a notional of € 7 million. To demonstrate the payoff of this option two different price trends are considered. First, assume that in 1 year BMW's share price will be trading at € 50 and VW's share price at € 63. In this case the payout on the outperformance option should be 35 % of € 7 million, as BMW went up 25 % and VW went down 10 % and therefore BMW outperformed VW by 35 %. Secondly, assume that in 1 year BMW's share price is at € 50 and the share price of VW is at € 91. In this case the payout on the outperformance option is zero as VW outperformed BMW by 10 %.

Another way to think about the payout of the outperformance option is to say that the holder of the outperformance option gets paid the greater of zero and the notional growth of BMW shares minus the notional growth/decline of VW shares. Viewing it like this makes it intuitive that the outperformance option is nothing more than a composite option, because the payout of a composite call option in the currency of the underlying stock is equal to the notional growth of the share minus the

notional growth/decline of the composite currency, quoted in the local currency. So, if one views a currency as an asset that is quoted in the currency of the underlying stock of a composite option, one sees that the definition of the composite option and the outperformance option are exactly equivalent. In other words, a composite option is a type of outperformance option where the performance of the stock is measured against a specific currency.

Another way to look at the outperformance option on two stocks is to say that one is allowed to swap the stock the investor thinks will underperform for the other stock in the ratio of the two stock prices at inception of the trade.

21.2 OUTPERFORMANCE OPTION DESCRIBED AS A COMPOSITE OPTION

Sub-section 21.1 made clear that an outperformance option is nothing more than a composite option. Although this is the case, it might not be easy to see how this translates into a comparison between the actual outperformance option where the strike is defined and a composite option. This is because with a composite option one is only interested in the payout of the option in the composite currency, and to compare the outperformance option with the composite option one needs to convert this payout to the currency of the underlying stock.

Before actually going into the pricing of an outperformance option it is good to first clearly state the payoff at time t for an outperformance option on BMW, BMW_t , versus VW, VW_t . This payoff is the following

$$C_t^{out} = \max \left[\left(\frac{BMW_t}{BMW_0} - \frac{VW_t}{VW_0} \right), 0 \right] \cdot N, \quad (21.1)$$

where N is the notional of the outperformance option, which was taken to be 7 million. By multiplying equation 21.1 by

$$\frac{BMW_0}{VW_t} \cdot \frac{VW_t}{BMW_0}$$

equation 21.1 can be rewritten as

$$C_t^{out} = \max \left[\left(\frac{BMW_t}{VW_t} - \frac{BMW_0}{VW_0} \right), 0 \right] \cdot VW_t \cdot \frac{N}{BMW_0}. \quad (21.2)$$

To see how a price for C_t^{out} can be determined, one just has to be able to price

$$C_t^{out,VW} = \max \left[\left(\frac{BMW_t}{VW_t} - \frac{BMW_0}{VW_0} \right), 0 \right]. \quad (21.3)$$

However, closer investigation of equation 21.3 gives that $C_t^{out,VW}$ expresses nothing more than a call option on *BMW* where the payout is in a different currency. The currency is obviously one VW share. In this currency the value of one BMW share at time t is $\frac{BMW_t}{VW_t}$, and as the initial value of a BMW share in the VW currency is $\frac{BMW_0}{VW_0}$ one sees that equation 21.3 is the payoff of an ATM call option on BMW in the VW currency. In turn, equation 21.3 can be viewed as a call option on a newly defined stock $F_t = \frac{BMW_t}{VW_t}$, which according to Appendix A can be described by the following process

$$\frac{dF_t}{F_t} = \sigma_{out} dW_t, \quad (21.4)$$

where

$$\sigma_{out}^2 = \sigma_{BMW}^2 + \sigma_{VW}^2 - 2\rho\sigma_{BMW}\sigma_{VW}. \quad (21.5)$$

The Black–Scholes formula gives a price for a call option on a process F_t , which is defined by equations 21.4 and 21.5, and has a strike price F_0 .

What makes it difficult to compare the outperformance option with the composite option is that with the composite option one is only interested in the value of the composite option in the composite currency, whereas with the outperformance option one has to convert the price of the option in the currency of VW back into a sensible euro price. In other words, in the case of the outperformance option on BMW versus VW, one needs to multiply $C_t^{out,VW}$ by VW_t to get the actual value of the outperformance option on BMW versus VW. Multiplying this by the number of options, N/BMW_0 , one exactly gets the payoff of the outperformance option as described by equation 21.2.

21.3 CORRELATION POSITION OF THE OUTPERFORMANCE OPTION

Equation 21.5 makes clear that the owner of an outperformance option has a short position in the correlation between the two assets. In other words, the more inversely correlated the two stocks in the

outperformance option are, the more valuable is this outperformance option. This is intuitive as the holder of an outperformance option stands to receive the highest payout if the two stocks move in opposite directions, i.e. the outperforming stock moves up and the other stock moves down.

21.4 HEDGING OF OUTPERFORMANCE OPTIONS

Delta hedging the outperformance option is very similar to the hedging of the composite option. Firstly, one can easily find a delta on the composite BMW call in the VW currency, as described by equation 21.4. So, if a trader is long a BMW outperformance option versus VW, he is basically long a composite call option on BMW in the VW currency on a notional that is equal to $VW_t \cdot \frac{N}{BMW_0}$. The parameters of equation 21.4 provide a delta, δ_{out} , from which he can calculate the euro notional of BMW he would need to sell

$$\frac{\delta_{out} \cdot VW_t \cdot \frac{N}{BMW_0}}{BMW_t}. \quad (21.6)$$

According to Table 20.1 the trader would then need to buy the VW currency in the same notional as his delta hedge on BMW. Buying the VW currency obviously means that the trader needs to buy VW shares and is therefore effectively paying (selling) euros in the same notional as his delta hedge on BMW. However, the hedge on the premium paid for the outperformance option works differently to the hedge for the composite option. Equation 21.4 shows that the financing of the outperformance option should be flat, which means that the dollar amount

Table 21.1 Hedges on the outperformance option and net euro cash flow

Security	Amount bought or sold	Euro cash flow
Outperformance option	1	$-C_T^{out,VW} \cdot VW_t \cdot \frac{N}{BMW_0}$
BMW share	$-\frac{\delta_{out} \cdot VW_t \cdot \frac{N}{BMW_0}}{BMW_t}$	$+\delta_{out} \cdot VW_t \cdot \frac{N}{BMW_0}$
VW share	$+\frac{\delta_{out} \cdot VW_t \cdot \frac{N}{BMW_0}}{VW_t}$	$-\delta_{out} \cdot VW_t \cdot \frac{N}{BMW_0}$
VW share	$-\frac{C_T^{out,VW} \cdot VW_t \cdot \frac{N}{BMW_0}}{VW_t}$	$C_T^{out,VW} \cdot VW_t \cdot \frac{N}{BMW_0}$
Net cash flow		0

spent and received should be the same. This means that to hedge this outperformance option fully the trader would also need to buy euros to sell the VW currency on the premium paid for the outperformance option. In other words, sell VW shares and therefore buy euros on the premium paid for the outperformance option. Table 21.1 summarises the full hedge of the outperformance option on BMW over VW and also shows that the actual cash euro value adds up to zero. Obviously, when hedging an outperformance option, a trader would always net out the two hedges on VW in Table 21.1. In Table 21.1 the two hedges on VW are shown separately for the sake of clarity, as each line serves a different aspect of the hedge of the outperformance option.

Best of and Worst of Options

The *best of* option is an option where the investor receives the gain on the best performing stock among a predefined number of stocks. In other words, the investor is long a call on the best performing stock among the predefined number of stocks. The *worst of* option is an option where the investor is long a put on the worst performing stock amongst a predefined number of stocks. The pricing of either the best of or the worst of option is typically done by using a Monte Carlo process. There is even a closed form solution for the best of and worst of options. However, since both the best of and worst of options are, in practice, always priced as a Monte Carlo process, this section focuses on explaining the risks associated with such options. The main risk of either a best of or a worst of option is the correlation between the underlying assets on which the best of or worst of option is based. Luckily, the correlation risk can be easily derived from an outperformance option.

22.1 CORRELATION RISK FOR THE BEST OF OPTION

Consider an investor who buys a best of option between VW and BMW expiring in 1 year. This means that after 1 year the investor gets paid the greater of the percentage gain in VW and BMW. In other words the investor is long a call option on either VW or BMW, whichever has performed better after one year. Mathematically, the payoff at maturity looks like this:

$$C_T^{\text{Best of VW/BMW}} = \max \left[\frac{VW_T - VW_0}{VW_0}, \frac{BMW_T - BMW_0}{BMW_0}, 0 \right]. \quad (22.1)$$

The best way to understand the correlation risk embedded in this best of option is to view this option as an ‘at the money’ call on VW plus an outperformance option of BMW over VW. Unfortunately this representation is only equal to the best of option under the condition that if VW is down after 1 year it has at least outperformed BMW. In other words, if VW is lower after one year and BMW has outperformed VW,

the representation of an ‘at the money’ call on VW plus an outperformance option of BMW over VW gives a payout, whereas the best of option does not give a payout at all. Nonetheless, this representation proves very useful in determining whether the seller of a best of option is long or short correlation. Mathematically, the representation of an ‘at the money’ call on VW plus an outperformance option on BMW over VW looks like this

$$\begin{aligned}
 C_T^{Best\ of\ VW/BMW} &= \max \left[\frac{VW_T - VW_0}{VW_0}, \frac{BMW_T - BMW_0}{BMW_0}, 0 \right] \\
 &\leq \max \left[\frac{VW_T - VW_0}{VW_0}, 0 \right] \\
 &\quad + \max \left[\left(\frac{BMW_T}{VW_T} - \frac{BMW_0}{VW_0} \right), 0 \right]. \quad (22.2)
 \end{aligned}$$

Whereas the payoff of an ‘at the money’ call on VW plus an outperformance option of BMW over VW is greater than the best of option of BMW and VW, the payoff of a long ‘at the money’ forward on VW plus an outperformance option on BMW over VW is smaller than the payoff of the best of option between BMW and VW. Mathematically, the payoff of the best of option is therefore encapsulated by these two representations and looks like this

$$\begin{aligned}
 &\max \left[\frac{VW_T - VW_0}{VW_0}, 0 \right] \\
 &- \max \left[\frac{VW_0 - VW_T}{VW_0}, 0 \right] + \max \left[\left(\frac{BMW_T}{VW_T} - \frac{BMW_0}{VW_0} \right), 0 \right] \\
 &\leq C_T^{Best\ of\ VW/BMW} \\
 &\leq \max \left[\frac{VW_T - VW_0}{VW_0}, 0 \right] \\
 &\quad + \max \left[\left(\frac{BMW_T}{VW_T} - \frac{BMW_0}{VW_0} \right), 0 \right]. \quad (22.3)
 \end{aligned}$$

Since any outperformance option is short correlation, each side of inequality 22.3 is short correlation. The first term in formula 22.3 at a correlation of -1 is greater than the third term at a correlation of 1 . Also at a correlation of -1 , $C_T^{Best\ of\ VW/BMW}$ is equal to the first term in formula 22.3. At a correlation between BMW and VW of 1 , $C_T^{Best\ of\ VW/BMW}$ is

equal to the third term of formula 22.3. This means that the value of $C_T^{\text{Best of VW/BMW}}$ increases as the correlation goes down. Hence, the holder of a best of option is short correlation. In other words, the lower the correlation between the constituents of a best of option, the higher the value of this best of option.

22.2 CORRELATION RISK FOR THE WORST OF OPTION

Consider an investor who buys a worst of option between VW and BMW expiring in 1 year. This means that after 1 year the investor gets paid the greater of the percentage decline in VW and BMW. In other words the investor is long a put option on either VW or BMW, whichever has performed worse after one year. Mathematically, the payoff at maturity looks like this:

$$C_T^{\text{Worst of VW/BMW}} = \max \left[\frac{VW_0 - VW_T}{VW_0}, \frac{BMW_0 - BMW_T}{BMW_0}, 0 \right]. \quad (22.4)$$

To determine the correlation risk associated with a worst of option, the same analysis as in sub-section 22.1 is needed. The payoff of the worst of option has a lower bound equal to a short forward position in VW plus an outperformance option of VW over BMW and an upper bound equal to a long 'at the money' put position plus an outperformance option of VW over BMW. Mathematically this can be presented as

$$\begin{aligned} & \max \left[\frac{VW_0 - VW_T}{VW_0}, 0 \right] \\ & - \max \left[\frac{VW_T - VW_0}{VW_0}, 0 \right] + \max \left[\left(\frac{VW_T}{BMW_T} - \frac{VW_0}{BMW_0} \right), 0 \right] \\ & \leq C_T^{\text{worst of VW/BMW}} \\ & \leq \max \left[\frac{VW_0 - VW_T}{VW_0}, 0 \right] \\ & + \max \left[\left(\frac{VW_T}{BMW_T} - \frac{VW_0}{BMW_0} \right), 0 \right]. \end{aligned} \quad (22.5)$$

Again, the first term in formula 22.5 at a correlation of -1 is greater than the third term at a correlation of 1 . Also at a correlation of -1 ,

$C_T^{\text{Worst of VW/BMW}}$ is equal to the first term in formula 22.5. At a correlation between BMW and VW of 1, $C_T^{\text{Worst of VW/BMW}}$ is equal to the third term of formula 22.5. This means that the value of $C_T^{\text{Worst of VW/BMW}}$ increases as the correlation goes down. Hence, the holder of a worst of option is also short correlation. In other words, the lower the correlation between the constituents of a worst of option, the higher the value of this worst of option.

22.3 HYBRIDS

Although hybrids have not explicitly been discussed so far, the main products used to create a hybrid have already been discussed. It is important to understand that hybrids in themselves are not exotic but it is the product that can comprise underlyings from different asset classes that makes a structure exotic. For example, the best of and worst of options do not necessarily have to compare the performances of underlyings from the same asset class, but can also compare performances from different asset classes. The same holds true for the outperformance option and even the basket option, which will be discussed in Chapter 24.

Variance Swaps

In recent years variance swaps have become increasingly popular as they are a better way to express a view on the volatility of an underlying stock or index than regular options. When betting on the volatility of a certain underlying, options have the disadvantage of having vega and gamma exposures concentrated around the strike price of the option. This means that, if the stock price moves away from the strike price during the term of the option, any movement in the underlying is irrelevant with respect to the profit of that option and, because the vega has disappeared as well, changes in implied volatility do not influence the profitability of that option either. Therefore, executing a pure bet on the volatility of an underlying through options is far from perfect because one can be right on the volatility of the underlying during the whole term of the option and still lose money. For example, a trader wants to take advantage of cheap volatility by buying an option but the stock drifts away from the strike in a very unvolatile manner after which it becomes very volatile, bringing the realised volatility of the option during the term of the option far above the implied volatility for which the trader initially bought his option. In this case the trader obviously loses money despite being right on the volatility. To get around this local characteristic of volatility exposure for options, variance swaps were introduced.

A variance swap is a contract that pays the difference between the **annualised** variance of an underlying and the annualised variance strike agreed upon at inception of the trade. Its payoff at maturity is equal to

$$(\sigma_R^2 - \sigma_K^2) \times N, \quad (23.1)$$

where

- σ_R is the annualised realised volatility during the term of the variance swap;
- σ_R^2 is the annualised realised variance during the term of the variance swap;
- σ_K is the annualised volatility strike;

- σ_K^2 is the annualised variance strike;
- N is the notional of the variance swap, also called variance notional, in euros per annualised volatility point squared.

As with any swap the cash flow at inception is zero. In other words, when pricing a variance swap one solves for variance strike, σ_K^2 , that makes the expected payoff of the swap at maturity zero.

The following sub-sections will show how to replicate a variance swap with options and how it can be shown that, at any moment, the vega is constant in stock price and decreases linearly over time and the gamma of a variance swap is constant over time and although the gamma is not constant in stock price the gamma cash is. This is in stark contrast to the concentrated gamma around the strike of a single option.

23.1 VARIANCE SWAP PAYOFF EXAMPLE

Suppose a trader has sold a 1 year variance swap on BMW at an implied volatility level of 20 % and on a notional of € 5 million. After 1 year the realised volatility appears to have been 18 %. In this case the trader will receive an amount equal to

$$-(0.18^2 - 0.20^2) * 5\,000\,000 = 38\,000. \quad (23.2)$$

23.2 REPLICATING THE VARIANCE SWAP WITH OPTIONS

The variance swap can be understood more easily by replicating it with options. In this sub-section it will be shown that a portfolio of options consisting of $1/K^2$ options of each strike replicates the payoff profile of a variance swap to a factor $2/T$. In other words, multiplying this portfolio by $2/T$ exactly gives the payoff of the variance swap.¹

Consider a portfolio of options $\Pi(S_t)$ where each strike has a weighting of $1/K^2$. Since out of the money options are most liquid, all the options in the portfolio with a strike lower than the cutoff point S^c are put options and the options in $\Pi(S_t)$ with strikes larger than S^c are call options. S^c will be around the ‘at the money’ level. At expiration, T ,

¹ The following proof is based on the publication *More Than You Ever Wanted To Know About Volatility Swaps* (Demterfi *et al.*, 1999).

Appendix B proves that the value of this portfolio $\Pi(S_T)$ is:

$$\begin{aligned}\Pi(S_T) &= \int_0^{S^c} \frac{1}{K^2} \max[K - S_T, 0] dk + \int_{S^c}^{\infty} \frac{1}{K^2} \max[S_T - K, 0] dk \\ &= \frac{S_T - S^c}{S^c} - \ln\left(\frac{S_T}{S^c}\right).\end{aligned}\quad (23.3)$$

Appendix B also shows that the value of the portfolio Π for any time t is equal to²

$$\begin{aligned}\Pi(S_t) &= \int_0^{S^c} \frac{1}{K^2} [K e^{-r(T-t)} N(-d_2) - S_t N(-d_1)] dk \\ &\quad + \int_{S^c}^{\infty} \frac{1}{K^2} [S_t N(d_1) - K e^{-r(T-t)} N(d_2)] dk \\ &= \frac{S_t - S^c}{S^c} - \ln\left(\frac{S_t}{S^c}\right) + \frac{\sigma^2(T-t)}{2},\end{aligned}\quad (23.4)$$

where σ is the implied volatility expressed annually and is assumed to be the same for all different strikes. Because of the existence of skew this is not the case in practice, which will be elaborated on in sub-section 23.8. The variance exposure of $\Pi(S_t)$ can be obtained by differentiating equation 23.4 with respect to σ^2 . This gives the variance exposure φ as

$$\varphi = \frac{\partial \Pi(S_t)}{\partial \sigma^2} = \frac{T-t}{2}.\quad (23.5)$$

Equation 23.5 already shows that the variance exposure of the portfolio $\Pi(S_t)$ is constant and independent of the stock price. However, to make sure that the variance exposure at inception, i.e. $t = 0$, is equal to €1 per volatility point squared, one needs to have a quantity of $2/T$ of the portfolio $\Pi(S_t)$. Equation 23.4 makes it easy to prove that

$$\frac{2}{T} \Pi(S_t) = \frac{2}{T} \left[\frac{S_t - S^c}{S^c} - \ln\left(\frac{S_t}{S^c}\right) \right] + \frac{\sigma^2(T-t)}{T}\quad (23.6)$$

does indeed replicate the payoff of a variance swap. Indeed, if a trader buys $2/T$ of $\Pi(S_t)$ at $t = 0$ the fair value is

$$\frac{2}{T} \Pi(S_0) = \frac{2}{T} \left[\frac{S_0 - S^c}{S^c} - \ln\left(\frac{S_0}{S^c}\right) \right] + \sigma^2.$$

² This can be shown by integrating over K the value of the $1/K^2$ put option from 0 to S^c plus integrating over K the value of the $1/K^2$ call option from S^c to infinity.

If the trader delta hedges the above portfolio perfectly and, over the term of the options in the portfolio, the annualised realised volatility appears to be σ_R , the real value of $2\Pi(S_0)/T$ is

$$\frac{2}{T}\Pi(S_0) = \frac{2}{T} \left[\frac{S_0 - S^c}{S^c} - \ln \left(\frac{S_0}{S^c} \right) \right] + \sigma_R^2. \quad (23.7)$$

This means that the trader's payoff at maturity is equal to

$$\sigma_R^2 - \sigma^2. \quad (23.8)$$

Expression 23.8 is exactly the payoff formula for a variance swap when σ_K is equal to the implied volatility, σ , of the individual options in the portfolio, which was assumed to be constant for all different strikes. Since the implied volatility is the trader's best estimate of the future realised volatility, the implied volatility should be taken as the fair variance strike. Another way of putting it is that, if σ_K is equal to the implied volatility, σ , the variance swap can be replicated by $2/(T \cdot K^2)$ number of options per strike.

23.3 GREEKS OF THE VARIANCE SWAP

The Greeks of the variance swap can easily be derived from equation 23.6. Since equation 23.7 assumes that all the options in portfolio $\Pi(S_t)$ are being delta hedged, the delta of the variance swap cannot be derived from equation 23.6. However, assuming that the implied volatility is constant in strike, the delta of a variance swap is obviously zero as the variance payoff is independent of stock price. The fact that a variance swap does have a so-called skew delta will be discussed in sub-section 23.8.

To get to the expressions for vega, gamma and theta of a variance swap one just has to differentiate equation 23.6 to its respective variables.

The sensitivity of a variance swap with respect to the implied volatility, i.e. the vega, ν , of a variance swap is

$$\nu = \frac{\partial \left[\frac{2}{T}\Pi(S_t) \right]}{\partial \sigma} = 2 \cdot \frac{\sigma(T-t)}{T}. \quad (23.9)$$

Equation 23.9 shows that the vega of a long variance swap position is always positive, but also that the vega changes with the changing implied volatility. The speed at which the vega changes with the changing implied volatility can be derived by taking the derivative of ν with respect to

σ . The fact that the vega is not constant in implied volatility is called convexity and is elaborated on in sub-section 23.9.

The exposure, ϕ , of a variance swap to variance, σ^2 , is equal to

$$\phi = \frac{\partial \left[\frac{2}{T} \Pi(S_t) \right]}{\partial \sigma^2} = \frac{T - t}{T}. \quad (23.10)$$

Equation 23.10 shows that the variance exposure of a variance swap decreases linearly over time and is zero at maturity and 1 at $t = 0$.

The sensitivity of a variance swap with respect to time, i.e. the time decay or theta, θ , of a variance swap is

$$\theta = \frac{\partial \left[\frac{2}{T} \Pi(S_t) \right]}{\partial t} = -\frac{\sigma^2}{T} \quad (23.11)$$

Equation 23.11 shows that the time decay of the variance swap stays constant over time. Also, the theta of a variance swap is linearly dependent on the variance. This means that the cumulative theta at expiration is equal to the full variance. This is what one would expect since it shows that, if the stock does not move during the term of the variance swap, the seller of this variance swap will collect the full variance strike level times the variance notional, N .

The rate of change of the delta of the variance swap with respect to the stock price, i.e. the gamma of the variance swap, is equal to

$$\gamma = \frac{\partial^2 \left[\frac{2}{T} \Pi(S_t) \right]}{\partial^2 S_t} = \frac{2}{T \cdot S_t^2} \quad (23.12)$$

Equation 23.12 shows that the gamma of a variance swap is constant through time. It also shows that gamma is not constant in stock price. However, gamma cash, which is defined as $\gamma \cdot S_t^2 / 100$ (see equation 4.2), is constant in both stock price and time. **The fact that gamma cash is constant through time and in stock price means that the profit on a given percentage move for a long variance position is always the same regardless of the stock price or when the move occurs. This is an important difference from a single option, where both the stock price and the time that a move occurs matter significantly for the profit on a move of a long option position. Indeed, if the stock price is close to the strike price and the option is close to expiration, the profit on a move is far greater than when the option has a long time to maturity or when the stock price is not close to the strike price.**

One would expect that the profit formula, as derived in Chapter 3, also holds for variance swaps. This can be seen by substituting equation 23.11 for theta and equation 23.12 for gamma into the following formula.

$$\frac{1}{2}\gamma\sigma^2S_t^2 + \theta = 0. \quad (23.13)$$

23.4 MYSTERY OF GAMMA WITHOUT DELTA

It is paradoxical that a variance swap does not have a delta where it does have gamma on it. The reason is that a variance swap does accumulate delta cash in between the time intervals over which the volatility is measured, but resets to zero after each time interval. The delta cash within each time interval is given by gamma cash multiplied by the percentage move. To understand this one has to go back to the root of how realised volatility is measured and over which time intervals. These time intervals are agreed upon at inception of the variance swap transaction and are per default defined as the close to close volatility, i.e. the one day volatility. The following example clarifies how the time intervals over which the volatility is measured impact the relationship between delta cash and gamma cash.

Consider an investor who buys a 2 day variance swap on BMW where $\sigma_K = 32$, i.e. an annualised volatility of 32 %. BMW happens to move 1 % on both of these days. Assuming that a year has 256 days, the annualised realised volatility is therefore 16 and because the variance swap is bought at 32, the investor's break-even volatility is 2 % per day. Suppose that the gamma cash on the variance swap is 1 million euros. A gamma cash of 1 million and a break-even volatility of $z = 2\%$ a day give a theta of 20 thousand euros a day. Indeed, since the absolute break-even volatility in equation 3.2 is equal to $y = zS_t$, equations 3.2 and 4.2 give that

$$\theta = \frac{\gamma \cdot z^2 S_t^2}{2} = \frac{\gamma_c \cdot z^2}{200} = 20\,000. \quad (23.14)$$

Although the delta of the variance swap resets to zero at the end of each business day, one finds that the variance swap does have a delta during this business day. This means that at the start of each business day the profit of a variance swap is independent of the direction of the move. However, once the stock moves during the business day, the profit of the variance swap starts to depend on the direction of the move that follows. For example, if the buyer of the BMW variance swap snapshots

his position halfway through the first business day when BMW is down 0.5 %, he finds himself having a short delta cash position of 500 thousand euros. In other words, when BMW is already down 0.5 % the holder of the variance swap benefits more if the stock goes down even more than when it goes back up again. This is logical since the holder of the variance swap wants to have as large a move as possible. The profit on the variance swap when it goes down another 0.5 % to finish the first business day down 1 % can now be calculated as either:

1.

$$\frac{1^2}{200} \cdot \gamma_c = 5000$$

or

2. gamma profit on the first 0.5 % move + profit on 500 thousand delta cash for the second 0.5 % move (delta profit) + gamma profit on the second 0.5 % move (equals the first 0.5 % move), mathematically

$$\frac{0.5^2}{200} \cdot \gamma_c + 0.5 \times 5 \cdot 10^5 + \frac{0.5^2}{200} \cdot \gamma_c = 5000$$

Obviously, the profit on the move for the second business day is also 5000 euros as the gamma cash is constant through time and stock price. In other words, the investor has made 10 000 euros on the two respective moves, but lost 40 000 euros in theta, i.e. an overall loss of 30 000 euros.

23.5 REALISED VARIANCE VOLATILITY VERSUS STANDARD DEVIATION

The realised volatility, σ_R , as defined in the payoff of the variance swap, is not defined as a regular standard deviation. The volatility in the payoff of the variance swap is really defined as the square root of the lognormal returns squared times the annualisation factor. Mathematically,

$$\sigma_R^2 = \sum_{i=1}^{N_d} \left[\ln \left(\frac{S_i}{S_{i-1}} \right) \right]^2 \cdot \frac{256}{N_d}, \quad (23.15)$$

where,

- S_i is the stock price at the close of business on day i ;
- $256/N_d$ is the annualisation factor;
- N_d is defined as the number of days during the term of the variance swap.

The reason that the volatility in the variance payoff is defined according to equation 23.15 is to make sure that every move contributes to the variance payoff regardless of any trend the stock price might have. Indeed, the standard deviation reduces when a stock has a specific trend. For example, a stock that goes up by 1 % for 10 days in a row has a standard deviation of 0. In other words, the standard deviation of the lognormal returns compensates for any trend the stock might have by subtracting the average return from each return. Mathematically the standard deviation, s_d , of the lognormal returns is defined as

$$s_d = \sum_{i=1}^{N_d} \left[\ln \left(\frac{S_i}{S_{i-1}} \right) - \frac{\sum_{i=1}^{N_d} \ln \left(\frac{S_i}{S_{i-1}} \right)}{N_d} \right]^2 \cdot \frac{256}{N_d}. \quad (23.16)$$

23.6 EVENT RISK OF A VARIANCE SWAP VERSUS A SINGLE OPTION

It has already briefly been discussed that the event risk on a variance swap is far greater than the event risk on a single option, except for single options close to expiry and close to the strike price. The reason being that a variance swap has a constant gamma cash whereas the gamma cash of a single option is concentrated around the strike price. This means that if a large move occurs the profit impact is far greater on a variance swap than for a single option. The single option first needs to be relatively close to the strike for the move to have any impact at all. Even then the single option loses its gamma quickly for very large moves as it moves away from the strike price. Another very distinct feature of a variance swap is that the profit impact for a specific move does not depend on when the move occurs during the term of the option. A 10 % move that occurs a year away from maturity has the same profit impact as a 10 % move that occurs 1 day before maturity. This can easily be seen from the profit formula in combination with the fact that the gamma cash on a variance swap is constant. The profit formula is stated below. In this formula z is the percentage movement of the underlying stock price.

$$\frac{\gamma_c \cdot z^2}{200}$$

The fact that the profit on a move is independent of time for a variance swap is in stark contrast to a single option, where the timing of the move heavily affects the profitability of the move. This is for the simple reason

that the gamma on a single option changes quickly over time. Assuming that the option is at the money, a 10 % move that occurs a year away from maturity is far less profitable than a move that occurs a day before expiry. Indeed, the gamma of a 1 day option is a multiple of a 1 year option.

23.7 RELATION BETWEEN VEGA EXPOSURE AND VARIANCE NOTIONAL

The standard for trading variance swaps is to agree on the vega exposure of the variance swap rather than the variance notional. In order to calculate the variance notional associated with the vega exposure of a variance swap, one has to derive the variance notional needed to have the change in volatility times vega equal the change in variance times the variance notional, N . In other words, if the variance notional is N , the variance strike is σ_K^2 and the change in volatility is y , one wants to solve the following equation

$$[(\sigma_K + y)^2 - \sigma_K^2] \times N = \text{vega} \times y \times 100. \quad (23.17)$$

The factor 100 in the right hand side of equation 23.17 stems from the fact that volatility is expressed as a percentage. Since y is a small number, Taylor³ expansion on $(\sigma_K + y)^2$ reduces equation 23.17 to (square terms in y or higher are negligibly small)

$$\begin{aligned} \sigma_K^2 \cdot N + 2\sigma \cdot y \cdot N - \sigma_K^2 \cdot N &= \text{vega} \cdot y \cdot 100 \\ \Downarrow \\ N &= \frac{\text{vega}}{2\sigma} \cdot 100 \end{aligned} \quad (23.18)$$

23.8 SKEW DELTA

All the derivations and conclusions about variance swaps in the previous sub-sections are based on the assumption that the implied volatility is constant in strike price. Chapter 5 shows that this is not a valid assumption as the effect of skew results in higher implied volatilities for lower strike prices. This skew effect means that σ_K is higher than the ‘at the

³ If $f(y)$ is a function, Taylor expansion can rewrite the function $f(y)$ as

$$f(y) = f(0) + f'(0)y + \frac{f''(0)y^2}{2!} + \frac{f'''(0)y^3}{3!} + \dots$$

money' volatility. More importantly, it means that the variance swap has a delta after all, because the buyer of a variance swap benefits from the market going down as the market will likely become more volatile and the 'at the money' implied volatility will go up. The delta on a variance swap induced by the effect of skew is better known as the *skew delta*. The *skew delta* of a variance swap can easily be quantified and is clarified by the following example.

Consider a trader who sells a 1 year variance swap on BMW at $\sigma_K = 20$ and the vega he sells is 25 thousand euros. The skew is such that every 1 % down move in BMW's share price translates to a 0.20 % increase in the 'at the money' implied volatility. Since the trader is short 25 thousand euros of vega, each 1 % down move in the BMW price translates to a 5 thousand euros loss on the variance swap (σ_K goes up to 20.20 and hence a loss of $0.2 \times 25 = 5$ thousand euros). To hedge this exposure to the dynamics of the variance swap as a result of skew, the trader has to sell 500 thousand euros delta cash of BMW, which serves as his *skew delta*. This *skew delta* will exactly offset any losses on the BMW variance swap as a result of BMW going down by 1 %, as he will make 5 thousand euros on the *skew delta* for a 1 % down move ($500 \times 0.01 = 5$).

23.9 VEGA CONVEXITY

The vega of a variance swap is not constant in implied volatility. In other words, whenever the implied volatility changes the vega exposure of the variance swap changes. This can be seen by taking the derivative of v in equation 23.9 with respect to σ . This results in:

$$\frac{\partial v}{\partial \sigma} = \frac{\partial^2 \left[\frac{2}{T} \Pi(S_t) \right]}{\partial^2 \sigma} = 2 \cdot \frac{(T - t)}{T}. \quad (23.19)$$

Equation 23.19 shows that, if the implied volatility goes up, the vega exposure of a variance swap goes up as well.⁴ By the same token, if the implied volatility goes down, the vega of a variance swap goes down as well. Because of skew the 'at the money' implied volatility goes up when the market goes down, which in turn translates into a higher vega exposure. This means that the vega exposure goes up when the

⁴ Vomma is positive.

underlying goes down and vice versa.⁵ This principle is referred to as vega convexity. In other words, if the market goes down, the vega of a variance swap goes up and if the market goes up, the vega of a variance swap goes down. This can also be seen from formula 23.18. On any given day, when the implied volatility goes up, the vega needs to go up for the variance notional to remain constant.

⁵ Vanna is negative.

Dispersion literally means scattering or spreading widely. In trading, dispersion means that one trades the basket volatility against the volatilities of the constituents in that basket. The reason that being short volatility on the basket and long the volatilities of the basket constituents is called dispersion is because this trade makes money when the constituents are heavily inversely correlated, i.e. -1 . This can be seen by considering a basket with two constituents A and B with equal weights. If both A and B are very volatile, in which case being long volatility on these names makes money, and A and B are perfectly inversely correlated, i.e. A moves opposite to B , which means that, *de facto*, the basket does not move and therefore being short volatility on the basket makes money as well. This section discusses trading dispersion in more detail and shows how the volatility of the basket is derived from its constituents. More importantly, this chapter shows that being long dispersion means being short correlation plus long volatility rather than just being short correlation.

24.1 PRICING BASKET OPTIONS

Basket options are typically priced as a Monte Carlo process. However, to better understand the theory of dispersion it is imperative to be able to use the Black–Scholes formula to price a basket option. Unfortunately a basket option is not lognormally distributed even though its constituents are. Therefore it is impossible to use the Black–Scholes formula to accurately price a basket option. Nonetheless, one can use $\ln\left(\prod_{i=1}^n S_{i,t}^{w_i}\right)$, which is lognormally distributed, to estimate $\ln\left(\sum_{i=1}^n w_i S_{i,t}\right)$. This approximation enables one to use the Black–Scholes formula to price a basket option. Although this method will never provide the most accurate price for the basket option, it does give the opportunity to derive the specific features associated with a basket option and hence the features that are so pivotal in understanding dispersion trading.

24.2 BASKET VOLATILITY DERIVED FROM ITS CONSTITUENTS

Consider a basket, B_t , with n constituents, $S_{i,t}$, and each constituent has a weight w_i . Mathematically this basket looks like

$$B_t = \sum_{i=1}^n w_i S_{i,t}, \quad \sum_{i=1}^n w_i = 1. \quad (24.1)$$

If the ‘at the money’ implied volatility of each constituent $S_{i,t}$ is σ_i ,¹ statistics show that the ‘at the money’ implied volatility, σ_b , of the basket is determined by²

$$\sigma_b^2 = \text{Var} \left[\ln \left(\sum_{i=1}^n w_i S_{i,t} \right) \right] \quad (24.2)$$

$$\approx \text{Var} \left[\ln \left(\prod_{i=1}^n S_{i,t}^{w_i} \right) \right] \quad (24.3)$$

$$\begin{aligned} &= \sum_{i=1}^n \sum_{j=1}^n w_i w_j \text{Cov}(\ln(S_{i,t}), \ln(S_{j,t})) \\ &= \sum_{i=1}^n \sum_{j=1}^n w_i w_j \sigma_{ij} \\ &= \sum_{i=1}^n \sum_{j=1}^n w_i w_j \sigma_i \sigma_j \rho_{i,j} \\ &= \sum_{i=1}^n w_i^2 \sigma_i^2 + \sum_{i=1}^n \sum_{j=1, j \neq i}^n w_i w_j \sigma_i \sigma_j \rho_{i,j}, \end{aligned} \quad (24.4)$$

where

- σ_{ij} is the covariance between $\ln(S_i)$ and $\ln(S_j)$;
- σ_i is the implied volatility of S_i ;
- ρ_{ij} is the correlation between $\ln(S_i)$ and $\ln(S_j)$.

¹ This notation is purely for the purposes of simplicity. Unfortunately, it can be a bit deceiving as it seems as if the implied volatility is independent of time, which it obviously is not. It would be better to assume that the implied volatility equals $\sigma_i(T-t)$, but for simplicity the $(T-t)$ part is left out.

² Equation 24.2 makes implicit use of Jensen’s inequality, which states that

$$\sqrt[n]{\prod_{i=1}^n (S_{i,t})^{w_i}} \leq \frac{1}{n} \sum_{i=1}^n w_i S_{i,t} \quad \sum_{i=1}^n w_i = 1.$$

and therefore uses $\text{Var}[\ln(\prod_{i=1}^n S_{i,t}^{w_i})]$ to estimate $\text{Var}[\ln(\sum_{i=1}^n w_i S_{i,t})]$.

24.3 TRADING DISPERSION

In practice it is very easy to trade dispersion. Consider a basket with 10 equally weighted stocks. Suppose an investor wants to be long the volatilities of the constituents in respect to their weighting and short volatility on the basket. The way the investor achieves this is by buying vega on the stocks according to their weights in the basket and selling vega on the basket itself. The investor can either buy vega through buying options or through buying variance swaps (see Chapter 23) on the stocks. In order to be short vega on the basket he has to sell options on the basket or sell variance swaps on the basket. This investor wants to be short one million euros of vega on the basket and therefore long € 100 thousand vega on each of the individual stocks. Suppose that the weighted average implied volatility of the stocks is 25 and the implied volatility of the basket is 17. In this case the investor makes money if the spread between the weighted average implied volatility of the stocks and the implied volatility of the basket widens.³

24.4 QUOTING DISPERSION IN TERMS OF CORRELATION

Dispersion is typically quoted in terms of correlation. Equation 24.2 shows that one can express the volatility of the index or basket in terms of its constituent volatilities and the correlations between the stocks within the basket. In order to measure dispersion in terms of one correlation number, one assumes that $\rho_{i,j}$ is constant in all i and j and is equal to ρ_b . From equation 24.2 one can derive this correlation ρ_b as

$$\rho_b = \frac{\sigma_b^2 - \sum_{i=1}^n w_i^2 \sigma_i^2}{\sum_{i=1}^n \sum_{j=1, j \neq i}^n w_i w_j \sigma_i \sigma_j}. \quad (24.5)$$

24.5 DISPERSION MEANS TRADING A COMBINATION OF VOLATILITY AND CORRELATION

Consider an investor who is long dispersion. In other words, he is long vega on the stocks in the basket according to their weights and short

³ Sub-section 24.5 explains that a widening of the spread between the weighted average volatility and the implied volatility of the index is not always a recipe for making money on a long dispersion position. Namely, when the overall volatility goes down, one can even lose money on a widening of the spread. To make sure that a long dispersion position benefits from a widening spread, dispersion is often traded as a ratio where the vega position on the index is larger than the cumulative vega position on the stocks. This is discussed in sub-section 24.6.

vega on the basket itself. Mathematically this means that this investor has the following position:

$$\sum_{i=1}^n w_i \sigma_i - \sigma_b. \quad (24.6)$$

Equation 24.2 shows that this position is equivalent to

$$\sum_{i=1}^n w_i \sigma_i - \sigma_b = \sum_{i=1}^n w_i \sigma_i - \sqrt{\sum_{i=1}^n \sum_{j=1}^n w_i w_j \sigma_i \sigma_j \rho_{i,j}} \geq 0 \quad (24.7)$$

Equation 24.7 shows that being long dispersion means being long the volatilities of the individual stocks. Indeed, this is clear since $-1 \leq \rho_{i,j} \leq 1$ and therefore

$$\begin{aligned} \sum_{i=1}^n w_i \sigma_i &= \sqrt{\sum_{i=1}^n \sum_{j=1}^n w_i w_j \sigma_i \sigma_j} \\ &\geq \sqrt{\sum_{i=1}^n \sum_{j=1}^n w_i w_j \sigma_i \sigma_j \rho_{i,j}}. \end{aligned}$$

In other words, everything being equal (correlations and weights remain unchanged), if the volatilities of the individual stocks go up, a long dispersion position makes money. Knowing that a long dispersion position is like being long the individual stock volatilities, equation 24.7 can now be rewritten in terms of ρ_b , the constant correlation in which dispersion is typically measured, to see what the sensitivity is to correlation for a long dispersion position.

$$\sum_{i=1}^n w_i \sigma_i - \sigma_b = \sum_{i=1}^n w_i \sigma_i - \sqrt{\sum_{i=1}^n w_i^2 \sigma_i^2 + \sum_{i=1}^n \sum_{j=1, j \neq i}^n w_i w_j \sigma_i \sigma_j \rho_b} \quad (24.8)$$

Equation 24.8 shows that a long dispersion position is short correlation. In other words, if the correlation goes down, a long dispersion position makes money.

In summary, a long dispersion position is first and foremost a short correlation position. This means that, the more inversely correlated the individual constituents are, the more more profitable is a long dispersion position. However, it is important to recognise that a long dispersion

position also has a long exposure to the individual stock volatilities. To make sure that a dispersion position is solely a short correlation position, traders typically do not trade the cumulative vega of the stocks in equal size to the vega of the index or basket. The following sub-section shows how to ratio the vega of each leg of the dispersion to only have an exposure to correlation.

24.6 RATIO'D VEGA DISPERSION

As stated in the previous sub-section, dispersion is not a real short correlation trade. Indeed, a decrease in correlation can be offset by a decrease in volatility of the individual stocks. For this reason traders typically ratio the size in which they trade the vega on the index compared to the vega on the individual stocks. Determining the required ratio to establish a sole correlation position and therefore take away the residual long volatility position is rather simple and intuitive. This ratio is indeed nothing more than the ratio of the weighted average stock volatility level and the index volatility level. Consider the example in sub-section 24.3 where the weighted average stock volatility is 25 and the index volatility is 17. To effectively establish a sole short correlation position one has to trade the dispersion in a ratio of $1:\frac{25}{17}$, meaning that one has to be short 25/17 times the amount of vega on the index compared to the cumulative single stock position. In other words, if the cumulative single stock position is € 1 million vega, a short position of € 1.47 million vega in the index effectively gives a pure short correlation position.

When trading a dispersion through variance swaps one has to be careful to put the right ratio in place. Since a dispersion through variance swaps is equivalent to the position

$$\sum_{i=1}^n w_i \sigma_i^2 - \sigma_b^2, \quad (24.9)$$

the ratio should first of all be expressed in variance notionals rather than vega. The ratio in which the variance notionals should be traded is then derived by the square of the ratio between the weighted average single stock volatility and the index volatility. This means that if the weighted average single stock volatility is 25, the index volatility is 17 and the cumulative single stock variance notional is € 200 million, the variance notional on the index should be $200 \cdot (25/17)^2 = € 432.53$ million.

24.7 SKEW DELTA POSITION EMBEDDED IN DISPERSION

Any dispersion position has an embedded skew delta position. This stems from the fact that the skew on an index is almost always steeper than the skew of any of the individual stocks. In turn, the reason for this is that an individual stock might be volatile on the upside as well as on the downside, whereas the market as a whole becomes more volatile when the market goes down and less volatile when the market goes up. The fact that an index has a steeper skew than each of the individual stocks results in the fact that a long ratio'd dispersion position makes money as the market goes up and loses money if the market goes down, all other variables being equal. Indeed, whenever the market goes up, the short index implied volatility within the long ratio'd dispersion position goes down more than the implied volatility of the individual stocks. This is due to a steeper index skew than the skew of the individual stocks. When the market goes down, the steeper index skew results in the reverse, i.e. the short index position loses more than the long single stock volatilities make. In summary, *a long dispersion is effectively long delta as a result of a steeper skew on the index or basket than on the individual stock constituents*. This means that a long dispersion position has an embedded long skew delta position, which can be hedged by selling some delta on, for example, the index.

Engineering Financial Structures

This chapter discusses different retail investors and how to financially engineer structures that suit each investor's needs. The structures discussed in this chapter are typically not complicated or even exotic, but, since financially engineered retail products are such an important part of investor demand, this chapter shows how notes can be structured to achieve the financial needs of different investor profiles. There are basically three different types of retail investors: the risk averse retail investor who wants to invest his money in capital protected products; the retail investor who wants some exposure to the equity markets, but still wants to have a balanced and well thought out risk profile; the retail investor who uses financial structures to leverage his views.

25.1 CAPITAL GUARANTEED PRODUCTS

A capital guaranteed product is a suitable product for the risk averse investor who does want to have some exposure to the equity markets. The principle of the capital guaranteed product is that the investor buys a zero coupon bond and, since the present value of a zero coupon bond is less than 100 %, the difference between the present value of the zero coupon bond and 100 % is invested in buying an option that gives exposure to the equity markets. In other words, with a capital guaranteed product, the investor gives up a coupon to subsequently invest it in either a call or a put option. For example, a risk averse investor who is bullish on BMW could buy a capital guaranteed product where, instead of receiving a coupon on the bond, he invests the coupon to buy participation in a call option. Obviously, if the investor were bearish on BMW, he would give up the coupon to buy participation in a put option.

Consider a risk averse investor who is bearish on BMW for the coming year and therefore wants to buy a capital guaranteed structure that increases in value if BMW goes down. Suppose that a one year zero coupon bond is worth 95.5 %, and a one year 'at the money' put on BMW is worth 5.625 %. Therefore the investor buys a zero coupon

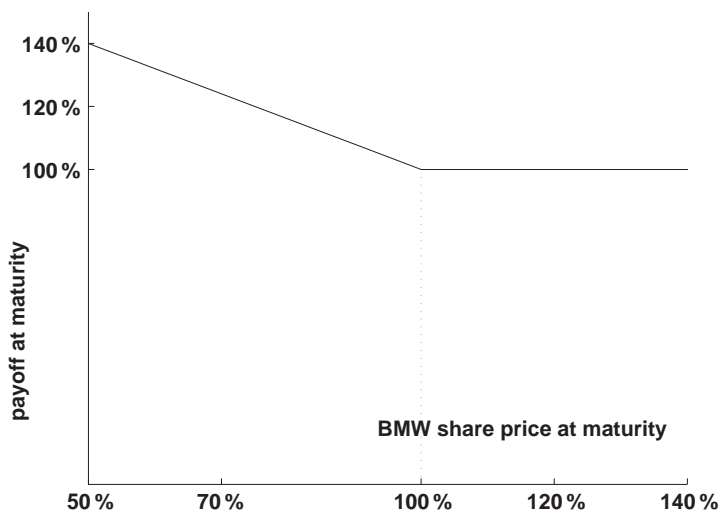


Figure 25.1 Payoff at maturity of a bearish capital guaranteed note, where the investor profits in 80 % of BMW's decline

bond and 80 % $((100 - 95.5)/5.625)$ participation in an 'at the money' put. The investor has € 10 million to invest and BMW's share price is trading at € 40. This means that the investor buys € 10 million worth of zero coupon bonds and € 8 million worth of BMW 'at the money' puts, which means that the capital guaranteed note has 200 thousand 'at the money' puts embedded in it. Since the investor only has 80 % participation in the 'at the money' put, the investor profits in only 80 % of BMW's decline. This means that, if after one year BMW has gone down by 50 %, the investor only makes 40 % on the invested notional and therefore gets paid € 14 million after one year, i.e. € 10 million from the zero coupon bond and € 4 million from the embedded put. The payoff profile after 1 year looks like Figure 25.1.

25.2 ATTRACTIVE MARKET CONDITIONS FOR CAPITAL GUARANTEED PRODUCTS

The market conditions under which capital guaranteed products look most attractive are when the option participation and therefore the exposure to the equity markets are largest. This can be achieved in two ways. Either the present value of the zero coupon bond is low, in which case

there is more premium available to invest in the option and therefore the participation in the equity markets is larger. Or the option premium is low, which also increases the participation in the equity markets. Obviously, the present value of the zero coupon bond is lower when interest rates are higher. Option premiums are lower when first of all the implied volatility is lower. For a bearish capital guaranteed product, the option premium is also lower when the dividend yield of the underlying is lower, since a bearish capital guaranteed product has a put embedded in it. A bullish capital guaranteed product comprises a zero coupon bond plus a participation in a call option and hence the participation will be larger for a higher dividend yield. In summary, these are the market conditions that make a bullish capital guaranteed product look most attractive, i.e. the call option participation is highest:

- The higher the interest rate the lower the present value of the zero coupon bond and therefore the more premium available to invest in the call option participation embedded in the bullish capital guaranteed product.
- The lower the implied underlying volatility the lower any option premium and thus the lower the call option premium, which results in a higher participation in a call option for the bullish capital guaranteed product.
- The higher the dividend yield the lower the call option premium and therefore the higher the participation in the call option embedded in the bullish capital guaranteed product.

For a bearish capital guaranteed product the most attractive market conditions can be summarised as follows:

- The higher the interest rate the lower the present value of the zero coupon bond and therefore the more premium available to invest in the put option participation embedded in the bullish capital guaranteed product.
- The lower the implied underlying volatility the lower any option premium and thus the lower the put option premium, which results in a higher participation in a put option for the bearish capital guaranteed product.
- The lower the dividend yield the lower the put option premium and therefore the higher the participation in the put option embedded in the bearish capital guaranteed product.

25.3 EXPOSURE PRODUCTS FOR THE CAUTIOUS EQUITY INVESTOR

The reverse convertible, as discussed in Chapter 15, is typically the type of investment that a cautious equity investor who has some appetite for risk would buy. Since the reverse convertible is already discussed in Chapter 15, this sub-section takes the opportunity to show yet another iteration of a financial note that the cautious equity investor with some risk appetite would invest in. This product is sometimes referred to as an *airbag* and, like the reverse convertible, this note gives exposure to the equity markets through a down-and-in put or regular put. As an additional feature, the coupon on this note is conditional on a certain lower bound of the stock price and therefore pays a higher coupon than the reverse convertible. In effect, the note is slightly leveraged on the downside as the conditional coupon means that the investor not only loses money through the put or down-and-in put on the downside, but can also lose his coupon if the underlying devalues too much. This conditional coupon is often priced as a down-and-out put. The pricing is shown by the following example.

Consider an investor who wants to buy a 1 year reverse convertible on BMW with a conditional coupon. Suppose that the one year 100/80 % knock-in reverse convertible pays a bullet coupon at maturity of 7.5 %. The investor is prepared to lose his coupon if, at any time during the life of the note, the stock drops below 85 % of BMW's initial share price. Because the coupon is conditional, the coupon on this note is higher than 7.5 %. Unlike the reverse convertible, where one solves for the coupon that makes the fair value or re-offer 100 %, the reverse convertible with conditional coupon does not solve for the coupon directly but solves for it indirectly through the redemption of the zero coupon bond and the rebate of a zero¹ strike down-and-out put with rebate. This rebate is equal to the excess redemption over 100 % on the zero coupon bond. In other words, the excess redemption over 100 % and the rebate of the down-and-out put are used as variables to make the reverse convertible with conditional coupon worth 100 %. The conditional coupon appears to be 10 %, which means that the following structure is worth 100 %.

- The price of a 1 year zero coupon bond with a redemption of 110 % (trader sells) instead of 100 %. **Minus**

¹ The fact that it is a zero strike down-and-out put with rebate means that it only gives a payout when it knocks out and therefore pays the rebate.

- The price of a 1 year 100/80 % down-and-in put (trader buys), which the trader prices as a 1 year 100/77 % down-and-in put because of the delta change over the barrier. Like the reverse convertible, the down-and-in put is used to enhance the coupon. **Minus**
- The price of a 0/85 % down-and-out put with a 10 % rebate (trader buys), which he prices as a 0/83 % down-and-out put because of the delta change over the barrier (the trader makes money because of the rebate as BMW's stock price goes down and through the barrier). Pricing the *barrier shifted* down-and-out put decreases the price of the down-and-out put with rebate and therefore the trader buys it at lower premium. The rebate has a delayed payment and is paid at maturity of the full structure. The rebate cancels out the excess redemption over 100 %, in this case 10 %, from the zero coupon bond and makes sure that the investor only receives 100 % at maturity if at any point the 85 % barrier is breached. It is important to understand that the rebate and the excess redemption over 100 % from the zero coupon bond are the same and therefore in solving the conditional coupon, the rebate of the down-and-out put and the redemption of the zero coupon bond are linked.

The payoff profile of the reverse convertible with conditional coupon is shown in Figure 25.2. The dotted payoff is to show that this only

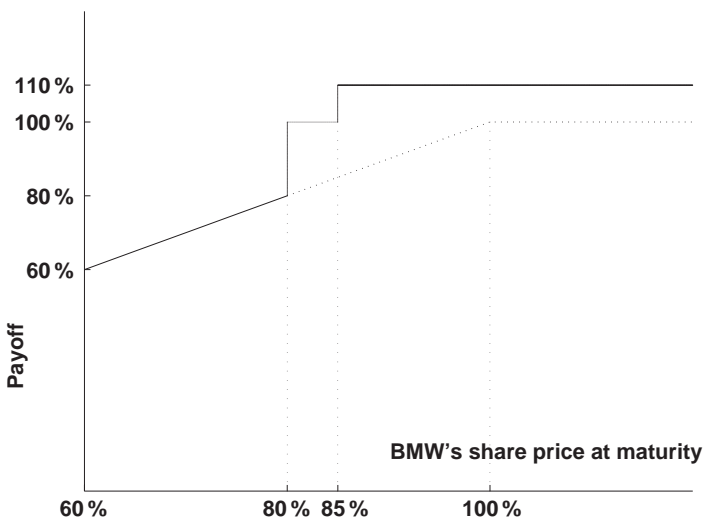


Figure 25.2 Payoff at maturity of a reverse convertible with a 10 % conditional coupon, where the coupon is conditional on BMW >85 % and the investor is short a 100/85 % knock-in put on BMW

occurs when the down-and-in put has actually knocked in. Obviously, the market conditions under which the conditional coupon is highest are when interest rates are high and when the implied volatility and the dividend yield of the underlying stock are high.

There are many iterations of the reverse convertible note with conditional coupon. For example, instead of being protected on the downside through the down-and-in put, one can think of a note where the investor is fully exposed to the downside of BMW. The 10 % coupon is still conditional on BMW never breaching 85 % of BMW's initial level. In return for giving up the protection through the down-and-in put and therefore participating in the full downside of BMW's share price, the investor participates in the full upside over 110 % up to 135 %. This structure can be priced as follows and is also worth 100 %. Although the structure looks complicated, it can still be priced as a series of options and barrier options. However, the pricing is quite convoluted but it shows the potential of creating innovative structures by financially engineering combinations of existing options rather than creating new option payoffs. Therefore, the pricing below can better be seen as an exercise than an addition to theory.

- A zero strike call which the trader sells. This zero strike call is worth 100 % minus the dividend yield and ensures that the investor is exposed to both the downside and the upside.
- A 110/100 % put spread which the trader also sells, where both puts knock out at 85 %. The 110/100 % put spread, which knocks out at 85 %, together with the zero strike call ensure that the investor receives a 10 % coupon provided that the share price never breaches 85 %, while at the same time fully exposing the investor to the downside.
- A 135 % call which the trader buys and ensures that the investor's participation to the upside is capped to 135 %.
- A 100/110 % call spread which the trader buys. However, both calls in the call spread are down-and-in calls and only knock in when BMW's share price breaches 85 %. This ensures that the investor does not participate in the first 10 % upside in BMW's share price when the 85 % level has been breached. When the 85 % has not been breached, the zero strike call in combination with the 110/100 % put spread ensures that the investor receives a 10 % coupon while being fully exposed to the downside, i.e. up to 100 % the payout to the investor is 10 % higher than one would expect based on the zero strike call.




Figure 25.3 Payoff at maturity of a reverse convertible with a 10 % conditional coupon, where the coupon is conditional on $\text{BMW} > 85\%$ and the investor has full exposure to the downside as well as the upside from 110 % up to 135 %

Therefore this 100/110 % call spread need only be live once the 85 % level has been breached. Hence, they are down-and-in calls.

Theoretically, one could apply a downward barrier shift to the 100/85 % down-and-in call. However, the value change over the barrier is minimal and the exposure to any delta change is also partially offset by the short 110/85 % down-and-in call.

Figure 25.3 shows the payoff profile of the above iteration of the reverse convertible with conditional coupon. The dotted line shows the payoff to the investor when the 85 % barrier has been breached, in which case the investor loses his conditional coupon and hence the 10 % parallel shift in payout.

25.4 LEVERAGED PRODUCTS FOR THE RISK SEEKING INVESTOR

The real risk seeking investor is prepared to give up any coupon for the sake of leveraging his view on a stock or the market. The perfect note such an investor would invest in is a structure where the investor buys a bond and, when he is bullish, he sells a put to gain additional




Figure 25.4 Payoff at maturity of a geared bullish note where the investor is fully exposed to the downside and participates in 150 % of the upside

participation in a call. When the investor is bearish, he sells a call option to gain additional participation in a put option.

As an example, consider an investor who is bullish on BMW and therefore sells a 1 year ‘at the money’ put in order to gear his exposure to the upside. In other words, the investor buys a 1 year zero coupon bond, sells a 1 year ‘at the money’ put and buys $X\%$ of a 1 year ‘at the money’ call. X is solved for, such that the full structure is worth 100 %. It appears that $X\%$ is 150 %, which means that the investor participates in 1.5 times the upside of BMW’s share price. The payoff profile of this note, which is sometimes referred to as a supertracker, is shown in Figure 25.4.

Now consider an investor who is bearish on BMW and therefore buys a bearish geared note. In this case the investor buys a 1 year zero coupon bond, sells a 1 year ‘at the money’ call and buys a 1 year $X\%$ geared ‘at the money’ put. When $X\%$ is equal to 233 %, the full bearish geared note appears to be worth 100 %.² The payoff profile of the bearish geared note is shown in Figure 25.5. The above shows that the bearish note works

² For the sake of completeness the numbers that are used in both the bullish and the bearish geared notes are that the 1 year zero coupon bond is worth 97 %, the 1 year ‘at the money’ call is worth 4 % and the 1 year ‘at the money’ put is worth 3 %.




Figure 25.5 Payoff at maturity of a geared bearish note where the investor is fully exposed to the upside and participates in 233 % of the downside

along the same lines as the bullish note with the caveat that the bearish note has some credit exposure to the investor that the bullish note does not have. Since the investor is owed 100 % at maturity and the put can never be worth more than 100 %, the bullish note gives no credit exposure to the investor in the note. However, with the bearish note the investor is short an 'at the money' call which is worth more than 100 % if BMW's share price is higher than 200 % and, since the investor is only owed 100 %, this bearish note does give some credit exposure to the investor. For this reason, the bearish note typically has a 200 % call embedded in it which the investor buys. This will slightly decrease the gearing on the put in the bearish note.

Appendix A

Variance of a Composite Option and Outperformance Option

It is relatively easy to prove that the variance of a composite option is equal to

$$\sigma_{compo}^2 = \sigma_S^2 + 2\rho\sigma_S\sigma_{FX} + \sigma_{FX}^2. \quad (A.1)$$

One just has to compare the movement of the underlying in the composite currency with the movement of a geometric basket option composed of the underlying in the local currency and the FX rate. The formula for this basket option is as follows:

$$F_t = S_{t,local}S_{t,FX}, \quad (A.2)$$

where

$$\frac{dS_{t,local}}{S_{t,local}} = r_{local}dt + \sigma_S dW_{1,t}, \quad (A.3)$$

and¹

$$\frac{dS_{t,FX}}{S_{t,FX}} = (r_{compo} - r_{local}) dt + \sigma_{FX} dW_{2,t}. \quad (A.4)$$

Since

$$\begin{aligned} Var[\ln(F_t)] &= Var[\ln(S_{t,local}S_{t,FX})] \\ &= Var[\ln(S_{t,local}) + \ln(S_{t,FX})] \\ &= (\sigma_S^2 + 2\rho\sigma_S\sigma_{FX} + \sigma_{FX}^2) t, \end{aligned} \quad (A.5)$$

and

$$E[\ln(F_t)] = E[\ln(S_{t,local}S_{t,FX})] \quad (A.6)$$

$$= r_{local}t + (r_{compo} - r_{local})t \quad (A.7)$$

$$= r_{compo}t, \quad (A.8)$$

¹ To get an intuitive feel for equation A.4, imagine an exchange rate with zero volatility and therefore the value of the exchange rate only changes because of differences in risk free interest rates between the currencies. In this case $S_{t+dt,FX} = S_{t,FX} \frac{1+r_{compo}dt}{1+r_{local}dt}$. Taylor expansion gives that this is equal to $S_{t,FX}(1 + (r_{compo} - r_{local})dt)$.

the movement of the stock in the composite currency can be modeled according to equation 20.1.

Assuming that another stock Y_t behaves according to the following process

$$\frac{dY_{t,local}}{Y_{t,local}} = r_{local}dt + \sigma_Y dW_{3,t}, \quad (A.9)$$

one can easily prove that the variance of the process $F_t = \frac{S_t}{Y_t}$ is equal to

$$\sigma_F^2 = \sigma_S^2 + \sigma_Y^2 - 2\rho\sigma_S\sigma_Y. \quad (A.10)$$

If the dividend yield is zero, the mean for F_t is zero. Namely,

$$\begin{aligned} Var[\ln(F_t)] &= Var\left[\ln\left(\frac{S_{t,local}}{Y_{t,local}}\right)\right] \\ &= Var[\ln(S_{t,local}) - \ln(Y_{t,local})] \\ &= \sigma_S^2 + \sigma_Y^2 - 2\rho\sigma_S\sigma_Y, \end{aligned} \quad (A.11)$$

and

$$\begin{aligned} E[\ln(F_t)] &= E[\ln(S_{t,local}) - \ln(Y_{t,local})] \\ &= 0 \end{aligned} \quad (A.12)$$

Appendix B

Replicating the Variance Swap

The variance swap can be replicated by $\frac{1}{K^2}$ options on all different strikes K . Since out of the money options are generally more liquid, a variance swap is replicated by both calls and puts. Let S^c be the cutoff above which the replicating option strikes are calls and below S^c the replicating option strikes are puts. Now it is easy to prove that the payoff of this replicating portfolio at maturity is:

$$\Pi(S_T) = \frac{S_T - S^c}{S^c} - \ln\left(\frac{S_T}{S^c}\right). \quad (\text{B.1})$$

Indeed, there are two possibilities; either $S_T > S^c$, in which case the puts are worthless and the payoff is merely the integration of all the calls, or $S_T < S^c$ and all the calls are worthless and the payoff is merely an integration of all the puts. First consider the scenario that $S_T > S^c$

$$\begin{aligned} \Pi(S_T) &= \int_0^{S^c} \frac{1}{K^2} \max[K - S_T, 0] dk + \int_{S^c}^{\infty} \frac{1}{K^2} \max[S_T - K, 0] dk \\ &= \int_{S^c}^{S_T} \frac{1}{K^2} (S_T - K) dk \\ &= \int_{S^c}^{S_T} -\frac{1}{K} + \frac{S_T}{K^2} dk \\ &= \left[-\ln(K) - \frac{S_T}{K} \right]_{S^c}^{S_T} \\ &= \frac{S_T - S^c}{S^c} - \ln\left(\frac{S_T}{S^c}\right) \end{aligned}$$

If $S_T < S^c$ one gets

$$\begin{aligned}
 \Pi(S_T) &= \int_0^{S^c} \frac{1}{K^2} \max[K - S_T, 0] dk + \int_{S^c}^{\infty} \frac{1}{K^2} \max[S_T - K, 0] dk \\
 &= \int_0^{S^c} \frac{1}{K^2} (K - S_T) dk \\
 &= \int_0^{S^c} \frac{1}{K} - \frac{S_T}{K^2} dk \\
 &= \left[\ln(K) + \frac{S_T}{K} \right]_0^{S^c} \\
 &= \frac{S_T - S^c}{S^c} - \ln\left(\frac{S_T}{S^c}\right)
 \end{aligned}$$

To prove that a variance swap can at any time be replicated by $1/K^2$ options on each strike, one not only needs an expression for $\Pi(S_T)$, the value of the replicating portfolio at maturity, but also for any t . Equation 23.4 shows that $\Pi(S_t)$ can be expressed as

$$\begin{aligned}
 \Pi(S_t) &= \int_0^{S^c} \frac{1}{K^2} [Ke^{-r(T-t)}N(-d_2) - S_t N(-d_1)] dk \\
 &\quad + \int_{S^c}^{\infty} \frac{1}{K^2} [S_t N(d_1) - Ke^{-r(T-t)}N(d_2)] dk \\
 &= \frac{S_t - S^c}{S^c} - \ln\left(\frac{S_t}{S^c}\right) + \frac{\sigma^2(T-t)}{2}.
 \end{aligned} \tag{B.2}$$

The proof for expression B.2 is quite prolonged. Nonetheless this appendix gives it to first of all allow the mathematically interested reader to check the validity of the replicating method and secondly it can serve as a point of reference.

The proof is split up into two parts. First it is shown that $\Pi(S_t)$ with a cutoff equal to S^c , $\Pi(S_t, S^c)$, minus $\Pi(S_t)$ with a cutoff of S_t , $\Pi(S_t, S_t)$, is equal to

$$\frac{S_t - S^c}{S^c} - \ln\left(\frac{S_t}{S^c}\right). \tag{B.3}$$

The second part of the proof shows that $\Pi(S_t, S_t)$ is equal to

$$\frac{\sigma^2(T-t)}{2}. \quad (\text{B.4})$$

Adding the first and second parts is obviously equal to the replicating portfolio of $1/K^2$ options per strike where a cutoff of S^c , $\Pi(S_t, S^c)$, is used and is in turn equal to the expression in equation B.2.

For simplicity, both parts are proved assuming that the interest rate is zero. The first part makes use of the following feature of a standard normal distribution

$$\begin{aligned} 1 - N(d_1) &= N(-d_1) \\ 1 - (N(d_2)) &= N(-d_2). \end{aligned}$$

Using this knowledge it is easy to prove the first part

$$\begin{aligned} \Pi(S_t, S^c) - \Pi(S_t, S_t) &= \int_0^{S^c} \frac{1}{K^2} [K - KN(d_2) - S_t + S_t N(d_1)] dk \\ &\quad - \int_0^{S_t} \frac{1}{K^2} [K - KN(d_2) - S_t + S_t N(d_1)] dk \\ &\quad + \int_{S^c}^{\infty} \frac{1}{K^2} [S_t N(d_1) - KN(d_2)] dk \\ &\quad + \int_{S_t}^{\infty} \frac{1}{K^2} [S_t N(d_1) - KN(d_2)] dk \\ &= \int_{S_t}^{S^c} \frac{1}{K^2} [K - KN(d_2) - S_t + S_t N(d_1)] dk \\ &\quad - \int_{S_t}^{S^c} \frac{1}{K^2} [S_t N(d_1) - KN(d_2)] dk \\ &= \int_{S_t}^{S^c} \frac{1}{K^2} [K - S_t] dk \\ &= \frac{S_t - S^c}{S^c} - \ln\left(\frac{S_t}{S^c}\right). \end{aligned}$$

The proof for the second part is quite daunting and is therefore only for the readers that are interested in the rigorous mathematics behind the replicating portfolio of the variance swap. The proof starts by substituting d_1 and d_2 into their respective positions and subsequently applying

several shifts to the integral.¹

$$\begin{aligned}
\Pi(S_t, S_t)^{\alpha=\sigma\sqrt{T-t}} &= \int_0^{S_t} \frac{1}{K^2} [KN(-d_2) - S_t N(-d_1)] dk \\
&\quad + \int_{S_t}^{\infty} \frac{1}{K^2} [S_t N(d_1) - KN(d_2)] dk \\
&= \int_0^{S_t} \frac{1}{K^2} \left[KN\left(-\frac{\ln(\frac{S_t}{K})}{\alpha} + \frac{\alpha}{2}\right) \right. \\
&\quad \left. - S_t N\left(-\frac{\ln(\frac{S_t}{K})}{\alpha} - \frac{\alpha}{2}\right) \right] dk \\
&\quad + \int_{S_t}^{\infty} \frac{1}{K^2} \left[S_t N\left(\frac{\ln(\frac{S_t}{K})}{\alpha} + \frac{\alpha}{2}\right) \right. \\
&\quad \left. - KN\left(\frac{\ln(\frac{S_t}{K})}{\alpha} - \frac{\alpha}{2}\right) \right] dk \\
&\stackrel{K=S_t e^l}{=} \int_{-\infty}^0 \left[S_t e^l N\left(\frac{l}{\alpha} + \frac{\alpha}{2}\right) - S_t N\left(\frac{l}{\alpha} - \frac{\alpha}{2}\right) \right] \frac{dl}{S_t e^l} \\
&\quad + \int_0^{\infty} \left[S_t N\left(-\frac{l}{\alpha} + \frac{\alpha}{2}\right) - S_t e^l N\left(-\frac{l}{\alpha} - \frac{\alpha}{2}\right) \right] \frac{dl}{S_t e^l} \\
&= \int_{-\infty}^0 \left[N\left(\frac{l}{\alpha} + \frac{\alpha}{2}\right) - e^{-l} N\left(\frac{l}{\alpha} - \frac{\alpha}{2}\right) \right] dl \\
&\quad + \int_0^{\infty} \left[e^{-l} N\left(-\frac{l}{\alpha} + \frac{\alpha}{2}\right) - N\left(-\frac{l}{\alpha} - \frac{\alpha}{2}\right) \right] dl \\
&= \int_0^{\infty} \left[N\left(-\frac{l}{\alpha} + \frac{\alpha}{2}\right) - e^l N\left(-\frac{l}{\alpha} - \frac{\alpha}{2}\right) \right] dl \\
&\quad + \int_0^{\infty} \left[e^{-l} N\left(-\frac{l}{\alpha} + \frac{\alpha}{2}\right) - N\left(-\frac{l}{\alpha} - \frac{\alpha}{2}\right) \right] dl \\
&= \int_0^{\infty} \left[(1 + e^{-l}) N\left(-\frac{l}{\alpha} + \frac{\alpha}{2}\right) \right. \\
&\quad \left. - (1 + e^l) N\left(-\frac{l}{\alpha} - \frac{\alpha}{2}\right) \right] dl
\end{aligned}$$

¹ This proof is thanks to Alex Boer.

$$\begin{aligned}
&= \int_0^\infty \left[\int_{-\infty}^{-\frac{l}{\alpha} + \frac{\alpha}{2}} \frac{(1 + e^{-l})}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \right. \\
&\quad \left. - \int_{-\infty}^{-\frac{l}{\alpha} - \frac{\alpha}{2}} \frac{(1 + e^{-l})}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \right] dl \\
&= \int_0^\infty \left[\int_{-\infty}^0 \frac{(1 + e^{-l})}{\sqrt{2\pi}} e^{-\frac{(x - \frac{l}{\alpha} + \frac{\alpha}{2})^2}{2}} dx \right. \\
&\quad \left. - \int_{-\infty}^0 \frac{(1 + e^{-l})}{\sqrt{2\pi}} e^{-\frac{(x - \frac{l}{\alpha} - \frac{\alpha}{2})^2}{2}} dx \right] dl \\
&= \int_0^\infty \left[\int_{-\infty}^0 \left(e^{\frac{l}{2}} + e^{-\frac{l}{2}} \right) e^{-\frac{l}{2}} e^{-\frac{(x^2 + \frac{l^2}{\alpha^2} + \frac{\alpha^2}{4} - \frac{2lx}{\alpha} + \frac{x}{\alpha} - l)}{2}} \right] \frac{dx dl}{\sqrt{2\pi}} \\
&\quad - \int_0^\infty \left[\int_{-\infty}^0 \left(e^{\frac{l}{2}} + e^{-\frac{l}{2}} \right) e^{\frac{l}{2}} e^{-\frac{(x^2 + \frac{l^2}{\alpha^2} + \frac{\alpha^2}{4} - \frac{2lx}{\alpha} - \frac{x}{\alpha} + l)}{2}} \right] \frac{dx dl}{\sqrt{2\pi}} \\
&\stackrel{x=-x}{=} \int_0^\infty \left[\int_0^\infty \left(e^{\frac{l}{2}} + e^{-\frac{l}{2}} \right) e^{-\frac{x^2}{2} - \frac{l^2}{2\alpha^2} - \frac{\alpha^2}{8}} \right. \\
&\quad \left. \times \left(e^{-\frac{xl}{\alpha} + \frac{x\alpha}{2}} - e^{-\frac{xl}{\alpha} - \frac{x\alpha}{2}} \right) \right] \frac{dx dl}{\sqrt{2\pi}} \\
&\stackrel{l=2\alpha l, \ x=2x}{=} \int_0^\infty \left[\int_0^\infty \left(e^{\alpha l} + e^{-\alpha l} \right) e^{-2x^2 - 2l^2 - \frac{\alpha^2}{8}} \right. \\
&\quad \left. \times \left(e^{\alpha x} - e^{-\alpha x} \right) e^{-4xl} \frac{4\alpha}{\sqrt{2\pi}} \right] dx dl \\
&= \int_0^\infty \left[\int_0^\infty \cosh(\alpha l) \sinh(\alpha x) e^{-2(x+l)^2} e^{-\frac{\alpha^2}{8}} \frac{16\alpha}{\sqrt{2\pi}} \right] dx dl \\
&\stackrel{p=\frac{x+l}{2}, \ q=\frac{x-l}{2}}{=} \int_0^\infty \left[\int_{-p}^p e^{-8p^2} \cosh(\alpha(p-q)) \sinh(\alpha(p+q)) \right. \\
&\quad \left. \times e^{-\frac{\alpha^2}{8}} \frac{32\alpha}{\sqrt{2\pi}} \right] dq dp \\
&= \int_0^\infty \left[e^{-8p^2} e^{-\frac{\alpha^2}{8}} \frac{32\alpha}{\sqrt{2\pi}} \int_{-p}^p \cosh(\alpha(p-q)) \right. \\
&\quad \left. \times \sinh(\alpha(p+q)) dq \right] dp
\end{aligned}$$

$$\begin{aligned}
&= \int_0^\infty \left[e^{-8p^2} e^{-\frac{\alpha^2}{8}} \frac{32\alpha}{\sqrt{2\pi}} \sinh(2\alpha p) \cdot p \right] dp \\
&= \frac{16\alpha}{\sqrt{2\pi}} \int_0^\infty \left[\left(e^{-8p^2+2\alpha p-\frac{\alpha^2}{8}} - e^{-8p^2-2\alpha p-\frac{\alpha^2}{8}} \right) p \right] dp \\
&= \frac{16\alpha}{\sqrt{2\pi}} \int_0^\infty \left[p e^{-\left(\sqrt{8}p-\frac{\alpha}{\sqrt{8}}\right)^2} - p e^{-\left(\sqrt{8}p+\frac{\alpha}{\sqrt{8}}\right)^2} \right] dp \\
&= \frac{16\alpha}{\sqrt{2\pi}} \int_0^\infty \left[p e^{-\frac{(4p-\frac{\alpha}{2})^2}{2}} - p e^{-\frac{(4p+\frac{\alpha}{2})^2}{2}} \right] dp \\
&\stackrel{p=\frac{z}{4}}{=} \frac{16\alpha}{\sqrt{2\pi}} \int_0^\infty \left[\frac{z}{4} e^{-\frac{(z-\frac{\alpha}{2})^2}{2}} - \frac{z}{4} e^{-\frac{(z+\frac{\alpha}{2})^2}{2}} \right] \frac{dz}{4} \\
&= \frac{\alpha}{\sqrt{2\pi}} \left(\int_{-\frac{\alpha}{2}}^\infty \left[\left(z + \frac{\alpha}{2} \right) e^{-\frac{z^2}{2}} \right] dz \right. \\
&\quad \left. - \int_{\frac{\alpha}{2}}^\infty \left[\left(z - \frac{\alpha}{2} \right) e^{-\frac{z^2}{2}} \right] dz \right) \\
&= \frac{\alpha}{\sqrt{2\pi}} \left(\int_{-\frac{\alpha}{2}}^\infty \left[z e^{-\frac{z^2}{2}} \right] dz - \int_{\frac{\alpha}{2}}^\infty \left[z e^{-\frac{z^2}{2}} \right] dz \right) \\
&\quad + \frac{\alpha}{\sqrt{2\pi}} \left(\int_{-\frac{\alpha}{2}}^\infty \left[\frac{\alpha}{2} e^{-\frac{z^2}{2}} \right] dz + \int_{\frac{\alpha}{2}}^\infty \left[\frac{\alpha}{2} e^{-\frac{z^2}{2}} \right] dz \right) \\
&= \frac{\alpha}{\sqrt{2\pi}} \left(\int_{-\frac{\alpha}{2}}^\infty \left[\frac{\alpha}{2} e^{-\frac{z^2}{2}} \right] dz + \int_{\frac{\alpha}{2}}^\infty \left[\frac{\alpha}{2} e^{-\frac{z^2}{2}} \right] dz \right) \\
&= \frac{\alpha}{\sqrt{2\pi}} \cdot \frac{\alpha\sqrt{2\pi}}{2} \\
&= \frac{\alpha^2}{2} = \frac{\sigma^2(T-t)}{2}.
\end{aligned}$$

Bibliography

- Aarts, P. (1999). Pricing the basket option and exploring its volatility structure, *Math thesis Utrecht University*.
- Demterfi, K., Derman, E., Kamal, M. and Zou, J. (1999). *More Than You Ever Wanted To Know About Volatility Swaps*, Quantitative Strategies Research Notes, Goldman Sachs.
- Derman, E. (1992). *Outperformance Options*. Goldman Sachs & Co.
- Durrett, R. (1996). *Probability: Theory and Examples*, 2nd ed. Duxbury Press, Belmont California.
- Grimmett, G.R. and Stirzaker, D.R. (1992). *Probability and Random Processes*, 2nd ed. Clarendon Press, Oxford.
- Hull, J.C. (1993). *Options, Futures, and other Derivative Securities*, 2nd ed. Prentice-Hall Inc.
- Hull, J.C. (1997). *Options, Futures and other Derivatives*, Prentice-Hall Inc.
- Lamberton, D. and Lapeyre, B. (1996). *Introduction to Stochastic Calculus Applied to Finance*, Chapman & Hall, London.
- Roelfsema, M.R. (2000). *Exploiting the Dependencies between Index and Stock Options*, Delft University Press.
- de Weert, F.J. (2002). Basket en skew arbitrage, *Math thesis Utrecht University*.
- de Weert, F.J. (2006). *An Introduction to Options Trading*, John Wiley & Sons Ltd, Chichester.
- Wilmott, P. (1998). *Derivatives: The Theory and Practice of Financial Engineering*, University Edition, John Wiley & Sons Ltd, Chichester.

Index

- ADR *see* American Depositary Receipt
- airbags
 - concepts 160–3
 - definition 160
- American Depositary Receipt (ADR)
 - 126
- American options
 - concepts 4, 46–8
 - definition 4
- An Introduction to Options Trading*
 - (author) 1
- analytically short skew, forward starting
 - options 74–5
- annualised variance, variance swaps
 - 139–49, 169–74
- appendices 167–74
- arithmetic Asian out options
 - see also* Asian...
 - concepts 105, 107–10
 - delta hedging 109–10
 - duration 108–10
 - geometric averages 108, 110
 - the Greeks 107–10
 - implied volatility 108, 110
 - Monte Carlo process 107
 - pricing 107–8, 110
- Asian in forwards
 - concepts 112–18
 - definition 112
 - gamma 113–15
 - hedging 113–15
 - Monte Carlo simulations 114–18
 - optional early termination 116–18
 - pricing 114–18
 - volatility 115
- Asian in options
 - concepts 105, 110–18
 - definition 105
 - delta hedging 110–12
 - pricing 111
- Asian options
 - concepts 105–18
 - definition 105
 - duration 108–10
 - implied volatility 107, 108
 - Monte Carlo simulations 107, 114–18
 - pricing 105–18
 - risk 105, 109–12
 - types 105
- Asian out options
 - see also* arithmetic...; geometric...
 - concepts 105–10
 - definition 105
 - pricing 105–8
 - risk 105, 109–10
- at the money options (ATM) 16, 29–37,
 - 41–4, 73–4, 81, 84–91, 93–6,
 - 99–104, 110, 119–23, 131, 136,
 - 140–2, 148–9, 152–8, 160–3, 164–5
- gamma 16
- quanto options 119–23
- reverse convertibles 87–91, 93–6,
 - 99–104, 160–3
- skew 29–32, 34, 73–4
- autocallable notes
 - concepts 93–8, 99–104
 - conditional coupons 94–7
 - definition 93
 - down-and-in put barrier options 93–6,
 - 101–4

- autocallable notes (*Continued*)
 - early redemption features 95–6, 97–8, 101–2
 - examples 93–4
 - interest/equity correlation factors 98
 - maturity dates 94
 - pricing 95–8, 99–104
 - reverse convertibles 93–6, 99–104
 - unconditional coupons 94, 97–8
- barrier options 57–69, 77–8, 87–90, 93–6, 160–3
 - see also* down-and-in . . . ; up-and-out . . .
 - autocallable notes 93–6
 - barrier shifts 60–3, 64, 67–9, 96, 104, 161–3
 - concepts 57–69, 78, 87–90, 93–6, 160–3
 - delta 58–65, 161–3
 - double barrier options 57–8, 68–9
 - the Greeks 57–65
 - ladder options 78
 - lookback options 79–80
 - maturity dates 62, 65–6
 - Monte Carlo simulations 57, 58, 59–60, 64, 66, 68–9
 - Newton–Raphson process 102
 - pricing 57–69, 77–8, 89–90, 95–8, 162–3
 - retail demand 57–8
 - reverse convertibles 87–90, 93–6, 99–104
 - risk 57–69
 - single barrier options 57–68
 - skew risk 57, 58, 66–8
 - time to maturity 62
 - types 57–65, 68–9
 - volatility 60–2
- barrier parity, knock-out/knock-in barrier options 65
- basket options 138, 151–6, 167–8
 - see also* dispersion
 - Black–Scholes options pricing formula 151
 - concepts 138, 151–6, 167–8
 - definition 151
 - hybrids 138
 - implied volatility 152–3
 - Monte Carlo simulations 151
 - pricing 151
 - vega 153–5
 - volatility 151–6
- bearish strategies 37–9, 63, 157–9, 164–5
- best of options
 - concepts 135–8
 - correlation risk 135–7
 - definition 135
 - hybrids 138
 - payoff profiles 135–6
 - pricing 135
- binomial trees
 - concepts 46–8
 - Monte Carlo simulations 46–7, 48
- Black–Scholes options pricing
 - formula
 - basket options 151
 - Brownian motion 45
 - concepts 6–10, 12–14, 45, 84–5, 107, 120–1, 131, 151
 - delta 12–14
 - example 7
 - forwards 10
 - implied volatility 8, 45, 107
 - parameters 6, 8–9, 10, 12–14, 107, 120–1
- BMW shares 29–30, 35–6, 39–40, 46–8, 49–51, 53–4, 58–60, 63–4, 71–2, 77, 79–80, 83–4, 87–9, 91, 93–8, 99–104, 109–16, 129–38, 140–9, 157–65
- Boer, Alex 172
- bonds
 - coupons 87–91, 93–8, 99–104, 157–8, 160–3
 - leveraged products for risk-seeking investors 157, 163–5
 - reverse convertibles 87–91, 93–6, 99–104, 160–3
 - zero coupon bonds 157–8, 160–3
- BP shares 119–23
- Brazilian real *see* BRL
- break-even levels, simple option strategies 35–44
- British pound
 - USD composite options 125–8
 - USD quanto options 119–23
- BRL/EUR exchange rate, skew 33–4
- Brownian motion 45, 121
- bullish strategies 35–6, 39–40, 63, 157–9, 163–5

- call options
 - see also* Greeks
 - chooser options 49–51
 - concepts 1–8, 10–21, 23–4, 35–44, 158–65
 - definition 3
 - forward prices 10
 - payoff profiles 4–5
 - pricing 6–8, 10
 - put–call parity 11–12
 - up-and-out call barrier options 16–17, 58, 63–5, 67–9, 78–80, 103
- call spread option strategy
 - concepts 35–7, 53–5, 85, 95–6
 - digital options 1, 53–5, 62, 95–8
 - examples 35–6, 53–4
 - payoff profiles 36–7
 - ratchet options 85
 - skew 36–7, 85
- callable reverse convertibles
 - concepts 99–102
 - definition 99
- capital guaranteed products for risk
 - averse investors 157–9
 - attractive market conditions 158–9
 - payoff profiles 158
- cautious equity investors, exposure equity products 157, 160–3
- central banks, interest rates 28
- Central Limit Theorem 46
- chooser options
 - concepts 49–51
 - definition 49
 - example 50–1
 - strategy rationale 51
 - terms 49–50
- cliquet options
 - see also* ratchet options
 - concepts 83–5
 - definition 83
 - forward starting options 83–5
 - risk 85
- collars
 - see also* risk reversal strategies
 - concepts 39–40
 - definition 39
- composite options
 - concepts 125–8, 129–30, 167–8
 - definition 125
 - delta hedging 126–8, 132–3
 - example 125–6
 - FX exposure 126–8
 - implied volatility 126
 - outperformance options 129, 130–1
 - pricing 127–8
 - quanto options 125
 - variance 126, 167–8
- computer programmes, Monte Carlo simulations 48
- conditional coupons, autocallable notes 94–7
- conditional double barrier options
 - see also* double barrier options
 - concepts 68–9
- continuous dividends 10
- conventional options
 - see also* Greeks
 - concepts 1, 3–21, 157–65
 - financially engineered structures 157–65
 - forward prices 10
 - forward starting options 71–2
 - interest rate exposures 21
 - pricing 1, 6–10, 12–14, 34, 45–8, 71–2, 105–7
 - strategies 35–44
 - variance swap replication 140–9, 169–74
- convexity principle, vega 148–9
- correlation
 - best of options risk 135–7
 - composite options risk 125–6, 131–2
 - dispersion 151–6
 - exotic options 1
 - interest/equity correlation factors 98
 - outperformance options 131–2, 136–7
 - quanto options risk 119–22
 - worst of options risk 135, 137–8
- cost of carry, definition 9
- coupons 87–91, 93–8, 99–104, 157–8, 160–3
 - autocallable notes 93–8, 99–104
 - interest rates 90–1
 - reverse convertibles 87–91, 93–6, 99–104, 160–3
- covariance 152
- curvature parameter, implied volatility 31–2
- delta
 - see also* dynamic hedging
 - arithmetic Asian out options 109–10

delta (*Continued*)

- Asian options 109–12
 - barrier options 58–65, 161–3
 - Black–Scholes options pricing formula 12–14
 - collars 40
 - composite options 126–8, 132–3
 - concepts 12–14, 21, 25–6, 29–32, 40, 58–63, 90–1, 109–15, 122–3, 126–8, 132–3, 156
 - definition 12, 25
 - dispersion 156
 - down-and-in put barrier options 58–63, 161–3
 - extreme values 12
 - forward starting options 72–3
 - gamma 14–16, 23–4, 25–6, 27, 29, 54, 60, 75, 79–80, 110, 113–15, 142–5
 - interest rates exposures 21
 - outperformance options 132–3
 - quanto options 122–3
 - skew effects 29–32, 147–8, 156
 - variance swaps 141–5, 147–8, 156
- delta cash
- concepts 25–6, 144–5
 - definition 25
 - example 26
- digital options
- call spread option strategy 53–5, 95–6
 - concepts 1, 53–5, 62, 95–8
 - definition 53
 - payoff profiles 53–4
 - strike-selection criteria 55
- DIP *see* down-and-in put barrier options
- discounts 84–5, 89–90
- dispersion
- see also* basket options
 - concepts 151–6
 - correlation 151–6
 - definition 151
 - delta 156
 - ratio's vega 155
 - skew delta 156
 - trading possibilities 153–5
 - variance swaps 153, 155
 - vega 153–5
 - volatility 151–5
- dividends 6–12, 39–40, 120–2, 159
- Black–Scholes options pricing formula 6–7, 9–10, 120–1
 - collars 40
 - concepts 9–12, 39–40, 159
 - continuous dividends 10
 - forward starting options 72
 - forwards 8–9, 10, 39–40
 - premiums 159
- double barrier options
- see also* barrier options
 - concepts 57–8, 68–9
 - conditional double barrier options 68–9
 - definition 68
 - examples 68–9
 - risk 68–9
 - types 68–9
 - unconditional double barrier options 69
- down-and-in call barrier options, concepts 58, 163
- down-and-in put barrier options
- see also* barrier options
 - airbags 160–3
 - autocallable notes 93–6, 101–4
 - barrier shifts 60–3, 67–8, 96, 104, 161–3
 - concepts 57–63, 66, 87–90, 93–6, 101–4, 160–3
 - definition 58
 - delta 58–63, 161–3
 - exposure equity products for cautious investors 160–3
 - gamma 60
 - maturity dates 62, 66
 - payoff profiles 58–60, 66–7
 - pricing 58–63, 66, 89–90, 95–8, 101–4, 162–3
 - reverse convertibles 87–90, 93–6, 101–4
 - risk 58–60
 - share-buying situations 63
 - skew 58, 66–8
 - time to maturity 62
 - volatility factors 60–2
- down-and-out call barrier options, concepts 58, 68, 103–4
- down-and-out put barrier options, concepts 58, 103–4, 161–3
- downside options 27–34, 39–40, 67–8, 80–1
- duration, arithmetic Asian out options 108–10

- DVO1 91
- dynamic hedging
 - see also* delta
 - concepts 14, 21, 23–4, 26, 27, 30–2, 90–1, 109–15, 122–3, 132–3
 - definition 14
- dynamically short skew, forward starting options 74–5
- early redemption features, autocallable notes 95–6, 97–8, 101–2
- equity/FX contrasts, skew 32–4
- EUR/BRL exchange rate, skew 33–4
- EUR/USD exchange rate, skew 32–4
- European options 3–4, 11–18, 47–8, 49–51, 53–5, 58–69, 95–8
 - definition 3–4, 49
 - put–call parity 11–12
- event risk, variance swaps 146–7
- exercise
 - American options 4, 46–8
 - definition 3–4
- exotic options
 - correlation components 1
 - definition 1, 53, 87
 - full spectrum 1
 - gamma sign changes 16
 - pricing 1, 34, 45, 49
 - risks 1
 - skew 34, 35, 58–60
 - vanna/vomma 19–21
 - vega 16–17
- expected fair value of the stock, forwards 4–6
- expiration dates *see* maturity dates
- exposure equity products for cautious investors 157, 160–3
- falling markets, volatility 27–8
- financially engineered structures
 - see also* retail investors
 - capital guaranteed products for risk averse investors 157–9
 - concepts 157–65
 - exposure equity products for cautious investors 157, 160–3
 - leveraged products for risk-seeking investors 157, 163–5
- fixed strike lookback options
 - see also* lookback options
 - concepts 79–80
 - floating strike lookback options
 - see also* lookback options
 - concepts 79, 80–1
- foreign exchange *see* FX
- forward starting options
 - analytically short skew 74–5
 - cliquet options 83–5
 - concepts 71–5, 80–1, 83–5, 111
 - conventional options 71–2
 - definition 71
 - delta 72–3
 - dividends 72
 - dynamically short skew 74–5
 - examples 71–5
 - the Greeks 72–3, 75, 111
 - hedging 72–3, 74–5
 - implied volatility 72–4
 - lookback options 80–1
 - maturity dates 75
 - pricing 71–5, 80–1, 83–5, 111
 - risk 71, 72–5, 85
 - skew risk 71, 72–5, 85
 - strike prices 71–5, 111
 - term structure of the skew 73–4
 - vega 74–5, 80–1
- forwards
 - Asian in forwards 112–18
 - Black–Scholes options pricing formula 10
 - concepts 4–21, 51
 - cost of carry 9
 - definition 4–6, 8
 - dividends 8–9, 10, 39–40
 - expected fair value of the stock 4–6
 - interest rates 10, 39–40
 - payoff profiles 5–6
 - pricing 4–5, 8–10
 - strike prices 6, 8–9, 12
- FX
 - composite options 125–8, 167–8
 - quanto options 119–23
 - skew equity contrasts 32–4
 - volatility 120–3, 126–8
- gamma 1, 14–16, 18, 23–9, 54, 60, 75, 79–80, 110, 113–15, 142–5, 146–7
 - arithmetic Asian out options 110
 - Asian in forwards 113–15
 - at the money options 16
 - concepts 1, 14–16, 18, 23–9, 54, 60, 75, 79–80, 110, 113–15

- gamma (*Continued*)
 - definition 14
 - delta 14–16, 23–4, 25–6, 27, 29, 54, 60, 75, 79–80, 110, 113–15, 142–5
 - down-and-in put barrier options 60
 - examples 23–4, 26
 - exotic options 16
 - hidden gamma 113–14
 - importance 16, 23–4
 - lookback options 79–80
 - profit-determination role 1, 16, 23–4, 27, 29, 54, 113–14, 143–4, 146–7
 - theta 23–4
 - variance swaps 142–5, 146–7
 - vega comparisons 18
 - gamma cash
 - concepts 25–6, 143–6
 - definition 25
 - gearing 39–40, 53–5, 66, 95–6, 164–5
 - see also* leveraged products ...
 - geometric Asian out options
 - see also* Asian ...
 - concepts 105–7
 - implied volatility 107, 108
 - geometric averages, arithmetic Asian out options 108, 110
 - geometric basket options 167–8
 - Girsanov theorem, concepts 7
 - the Greeks
 - see also* delta; gamma; theta; vega
 - arithmetic Asian out options 107–10
 - Asian in options 111
 - barrier options 57–65
 - concepts 1, 3, 12–21, 57, 63
 - forward starting options 72–3, 75, 111
 - variance swaps 140–7
 - GSK shares 125–8
 - hedging 8–9, 11–12, 14, 21, 23–4, 26, 27, 30–2, 53–5, 58–60, 63–4, 72–5, 90–1, 105, 109–10, 113–15, 122–3, 126–8, 132–3
 - hidden gamma, Asian in forwards 113–14
 - higher order derivatives, concepts 19–21
 - hybrids, concepts 138
 - implied volatility
 - see also* vega; volatility; vomma
 - arithmetic Asian out options 108, 110
 - Asian options 107, 108
 - basket options 152–3
 - Black–Scholes options pricing formula 8, 45, 107
 - call spread option strategy 36–7
 - composite options 126
 - concepts 8, 16–18, 19–21, 27–34, 36–7, 64–5
 - curvature parameter 31–2
 - definition 8
 - forward starting options 72–4
 - geometric Asian out options 107, 108
 - premiums 159
 - reverse convertibles 90
 - skew 27–34, 35, 36–7, 58, 66–8, 71, 72–5, 85, 141–2, 147–8, 156
 - smiles 29–34
 - strike prices 27–34, 35, 36–7, 58, 66–8, 71, 72–5, 85, 141–2, 147–8, 156
 - up-and-out call barrier options 16–17, 64–5, 67–8
 - variance swaps 141–9
 - in the money options 12–14, 16, 19, 27, 30–2, 50–1, 59–61, 63–4, 69, 110–12
 - indices
 - autocallable notes 93–8
 - options 28
 - skew 156
 - interest rate hedge, reverse convertibles 90–1
 - interest rates 6–7, 11–12, 18–19, 21, 27–8, 39–40, 90–1, 106–7, 126, 167
 - see also* risk free ...
 - central banks 28
 - concepts 6–7, 11–12, 18–19, 21, 27–8, 39–40, 90–1, 106–7, 126, 167
 - coupons 90–1
 - delta hedging 21
 - forwards 10, 39–40
 - future rates 27–8
 - liquid markets 7
 - reverse convertibles 90–1
 - zero coupon bonds 159
- interest/equity correlation factors, autocallable notes 98
- investment banks 85, 113–15
 - investors
 - barrier options 57–8
 - capital guaranteed products for risk averse investors 157–9

- cautious equity investors 157, 160–3
- exposure equity products for cautious investors 157, 160–3
- fears 28, 157–65
- financially engineered structures 157–65
- leveraged products for risk-seeking investors 157, 163–5
- risk attitudes 157–65
- risk averse investors 157–9
- risk-seeking investors 157, 163–5
- types 157–65
- Ito's lemma, concepts 7
- Jensen's inequality 108
- knock-in reverse convertibles
 - see also* reverse convertibles
 - concepts 87–90
- knock-out/knock-in barrier options
 - see also* barrier options
 - barrier parity 65
- ladder options
 - barrier options 78
 - concepts 77–8, 79–80
 - definition 77
 - example 77
 - lookback options 79–80
 - pricing 78, 79–80
- leveraged products for risk-seeking investors 157, 163–5
 - see also* gearing
- liquid markets, interest rates 7
- lognormal returns, standard deviation 145–6
- lookback options
 - barrier options 79–80
 - concepts 79–81
 - definition 79
 - fixed strike lookback options 79–80
 - floating strike lookback options 79, 80–1
 - forward starting options 80–1
 - gamma 79–80
 - ladder options 79–80
 - pricing 79–81
 - risk 80–1
 - skew 81
 - types 79
- maturity dates 3–21, 62, 65–6, 75
 - see also* time to maturity
 - autocallable notes 94
 - barrier options 62, 65–6
 - definition 3–4
 - down-and-in put barrier options 62, 66
 - forward starting options 75
 - reverse convertibles 88
 - skew 28–30
- Monte Carlo simulations 1, 45–8, 57, 59–60, 64, 66, 68–9, 71–3, 79–80, 107, 114–18, 135, 151
 - arithmetic Asian out options 107
 - Asian in forwards 114–18
 - Asian options 107, 114–18
 - barrier options 57, 58, 59–60, 64, 66, 68–9
 - basket options 151
 - binomial trees 46–7, 48
 - computer programmes 48
 - concepts 45–8
 - exotic options 45
 - pricing 45–8, 57, 58, 59–60, 64, 66, 68–9, 71–3, 79–80, 107, 114–18, 135, 151
 - principle 45–6
 - workings 48
- Newton–Raphson process, concepts 102, 117
- Normal distributions 7, 139–40, 141–2, 145–6
- notional, definition 119
- options
 - see also* call ...; conventional ...; exotic ...; Greeks; pricing; put ...
 - concepts 1, 3–21, 157–65
 - financially engineered structures 157–65
 - forward prices 10
 - indices 28
 - interest rate exposures 21
 - premiums 21, 31–2, 90, 123, 127–8, 132–3, 158–9
 - strategies 35–44, 51–69, 85, 157–65
- out of the money options 12–14, 27

- outperformance options
 - composite options 129, 130–1
 - concepts 129–33, 136–7, 167–8
 - correlation position 131–2, 136–7
 - definition 129
 - delta hedging 132–3
 - example 129–33
 - hybrids 138
 - payoff profiles 129, 130–1
 - pricing 129–31
 - variance 131, 167–8
- payoff profiles
 - best of options 135–6
 - call options 4–5
 - call spread option strategy 36–7
 - capital guaranteed products 158
 - collars 39–40
 - concepts 4–8
 - digital options 53–4
 - down-and-in put barrier options 58–60, 66–7
 - forwards 5–6
 - outperformance options 129, 130–1
 - put options 4–5
 - put spread option strategy 37–9
 - straddles 42, 44
 - strangles 43–4
 - supertrackers 164–5
 - variance swaps 139–40, 142
 - worst of options 137
- performance comparisons
 - best of options 135–8
 - outperformance options 129–33, 167–8
 - worst of options 135, 137–8
- Philips shares 23–4, 26
- pin risk, digital options 54
- price earnings ratios 27–8
- pricing
 - see also* Greeks
 - arithmetic Asian out options 107–8
 - Asian in forwards 114–18
 - Asian in options 111
 - Asian options 105–18
 - Asian out options 105–8
 - autocallable notes 95–8, 99–104
 - barrier options 57–69, 77–8, 89–90, 95–8, 162–3
 - basket options 151
 - best of options 135
 - binomial trees 46–8
 - Black–Scholes options pricing formula 6–10, 12–14, 84–5, 107, 120–1, 131, 151
 - call options 6–8
 - callable reverse convertibles 99–102
 - chooser options 49–51
 - cliquet options 83–5
 - composite options 127–8
 - concepts 1, 4–10, 12–14, 34, 45–8, 71
 - conventional options 1, 6–10, 12–14, 34, 45–8, 71–2, 105–7
 - dividends 6–9, 10–12
 - down-and-in put barrier options 58–63, 66, 89–90, 95–8, 101–4, 162–3
 - exotic options 1, 34, 45, 49
 - forward starting options 71–5, 80–1, 83–5, 111
 - forwards 4–5, 8–10
 - geometric Asian out options 105–7
 - ladder options 78, 79–80
 - lookback options 79–81
 - methods 1, 6–8
 - Monte Carlo simulations 45–8, 57, 58, 59–60, 64, 66, 68–9, 71–3, 79–80, 107, 114–18, 135, 151
 - outperformance options 129–31
 - put options 6–8, 10
 - put–call parity 11–12
 - puttable reverse convertibles 102–4
 - quanto options 119–22
 - reverse convertibles 89–90, 95–8, 99–104, 162–3
 - skew uses 34
 - up-and-out call barrier options 64–5
 - variance swaps 140–2
 - worst of options 135
- profit-determination role, gamma 1, 16, 23–4, 27, 29, 54, 113–14, 143–4, 146–7
- put options
 - see also* Greeks
 - chooser options 49–51
 - concepts 1–8, 10, 12–21, 23–44, 158–65
 - definition 3
 - down-and-in put barrier options 57–63, 66, 87–90, 93–6, 160–3
 - forward prices 10
 - payoff profiles 4–5

- pricing 6–8, 10
- up-and-out put barrier options 58, 65, 68
- put spread option strategy, concepts 37–9
- put–call parity, concepts 11–12
- puttable reverse convertibles
 - concepts 99, 102–4
 - definition 99
- quanto options
 - composite options 125
 - concepts 119–23
 - correlation risk 119–22
 - definition 119
 - delta hedging 122–3
 - FX volatility 120–3
 - pricing 119–22
- ratchet options
 - see also* cliquet options
 - concepts 83–5
 - definition 83
 - investment bank losses 85
- ratio's vega dispersion 155
 - see also* dispersion
- realised-variance/standard-deviation
 - contrasts, variance swaps 139–40, 145–6
- rebates, barrier options 64, 77–8, 97–8, 101–2, 160–3
- recessions 27–8
- references 175
- replication, variance swaps 140–9, 169–74
- retail investors
 - barrier options 57–8
 - capital guaranteed products for risk averse investors 157–9
 - cautious equity investors 157, 160–3
 - exposure equity products for cautious investors 157, 160–3
 - financially engineered structures 157–65
 - leveraged products for risk-seeking investors 157, 163–5
 - risk averse investors 157–9
 - risk-seeking investors 157, 163–5
 - types 157–65
- reverse convertibles
 - autocallable notes 93–6, 99–104
 - barrier options 87–90, 93–6, 99–104
 - callable reverse convertibles 99–102
 - concepts 87–91, 93–6, 99–104, 160–3
 - conditional coupons 161–3
 - coupon-magnitude market conditions 89–90
 - definition 87
 - demand determinants 89–90
 - down-and-in put barrier options 87–90, 93–6, 101–4
 - examples 87–9
 - expiry scenarios 88
 - exposure equity products for cautious investors 160–3
 - hedging 90–1
 - implied volatility 90–1
 - interest rate hedge 90–1
 - pricing 89–90, 95–8, 99–104, 162–3
 - puttable reverse convertibles 99, 102–4
 - risk 88, 90–1
 - volatility hedge 90–1
- risk
 - attitudes 157–65
 - concepts 1, 35–44, 51–69, 157–65
 - exotic options 1
- risk free interest rates 6–7, 11–12, 39–40, 126, 167
 - Black–Scholes options pricing formula 6–7, 121
 - concepts 6–7, 11–12
- risk management
 - barrier options 57–69
 - strategies 35–44, 51–69
- risk neutral portfolios, concepts 6–8
- risk reversal strategies
 - see also* collars
 - concepts 39–40
- risk averse investors, capital guaranteed products 157–9
- risk-seeking investors, leveraged products 157, 163–5
- Royal Dutch/Shell shares 8–9
- settlement prices
 - see also* strike prices
 - Asian options 105–18
 - definition 105
- short players, falling markets 28
- simple option strategies 35–44

- single barrier options
 - see also* barrier options
 - concepts 57–68
 - retail demand 57–8
- skew
 - see also* implied volatility
 - barrier options 57, 58, 66–8
 - call spread option strategy 36–7, 85
 - concepts 27–34, 35, 58, 66–8, 71, 72–5, 85, 141–2, 147–8, 156
 - definition 27
 - delta effects 29–32, 147–8, 156
 - dispersion 156
 - down-and-in put barrier options 58, 66–8
 - exotic options 34, 35, 58–60
 - forward starting options 71, 72–5, 85
 - FX/equity contrasts 32–4
 - indices 156
 - lookback options 81
 - maturity dates 28–30
 - options on indices 28
 - pricing uses 34
 - term structure of the skew 28–30, 73–4
 - up-and-out call barrier options 67–8
 - up-and-out put barrier options 68
 - variance swaps 141–2, 147–8, 156
- smiles 29–34
- standard deviation
 - lognormal returns 145–6
 - realised variance 139–40, 145–6
 - variance swaps 141–2, 145–6
- standard options *see* conventional options
- stochastic processes 121
 - see also* Brownian motion
- straddles
 - chooser contrasts 51
 - concepts 40–2, 44, 51
 - definition 40–1
 - example 41–2, 44
 - payoff profiles 42, 44
- strangles
 - concepts 42–4
 - definition 42
 - payoff profiles 43–4
- strategies 35–44, 51–69, 85, 157–65
 - bearish strategies 37–9, 63, 157–9, 164–5
 - bullish strategies 35–6, 39–40, 63, 157–9, 163–5
 - call spread option strategy 35–7, 53–5, 85
 - capital guaranteed products for risk averse investors 157–9
 - cautious equity investors 157, 160–3
 - chooser options 51
 - collars 39–40
 - exposure equity products for cautious investors 157, 160–3
 - financially engineered structures 157–65
 - leveraged products for risk-seeking investors 157, 163–5
 - put spread option strategy 37–9
 - retail investors 157–65
 - risk averse investors 157–9
 - simple strategies 35–44
 - straddles 40–2, 51
 - strangles 42–4
- strike prices
 - Asian options 105–18
 - barrier shifts 60–3, 64, 67–9, 96, 104, 161–3
 - Black–Scholes options pricing formula 6–7
 - call spread option strategy 35–7
 - cliquet options 83–5
 - concepts 3–8, 12, 16–18
 - definition 3–4, 105
 - forward starting options 71–5, 80–1, 83–5, 111
 - forwards 6, 8–9, 12
 - implied volatility 27–34, 35, 36–7, 58, 66–8, 71, 72–5, 85, 141–2, 147–8, 156
 - lookback options 79–81
 - put spread option strategy 37–9
 - put–call parity 12
 - skew 27–34, 35, 58, 66–8, 71, 72–5, 85, 141–2, 147–8, 156
 - variance swaps 139–49, 155, 169–74
- structured products
 - double barrier options 68
 - down-and-in put barrier options 57
 - financially engineered structures 157–65
 - hybrids 138
 - supertrackers 164–5
 - swaps 91, 139–49, 169–74
 - see also* variance swaps

- Taylor expansion 147, 167
 term structure of the skew
 see also skew
 concepts 28–30, 73–4
 forward starting options 73–4
 theta
 arithmetic Asian out options 110
 concepts 1, 18–20, 23–4, 75, 110, 142–5
 definition 18
 expected profits 24
 gamma 23–4
 variance swaps 142–5
 time, theta 18–20, 23–4, 75, 110, 142–5
 time to maturity 6–21, 23–4, 28–30, 62, 75, 110, 121, 142–5
 Black–Scholes options pricing formula 6–7, 121
 down-and-in put barrier options 62
 skew 28–30
 theta 18–20, 23–4, 75, 110, 142–5
 TomTom shares 41–2, 43–4

 uncertainty contrasts, volatility 27–8, 32–3
 unconditional coupons, autocallable notes 94, 97–8
 unconditional double barrier options
 see also double barrier options
 concepts 69
 underlying share prices
 Black–Scholes options pricing formula 6–8, 45, 121
 delta 12–14, 21, 23–4, 25–6, 29–32, 40, 58–63, 109–12, 122–3, 126–8, 132–3, 156
 gamma 14–16, 23–4, 25–6, 27, 29, 54, 60, 75, 79–80, 110, 113–15, 142–5
 variance swaps 139–49, 155, 169–74
 up-and-in call barrier options, concepts 58, 68
 up-and-in put barrier options, concepts 58
 up-and-out call barrier options
 barrier shifts 64, 68–9
 concepts 16–17, 58, 63–5, 67–9, 78–80, 103
 definition 63
 implied volatility 16–17, 64–5
 pricing 64–5
 rebates 64, 77–8

 risk 58, 63–5, 67–8
 skew 67–8
 vega exposure 16–17, 64–5
 up-and-out put barrier options
 concepts 58, 65, 68
 definition 65
 skew 68
 USD composite options 125–8
 USD quanto options 119–23
 USD/EUR exchange rate, skew 32–4

 vanna
 see also vega
 concepts 19–21
 definition 19
 variance
 see also volatility
 composite options 126, 167–8
 outperformance options 131, 167–8
 variance notional
 concepts 140–9, 155
 vega 147–8, 155
 variance swaps
 assumptions 147–8
 concepts 139–49, 169–74
 definition 139, 146
 delta 141–5, 147–8, 156
 dispersion trading 153, 155
 event risk 146–7
 gamma 142–5, 146–7
 the Greeks 140–7
 mystery of gamma/delta 144–5
 payoff profiles 139–40, 142
 pricing 140–2
 realised-variance/standard-deviation contrasts 139–40, 145–6
 replication 140–9, 169–74
 skew delta 141–2, 147–8, 156
 standard deviation 141–2, 145–6
 theta 142–5
 vega 142–5, 147–9, 153
 vega
 see also implied volatility; vanna; vomma
 arithmetic Asian out options 110
 basket options 153–5
 concepts 1, 16–18, 37, 64–5, 74–5, 80–1, 90–1, 110, 142–5, 147–9, 153–5
 convexity principle 148–9
 definition 16

- vega (*Continued*)
 - dispersion 153–5
 - exotic options 16–17
 - forward starting options 74–5, 80–1
 - gamma comparisons 18
 - ratio's vega dispersion 155
 - up-and-out call barrier options 16–17, 64–5
 - variance notional 147–8, 155
 - variance swaps 142–5, 147–9, 153
- volatility
 - see also* implied . . . ; variance
 - Asian options 107, 115
 - barrier shifts 60–2, 64, 67–9, 96, 161–3
 - basket options 151–6
 - Black–Scholes options pricing formula 6–8, 45, 107, 121
 - concepts 6–8, 16–18, 151–6
 - dispersion 151–5
 - down-and-in put barrier options 60–2
 - falling markets 27–8
 - FX 120–3, 126–8
 - uncertainty contrasts 27–8, 32–3
 - variance swaps 139–49, 169–74
 - volatility hedge, reverse convertibles 90–1
- vomma
 - see also* vega
 - concepts 19–21
 - definition 19
- VW shares 37–8, 129–33, 135–8
- worst of options
 - concepts 135, 137–8
 - correlation risk 135, 137–8
 - definition 135
 - hybrids 138
 - payoff profiles 137
 - pricing 135
- yield curves 91
- zero coupon bonds
 - see also* financially engineered structures
 - concepts 157–9, 160–5
 - interest rates 159

Index compiled by Terry Halliday